

ROBUST ANTI-WINDUP CONTROL AND ITS APPLICATION TO
PERMANENT MAGNET SYNCHRONOUS MOTOR
SPEED REGULATION

Thesis submitted for the degree of
Doctor of Philosophy
at the University of Leicester

by

Phil March BEng (Leicester)
Control and Instrumentation Research Group
Department of Engineering
University of Leicester

March 2009

Abstract

“Robust Anti-Windup Control and its Application to Permanent Magnet Synchronous Motor Speed Regulation”
- Phil March BEng

This thesis is concerned with the design of anti-windup compensators and similar augmentations to electronic control systems that aid the system to cope with isolated nonlinearities within the closed loop such as actuator saturation. The thesis builds on theoretical contributions in the literature regarding the synthesis of low order dynamic anti-windup compensators by presenting a successful industrial application of these techniques. A range of other anti-windup techniques are described and through simulation and mathematical analysis, the pros and cons of these designs are presented. Some subtle extensions to the recent optimal synthesis routines are also presented that can offer improved flexibility in tuning and greater performance for certain systems. The industrial application chosen is an Electrically Powered Hydraulic Steering (EPHS) system in which complex constraints are applied to the currents flowing in and voltages applied to a three phase Permanent Magnet Synchronous Motor (PMSM) contained within.

Acknowledgements

First and foremost, I would like to thank my Supervisor, Dr. Matthew Turner for his ever present support and direction during the course of my research. My many questions sent over email were always met by swift and concise responses and the guidance given during our regular tutorial sessions was always pertinent. I was afforded a lot of freedom to carve out my own path of research and am grateful to have been given this. My co-supervisor Prof. Ian Postlethwaite and internal examiner Prof. Declan Bates also gave insightful direction and encouragement along the way.

I must also show great honour and appreciation for my beautiful wife who has sacrificed so much in order for me to complete this thesis. I would not have been able to complete without her support and encouragement and she deserves an award for her patience and forbearance.

Members of Conekt who assisted in the PMSM control application work also deserve special thanks: Xiao-Dong Sun, Connel Williams, Adrian Szabo and Pete Scotson all provided expert knowledge that was invaluable for understanding the system and developing successful designs. At the implementation stage, the support of Chris Dixon with TargetLink coding, and David Moule with data logging and the programming of reference sequences were also invaluable.

Finally, I would like to thank the rest of the Control and Instrumentation Research Group at Leicester. Particular mention goes to Sonny Martin, Prathyush Menon and Sajjad Fekri, who all added their own brand of vibrancy and warmth to what, prior to refurbishment, could be quite a cold and dull laboratory.

Phil March
March 2009

Contents

1	Overview	7
1.1	Thesis Structure	8
2	Preliminaries and Nomenclature	9
2.1	System Theory	9
2.1.1	State-Space (Internal) Stability	10
2.1.2	Input-Output Stability	10
2.1.3	Stability of Interconnected Systems	11
2.2	Nomenclature	15
3	PMSM Control and its Application in EPHS Systems	18
3.1	Motivation for (and Characteristics of) EPHS	18
3.2	The 3 Phase Permanent Magnet Synchronous Machine	19
3.3	3 Φ PMSM Model	20
3.4	The d - q Axis PMSM Model	23
3.5	Vector Control	24
3.6	A Linear Control Structure	27
3.7	Field Weakening Operation	28
3.8	Saturation within the Control Loop	29
3.9	The Effect of Discretisation	33
4	Introduction to Anti-Windup	34
4.1	The Problem of Actuator Saturation	34
4.2	Solving the Problem of Actuator Saturation	36
4.3	Approaches to Anti-Windup Compensation	38
4.3.1	A Generic Anti-Windup Architecture	38
4.3.2	Conditional Integration	39

<i>Contents</i>	4
4.3.3 Ad-hoc Integrator Reset Methods	40
4.3.4 Back Calculation (and Tracking)	41
4.3.5 Conventional Anti-Windup	43
4.3.6 The Observer Approach	44
4.3.7 The Observer Technique	45
4.3.8 Hanus Conditioning	46
4.3.9 Internal Model Control (IMC)	48
4.4 Optimal Dynamic and Static Anti-Windup	49
4.4.1 The Weston and Postlethwaite Framework	50
4.4.2 Performance of Anti-Windup Compensators	53
4.4.3 Full Order Anti-Windup	54
4.4.4 Static Anti-Windup	56
4.4.5 Low Order Anti-Windup	57
4.5 Summary	59
5 Anti-Windup Design for EPHS Motor Speed Control	61
5.1 Single Axis Speed Controller Anti-Windup Design	61
5.1.1 Anti-Windup Designs	63
5.1.2 Performance Comparison	75
5.2 Current Demand Constraint Implementation	79
5.3 Models for Anti-Windup Design	85
5.3.1 Linear Model Dynamics	87
5.4 Anti-Windup Designs for the Complex Model	89
5.4.1 Nonlinear Model Behaviour without Saturation	89
5.4.2 Linear Stability Analysis	90
5.4.3 Nonlinear Model Behaviour without Anti-Windup	93
5.4.4 Anti-Windup Test Conditions	94
5.4.5 Anti-Windup Designs	98
5.4.6 Simulation Analysis	100
5.4.7 Perturbed Anti-Windup Designs	107
5.5 Experimental Validation	111
5.5.1 Experimental Setup	111
5.5.2 Experimental Results - Perturbed Designs	112

5.5.3	Experimental Performance Comparisons	114
5.6	Discussion	115
5.6.1	Robustness	120
5.6.2	Model accuracy	123
5.7	Summary	132
6	Override Control and its Application to EPHS Motor Control	133
6.1	Introduction	133
6.1.1	Classic Override Control Approaches	134
6.1.2	Override Compensation	135
6.2	Override Compensator Synthesis	136
6.2.1	Static Override Compensator Synthesis	138
6.2.2	Dynamic Override Compensator Synthesis	140
6.3	Application to Current Limitation in PMSM Speed Regulation	142
6.3.1	Current limit model	143
6.3.2	Linear simulation results	145
6.3.3	Nonlinear simulation results	149
6.3.4	Comparison with anti-windup compensation	161
6.3.5	Conclusion	163
7	Improving Anti-Windup Synthesis	167
7.1	A Critique of the Existing Algorithms	168
7.2	Frequency Weighted Anti-Windup	169
7.2.1	A Modified Design Procedure	169
7.2.2	Case Studies	172
7.2.3	Remarks on Tuning the Compensator	176
7.2.4	Conclusions and Future Work	177
7.3	Reduced Sector Bounds	177
7.3.1	Reduced Sector Full-Order Anti-Windup Compensator LMI Derivation	180
7.3.2	Reduced Sector Case Studies	181
7.3.3	Discussion	190
7.3.4	Conclusions	193
7.4	Overall Conclusions	193
8	Conclusion	195

A PMSM Control System	198
A.1 Linearisation of the $d - q$ Axis PMSM Model	198
A.2 Model Parameters	199
B Anti-Windup Compensation	201
B.1 State Space Matrices for Minimal Realisation of Static and Low Order Anti-Windup Compensators	201
B.2 Static Compensator Matrices	201
B.3 Low-Order Compensator Matrices	202
B.4 State-Space Matrices for BCAT Stability Check LMI	202
C Override Compensation	203
C.1 State-Space Matrices of $G_{cl}(s)$ and $\tilde{G}_{cl}(s)$ for Static and Dynamic Override Compensation	203
C.2 Linearisation of the Phase Advance Function	203
C.3 Linear Plant Model Matrices for PMSM Override Case Study	205

Chapter 1

Overview

In engineering systems, limits on the magnitude of certain signals are almost always present. These may be mechanical limits, such as the angle of an aileron on an aircraft, limits on voltage or current within an electrical system, or limits on the range of a variable in software. When present within feedback control systems, these nonlinear functions can cause degradation of performance and may even cause instability in an otherwise stable system. In order to improve performance and preserve stability, many different approaches to the control of such systems have evolved over the years, drawing ideas from different branches of control theory.

In this thesis, the fields of anti-windup compensation and override compensation are explored. In both of these approaches, the central idea is to “retro-fit” an existing controller with an additional element that becomes active during violation of a constraint. Due to their nonlinear nature, the effect of such constraints on the stability and performance properties of a control system is not trivial to analyse. As a result, many of the early designs gave no consideration to these problems and were used simply because they seemed to work. In more recent years, rigorous tools to assess the stability and performance properties of such nonlinear systems, and, moreover, to synthesise compensators that guarantee stability and optimise performance have become available. In this thesis, a selection of these modern methods are compared to some of the better known ad-hoc approaches and their adoption in industry is furthered by application to an industrial case study.

The industrial application of interest is the motor control system within an Electrically Powered Hydraulic Steering (EPHS) system for use in the automotive sector. As the system is cost sensitive there is a strong incentive to achieve maximum utilisation of the available voltage and current and therefore the system is heavily dependent upon anti-windup compensators to preserve performance during various saturation scenarios. The plant and controller of the EPHS system are both nonlinear and the constraints applied are both multivariable and time varying, thus the application of linear anti-windup compensation is challenging. However, a novel method of posing the current constraint as a time varying single-input-single-output constraint allows successful and appealing designs to be generated that are validated both in simulation and in practice.

1.1 Thesis Structure

Chapter 2 introduces elements of the system theory required for the analysis and design approaches discussed in the thesis and defines the various notational conventions used.

Chapter 3 starts the main body of the thesis and describes the EPHS system, the PMSM motor and associated models and strategies in detail. The objective is to give a control-oriented description of the system which will be used in most of the subsequent chapters and also to describe the way in which the nonlinearities present affect the system's performance and the associated controller design.

Chapter 4 introduces the concept of anti-windup and describes various strategies which can be found in the literature. In particular, certain "optimal" strategies, on which most of the anti-windup designs discussed later are based, are described in some detail.

Chapter 5 describes the application of various anti-windup strategies to the EPHS system. In this chapter a simple, largely linear, model is used first to select promising candidate designs, and then a more complex nonlinear model is used for the design, simulation and comparison of the selected compensation techniques. Experimental results showing the success of these designs are also given.

Chapter 6 shows how some of the saturation constraints present in the the PMSM system may be more naturally addressed in the alternative framework of override control. The concept of override control is briefly described and the design of several override controllers for the PMSM system is described. The results are analysed and compared using a complex sampled data model of the motor and practical constraints which affect the performance of the override controllers are highlighted.

Chapter 7 is a supplemental chapter on some modifications and extensions which may be made to the modern optimisation-based compensators used in most of the foregoing chapters. Some of the possible deficiencies of these compensators are identified and modified design techniques are presented. These modified techniques are compared to the standard ones using some simple academic examples taken from the literature.

Chapter 8 concludes the thesis and states the main contributions and identifies directions for future work.

Chapter 2

Preliminaries and Nomenclature

This thesis can be viewed primarily as the application and development of modern anti-windup techniques to PMSM speed control systems. For the PMSM application all the necessary details are introduced together in Chapter 3. However, in order to appreciate the anti-windup techniques discussed in other chapters it is necessary first to have an understanding of some concepts from system theory. While it is not possible to describe exhaustively all technical background, a brief overview of the required concepts is included here as the first section in this introductory chapter. In the subsequent nomenclature section, a collection of naming and other notational conventions used throughout the thesis are presented for reference purposes.

2.1 System Theory

We shall be concerned with two systems. The first is the general nonlinear state-space system, Σ_{NL} , in which $x \in \mathbb{R}^n$ is the state vector, $u \in \mathbb{R}^m$ is the input vector and $y \in \mathbb{R}^{n_y}$ is the output vector:

$$\Sigma_{NL} = \begin{cases} x &= f(x, u) \\ y &= h(x, u) \end{cases} \quad (2.1)$$

When the state equations are linear, we have the linear state-space system

$$\Sigma_L = \begin{cases} x &= Ax + Bu \\ y &= Cx + Du \end{cases} \quad (2.2)$$

This system has an associated transfer function

$$G(s) = C(sI - A)^{-1}B + D \quad (2.3)$$

and sometimes the following shorthand notation is used to denote a state-space realisation:

$$G(s) \sim \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array} \right] \quad G(s) \sim (A, B, C, D) \quad (2.4)$$

Typically we shall be concerned with two types of stability in the thesis: state-space stability and input-output stability.

2.1.1 State-Space (Internal) Stability

The system Σ_{NL} is described as globally *asymptotically stable* if, assuming $u \equiv 0$ we have

$$\lim_{t \rightarrow \infty} x(t) = 0 \quad \forall x \in \mathbb{R}^n \quad (2.5)$$

Note that, without loss of generality, we have assumed that the origin is the equilibrium point; a simple change of co-ordinates in the state-space can be used to state asymptotic stability for another point. A stronger form of stability is *exponential stability* which insists on a certain decay rate of the initial state, viz.

$$x(t) = \kappa e^{-\eta t} \|x(0)\| \quad \kappa, \eta > 0, \quad \forall x \in \mathbb{R}^n \quad (2.6)$$

Local stability, a concept not used much in this thesis, is defined similarly except the region of attraction is taken to be a subset, $\mathcal{X} \subset \mathbb{R}^n$, of the state-space. In the special case of linear systems, a locally asymptotically stable system is necessarily globally asymptotically stable and in this case asymptotic (actually exponential) stability is simply checked using the eigenvalues of the matrix A .

A concept strongly related to asymptotic stability of nonlinear systems is the notion of *Lyapunov Stability*. Σ_{NL} is said to be *globally stable in the sense of Lyapunov* if there exists an “energy function” $V(x)$ such that the following properties hold

$$V(x) > 0 \quad \forall x \in \mathbb{R}^n \quad (2.7)$$

$$\dot{V}(x) \leq 0 \quad \forall x \in \mathbb{R}^n \quad (2.8)$$

In this case, it follows that all states will converge to a set $\mathcal{B} := \{x \in \mathbb{R}^n : \|x\| < \beta\}$. Furthermore, if the second inequality is strengthened to $\dot{V}(x) < 0$ it follows that the system will be globally asymptotically stable. Proving stability of a nonlinear system by searching for a Lyapunov function $V(x)$ is normally called “the second method of Lyapunov” and allows *global* stability conditions to be proved. Note that this method is a *sufficient* condition for stability of Σ_{NL} ; it is possible that Σ_{NL} may be stable even if inequalities (2.7) and (2.8) are not satisfied. Inevitably this means that proving stability using Lyapunov’s second method is conservative.

2.1.2 Input-Output Stability

Let us again consider Σ_{NL} , but this time assume that $u \neq 0$. An alternative way of assessing stability of a system is by so-called bounded-input-bounded-output (BIBO) stability: that is given a bounded input, does there exist a bound on the output and, if so, what is that bound? There are several types of BIBO stability, but perhaps the most useful are based on the use of the \mathcal{L}_p norms, which measure the “size” of a signal, viz.

$$\|x\|_p = \left(\int_0^\infty \|x(t)\|^p dt \right)^{\frac{1}{p}} \quad (2.9)$$

where $p \in [1, \infty]$. A nonlinear system, Σ_{NL} is said to have an \mathcal{L}_p gain or \mathcal{L}_p norm, if there exist constants γ and β such that

$$\|y\|_p < \gamma \|u\|_p + \beta. \quad (2.10)$$

The smallest such γ is said to be the \mathcal{L}_p gain or induced \mathcal{L}_p norm of the system. A particularly useful case is when $p = 2$ and we consider the \mathcal{L}_2 gain which roughly corresponds to the RMS energy gain of the system. For linear systems it is noteworthy that the \mathcal{L}_2 gain is actually equivalent to the \mathcal{H}_∞ norm of the system; that is

$$\sup_{0 \neq u \in \mathcal{L}_2} \frac{\|\Sigma_L u\|_2}{\|u\|_2} = \sup_{\omega} \bar{\sigma}[\Sigma_L(j\omega)] \quad (2.11)$$

Note that a system may be BIBO stable without being asymptotically stable. For example an oscillator is certainly \mathcal{L}_∞ stable but is not asymptotically stable.

2.1.3 Stability of Interconnected Systems

To assess the stability of two or more systems that are connected in closed loops some additional tools are available. These are introduced by considering first how the interconnection of linear systems can be tackled, and then extending to the case where one of the systems in the loop is nonlinear.

The Small Gain Theorem

The small gain theorem [81, 45] is a sufficient condition to guarantee BIBO stability of a closed loop system and allows stability of a closed loop system to be assessed without full knowledge of the frequency domain behaviour of its constituent parts. Instead, only information regarding the gain of the subsystems is used and, regardless of phase, stability of the closed loop system is guaranteed if certain conditions relating to the gain of the subsystems are satisfied.




Figure 2.1: Feedback interconnection of 2 possibly nonlinear systems

In its most general form, the small gain theorem ensures finite gain \mathcal{L}_p stability of the feedback interconnection of a number of nonlinear systems. Consider the case depicted in Figure 2.1 for which two such subsystems are H_1 and H_2 . If H_1 and H_2 are both finite gain \mathcal{L}_p stable with \mathcal{L}_p gains γ_1 and γ_2 , a sufficient condition for finite gain \mathcal{L}_p stability of the feedback loop from inputs u_1 and u_2 to outputs y_1 and y_2 is given by (2.12). Typically for nonlinear systems, finite gain \mathcal{L}_2 stability is sought where γ_1 and γ_2 represent the \mathcal{L}_2 gain of the subsystems.

$$\gamma_1 \gamma_2 < 1 \quad (2.12)$$

For linear systems, the \mathcal{H}^∞ norm is used which is equivalent to the \mathcal{L}_2 gain. From its definition above, stability is guaranteed provided that $\|H_1\|_\infty \|H_2\|_\infty < 1$. However, with linear systems we really only require that $\|L(s)\|_\infty < 1$ where $L(s) = H_1(s)H_2(s)$ and this proves to be less conservative as $\|H_1\|_\infty \|H_2\|_\infty \geq \|L(s)\|_\infty$. This form of the small gain theorem applies only to linear systems and is still conservative because from a gain and phase margin perspective, this ensures that the phase margin is infinite and the gain margin is at least $1/\|L(s)\|_\infty$. Therefore, when full knowledge of the system dynamics are known it is preferable to use standard linear tools such as the theorems of Bode and Nyquist. The obvious exception being for multivariable systems for which phase is not well defined.

The small gain theorem finds its main use in linear robust control where uncertainty in the system is modelled as a norm bounded nonlinearity. Consider the control system model of Figure 2.2 where K represents a linear controller and G is a linear model of the plant. Error between the plant model and real system is represented by an additive uncertainty model consisting of a frequency weighting function, W , and a normalised uncertain element, Δ , with norm equal to 1. Ignoring exogenous inputs and outputs, the uncertain closed loop system can be represented by the interconnection of a single linear transfer function, $M(s)$, and the uncertain element, Δ , as shown in Figure 2.3. Applying the small gain theorem, provided that $\|M(s)\|_\infty < 1$, the uncertain closed loop system is stable and the nominal control system has been shown to be robustly stable to the given uncertainty model.




Figure 2.2: Closed loop system with additive uncertainty model




Figure 2.3: M-delta feedback interconnection

The approach used for uncertain linear systems can also be applied to linear systems with isolated nonlinearities i.e. a Lure type nonlinear system, by modelling the nonlinearities as a norm bounded uncertainty. This can

be represented by the closed loop interconnection of Figure 2.4 where ψ represents the nonlinear elements. However, this approach tends to be very conservative and so alternatives are generally used.




Figure 2.4: Lure type nonlinear system

The Circle Criterion

The circle criterion is an application of absolute stability theory to a Lure type nonlinear system. This is a system for which the forward path transfer function is linear time invariant and any nonlinear characteristics of the system are modelled as an isolated nonlinearity in the feedback path as shown in Figure 2.4. The nonlinearity, ψ , is required to be memory-less and static but may be time-varying. For the circle criterion we require that the nonlinearity satisfy a sector condition, described by (2.13) and depicted in Figure 2.5,

$$[\psi(t, y) - \alpha y][\psi(t, y) - \beta y] \leq 0 \quad (2.13)$$

such that the nonlinearity lies within lines of gradient α and β intersecting the origin. Here, we only consider




Figure 2.5: Global sector nonlinearity

the case where the sector bound holds globally i.e. the nonlinearity lies within lines of gradient α and β for $y \in [-\infty, \infty]$. For all nonlinearities satisfying such a sector condition, saturation and deadzone functions being two examples, the origin is an equilibrium point. If the origin can be shown to be globally asymptotically stable for any nonlinearity contained within the sector, the nonlinear system is said to be globally absolutely stable [45]. Proof of absolute stability for an appropriate sector bound is an appealing method of checking nonlinear

stability for a Lure type nonlinear system since an explicit description of the nonlinearity is not required and a potentially complex nonlinearity can be handled therefore with relatively simple mathematics. For the sector bounds to hold globally with the nonlinearities considered in this thesis, we require that $\alpha = 0$ and thus, the sector condition is simplified to that of (2.14).

$$\psi(t, y)'[\psi(t, y) - \beta y] \leq 0 \quad (2.14)$$

Theorem 2.1.1 *Provided that A is Hurwitz, (A, B) is controllable, (A, C) is observable and $\psi(\cdot)$ satisfies the sector condition globally, the nonlinear system is globally absolutely stable if $Z(s) = I + \beta C(sI - A)^{-1}B$ is strictly positive real [45].* \diamond

The circle criterion can also be addressed using Lyapunov theory as follows where a quadratic Lyapunov function is chosen, viz.

$$V(x) = x'Px \quad (2.15)$$

$$\dot{V}(x) = \dot{x}'Px + x'P\dot{x} = x'(PA + A'P)x - 2x'PB\psi \quad (2.16)$$

The nonlinearity is bound by the sector $[0, I]$ which is defined as (2.17) where $W > 0$ is a diagonal matrix. This function can be added to the right hand side of (2.16) to produce the inequality of (2.18). Enforcing negative definiteness of the right hand side of (2.18) will guarantee $\dot{V}(x) < 0$ and if the free variable $W > 0$ is chosen to be small, the conservatism introduced by inclusion of the sector bound will be limited. Thus, satisfaction of the circle criterion is equivalent to the following condition.

$$2\psi'W(y - \psi) \geq 0 \quad (2.17)$$

$$\dot{V}(x) \leq x'(PA + A'P)x - 2x'PB\psi - 2\psi'W(y - \psi) \quad (2.18)$$

Theorem 2.1.2 *The nonlinear closed loop system formed by connection of the observable and controllable linear system with state space matrices (A, B, C) and the static nonlinearity, ψ , bound by the sector $[0, I]$ is globally absolutely stable if there exist a matrix $P > 0$ and diagonal matrix $W > 0$ which satisfy $x'(PA + A'P)x - 2x'PB\psi - 2\psi'W(y - \psi) < 0$.* \diamond

This is a sufficient but not necessary condition for stability and so failure to find a feasible solution does not imply that the nonlinear system is unstable. In some cases for a stability guarantee to be found, an alternative Lyapunov function may be required. However, the quadratic function shown here is well proven and widely used as the form of the resulting stability condition can be manipulated simply into a Linear Matrix Inequality (LMI) framework or an Algebraic Riccati Equation that can be solved reliably. Extensions to the case where $\alpha \neq 0$ and/or systems for which A is not Hurwitz are possible by the application of a *loop transformation*, resulting in what is referred to as the multivariable circle criterion [45].

In the scalar case, the circle criterion has a useful graphical interpretation such that absolute stability can be checked simply by inspecting the Nyquist contour of the linear subsystem.

Theorem 2.1.3 *The scalar nonlinear closed loop system formed by connection of the observable and controllable linear system with state space matrices (A, B, C) and the static nonlinearity, ψ , bound by the sector (α, β) is globally absolutely stable if one of the following conditions is satisfied.*

- *If $0 < \alpha < \beta$, the Nyquist contour of $C(j\omega I - A)^{-1}B$ does not enter a disc centred about the real axis with intercepts at $(-1/\alpha, 0)$ and $(-1/\beta, 0)$ and encircles it m times in the counter-clockwise direction, where m is the number of eigenvalues of A with positive real parts.*
- *If $0 = \alpha < \beta$, A is Hurwitz and the Nyquist contour of $C(j\omega I - A)^{-1}B$ lies to the right of the vertical line $\Re(s) = -1/\beta$.*
- *If $\alpha < 0 < \beta$, A is Hurwitz and the Nyquist plot of $C(j\omega I - A)^{-1}B$ lies within a circle centred about the real axis with intercepts $(-1/\beta, 0)$ and $(-1/\alpha, 0)$.*

Benefits of the circle criterion over the small gain theorem applied to such nonlinear systems are that the nonlinearity is more tightly defined and should therefore reduce conservatism, and also that the stronger condition of asymptotic stability is sought rather than BIBO stability. However, as conditions for absolute stability go, the circle criterion is one of the most conservative. The reason why it is so widespread is its amenability to solution within a LMI or ARE framework, allowing it to be included within performance and or robustness optimisations.

The Popov Criterion

The Popov Criterion is another application of absolute stability theory to a Lure type nonlinear system and makes use of a Lyapunov function of the form

$$V(x) = x'Px + \eta \int_0^y \psi'(\sigma)\beta d\sigma \quad (2.19)$$

The inclusion of an integral term adds complexity to its use but the results achieved are generally less conservative than with the circle criterion. As with the circle criterion, extensions can be made to multivariable nonlinearities and sectors that do not include the horizontal axis. However, due to its complexity, it generally leads to intractable optimisations if included in synthesis problems. For this reason it is usually only employed to check the stability of a pre-designed nonlinear system, and is not considered further in the context of this thesis.

2.2 Nomenclature

The Euclidean norm of a vector $x \in \mathbb{R}^n$ is defined as

$$\|x\| = \sqrt{\sum_{i=1}^n x_i^2}$$

The state-space model of a single degree of freedom controller is defined using the following state space matrices.

$$K(s) \sim \left[\begin{array}{c|c} A_c & B_c \\ \hline C_c & D_c \end{array} \right]$$

For a two degree of freedom controller, the convention adopted is given below where matrices B_{cr} and D_{cr} are driven by the reference signal and B_c and D_c are driven by measured plant outputs.

$$K(s) \sim \left[\begin{array}{c|cc} A_c & B_{cr} & B_c \\ \hline C_c & D_{cr} & D_c \end{array} \right]$$

The state-space model of a linear plant is given below where the subscript p denotes ‘plant’, matrices B_p and D_p correspond to the control inputs and matrices B_{pd} and D_{pd} correspond to disturbance inputs.

$$G(s) \sim \left[\begin{array}{c|cc} A_p & B_{pd} & B_p \\ \hline C_p & D_{pd} & D_p \end{array} \right]$$

A collection of commonly used notational symbols and abbreviations are given in Tables 2.1 and 2.2.

Notation	Description
$A > 0$	Matrix A is positive definite
$A \geq 0$	Matrix A is positive semi-definite
$A < 0$	Matrix A is negative definite
$A \leq 0$	Matrix A is negative semi-definite
A'	The transpose of matrix A
A^{-1}	The inverse of matrix A
$\text{diag}(x_1, \dots, x_n)$	Diagonal matrix $I \times [x_1 \dots x_n]'$
I	The Identity matrix
$\Im(\cdot)$	The imaginary component of a complex number
j	The complex number $\sqrt{-1}$
\mathbb{N}	The set of natural numbers, $0, 1, 2, \dots$
\mathbb{Q}	The set of rational numbers, $\mathbb{Q} = \frac{m}{n}$, $m \in \mathbb{Z}$, $n \in \mathbb{Z}$, $n \neq 0$
\mathbb{R}	The set of real numbers, $-\infty, \infty$
$\mathbb{R}^{m \times n}$	The set of real m by n matrices
$\Re(\cdot)$	The real component of a complex number
s	Laplace operator
$\text{sign}(x)$	The signum function, $x/ x $
$1e-6$	1×10^{-6}
\mathbb{Z}	The set of all integers, $\dots, -3, -2, -1, 0, 1, 2, 3, \dots$

Table 2.1: Mathematical notation

Notation	Description
AW	Anti-Windup
BCAT	Back Calculation and Tracking
BIBO	Bounded input bounded output
CAW	Conventional (high gain) anti-windup
Dz(.)	Deadzone operator
e	Tracking error signal, $r - y$
IMC	Internal Model Control
L(s)	Loop transfer function, typically GK
MIMO	Multiple input multiple output
m	Number of control inputs to the plant
n_c	Number of controller states
n_d	Number of plant disturbance inputs
n_p	Number of plant states
n_r	Number of reference inputs
n_y	Number of plant outputs
OR	Override
PI	Proportional plus integral control
q	Number of constrained plant outputs (for OR control)
u	Control signal
\bar{u}	Saturation limit on control signal
\tilde{u}	Violation of control saturation limit, $u - \text{sat}(u)$
r	Reference demand
$S(s)$	Sensitivity function, typically $(I + GK)^{-1}$
sat(.)	Saturation operator
SISO	Single input single output
$T(s)$	Complementary sensitivity function, typically $GK(I + GK)^{-1}$
x	Generic state vector
x_c	Controller state vector
x_p	Plant model state vector
y	Output signal
z^{-1}	Delay operator

Table 2.2: Commonly used symbols and abbreviations

Chapter 3

PMSM Control and its Application in EPHS Systems

In this chapter the concept of Electrically Powered Hydraulic Steering (EPHS) is introduced and the requirements of the electric drive systems that power them are given. Motivation for the choice of Permanent Magnet Synchronous Motors (PMSM) for the EPHS application is given, together with an introduction to a typical speed control strategy and the ideas behind it. In addition, practical constraints associated with their application in an automotive system are given, highlighting the requirement for anti-windup conditioning and other such nonlinear modifications which are addressed in the remainder of the thesis.

3.1 Motivation for (and Characteristics of) EPHS

A conventional hydraulic power assisted steering (HPAS) system reduces the effort required to steer the hand wheel of a vehicle by applying hydraulic pressure in the appropriate direction at the steering rack. Hydraulic pressure is provided by a pump, driven by rotation of the engine. A complex rotary valve assembly mounted on the steering column is used to detect the torque applied by the driver and divert fluid flow to the rack according to the direction and magnitude of the applied torque. Two pipes are used to direct fluid flow to either side of the pinion, each with a return path to the pump. When no torque is applied by the driver, equal hydraulic fluid flow is provided to each side of the pinion, resulting in no assistance. When a torque is applied at the hand wheel by the driver, a torsion bar in the steering column flexes, causing the valve to rotate. This rotation reduces the fluid flow to one side of the pinion and increases it to the other, generating a pressure differential, assisting the driver to counteract the rack loads. The assistance provided is tuned by careful design of the valve, with maximum assistance provided when all fluid flow is diverted to one output channel.

The flow of fluid, and hence the hydraulic pressure, upstream of the valve is dependent upon the engine speed and so to limit the driver efforts during parking, the pump is designed to provide adequate fluid flow while the engine is idling. During driving the engine speed is higher and so the hydraulic pressure is also higher, generating more hydraulic pressure than required. The main problem with conventional hydraulic power steering is that the hydraulic pump draws power from the engine continuously, even when no assistance is required such as when the vehicle is stationary. In addition, as the engine speed increases the power drawn from the engine

increases further leading to further reductions in efficiency. There are two factors here which conflict with the requirements of the driver. Firstly, greatest assistance is required at parking where the road to tyre friction is highest but in this condition, the engine speed will be low, limiting the assistance available. Secondly, during motorway cruising the engine speed may be relatively high but due to the dynamics of the vehicle and the small steering angles applied, only a small amount of assistance is required. In this case, the power drawn from the engine is almost entirely wasted. In addition to highlighting the inefficiencies of HPAS, this also highlights a fundamental trade-off in the tuning of hydraulic power steering as low assistance is desired at high speeds to provide good road feedback but high assistance is required at low speeds to limit driver effort levels.

With Electrically Powered Hydraulic Steering (EPHS) systems, the hydraulic pump is powered via an electric motor rather than the engine. As such, the hydraulic pressure can be accurately controlled, providing high pressure when greatest assistance is required, and reducing pressure to save energy when less assistance is required. This can significantly improve the overall efficiency of the vehicle and also ensure that maximum assistance is available during parking even when the engine is idling. In addition to scheduling the hydraulic pressure with vehicle speed, further efficiency gains can be made by dynamically varying the hydraulic pressure according to the demands made by the driver. As an example, when driving in a straight line at 40 miles per hour (mph) the assistance torque requirements are low, allowing the motor speed to be reduced. When high torque is applied to the hand wheel in an evasive manoeuvre, the motor speed can be quickly increased to meet the assistance demands. To achieve this type of operation, the motor needs to be able to run at low speed but accelerate quickly when the assistance demand increases to provide the hydraulic pressure required to meet the demand. Sharp increases in the required assistance can be caused either by aggressive inputs at the hand wheel, or by disturbances from the road such as caused by driving over a pot-hole. In either case, the motor control system must be able to react very quickly to step inputs and also be robust to load disturbances when running at low speeds, preventing disturbances from stalling the motor. For such an application, a motor with high torque capability and low inertia is required and a control system tuned to provide fast dynamics. This is an area where permanent magnet synchronous motors (PMSMs) are ideally suited.

3.2 The 3 Phase Permanent Magnet Synchronous Machine

The three phase permanent magnet synchronous machine (3ϕ PMSM) is an AC motor which is becoming increasingly prominent in high torque applications, and is steadily replacing DC motors in many applications due to its high power density and torque to inertia ratio [8]. The 3ϕ motor is constructed rather like an inside-out DC motor with permanent magnet pole pairs on the rotor, and three phase windings around the stator as shown in Figure 3.1. Note that for smoother torque production as the machine rotates and higher torque capability, it is common to have multiple (P) permanent magnet pole pairs and corresponding stator phase windings. In the example shown in Figure 3.1 there are three pole pairs (P=3). It is possible to drive such a motor in a number of ways but for this work we consider only sinusoidal commutation.




Figure 3.1: Schematic of PMSM with 3 permanent magnet pole pairs

Windings a', b', c' and a'', b'', c'' are wired in parallel with windings a, b, c and exhibit the same orientation to the permanent magnet poles, albeit to different pole pairs. As such, the torque generated by the interaction of the magnetic field produced by the current in windings a, b, c and the adjacent permanent magnet poles is also mirrored by the other pole pair interactions. These torque contributions sum together, meaning that the power rating of the machine is proportional to the number of pole pairs: In effect there are P machines working in parallel. A benefit of this is that for analysis purposes we need only consider one pole pair and corresponding set of stator windings and simply multiply the torque produced by P . Other factors to take into account are the rotor position and velocity information as for each mechanical revolution, the stator sees P electrically equivalent orientations. To accommodate this, a variable called the rotor electrical position, θ_e , is defined as P times the rotor mechanical position, with the rotor electrical velocity ω_e being its time derivative. The resulting model in electrical coordinates is represented by the schematic of Figure 3.2 and will be used in the analysis to follow.

3.3 3Φ PMSM Model

The PMSM is modelled in 3ϕ coordinates, a, b, c , as a fifth order nonlinear system according to the following system of equations [67] and [97], a key to which is found in Table 3.1.




Figure 3.2: PMSM electrical model schematic and associated reference frames

$$\frac{di_a}{dt} = \frac{1}{L_s} [V_a - R_s i_a + \psi_f \omega_e \sin(\theta_e)] \quad (3.1)$$

$$\frac{di_b}{dt} = \frac{1}{L_s} \left[V_b - R_s i_b + \psi_f \omega_e \sin\left(\theta_e - \frac{2\pi}{3}\right) \right] \quad (3.2)$$

$$\frac{di_c}{dt} = \frac{1}{L_s} \left[V_c - R_s i_c + \psi_f \omega_e \sin\left(\theta_e - \frac{4\pi}{3}\right) \right] \quad (3.3)$$

$$\frac{d\theta_e}{dt} = P\omega_m = \omega_e \quad (3.4)$$

$$\frac{d\omega_m}{dt} = \frac{1}{J} [T_e - \text{load} - B\omega_m] \quad (3.5)$$

$$T_e = \frac{-K_e}{\sqrt{3}} \left[i_a \sin \theta_e + i_b \sin\left(\theta_e - \frac{2\pi}{3}\right) + i_c \sin\left(\theta_e - \frac{4\pi}{3}\right) \right] \quad (3.6)$$

$$\psi_f = \frac{K_e}{P\sqrt{3}} \quad (3.7)$$

V_a , V_b and V_c represent the instantaneous voltages applied to each winding on the stator. Applying sinusoids with a phase difference of $\frac{2}{3}\pi$ radians between them generates a voltage phasor which rotates around the rotor shaft. The current phasor resulting from the applied voltage waveforms induces an electric field that interacts with the magnetic field of the permanent magnets to produce torque.

Three of the five states in the model relate to the stator phase currents i.e. the currents flowing in the three windings, i_a , i_b and i_c . Since these currents can be resolved into a two dimensional phasor, there is a superfluous state in the model. It is possible therefore to produce a 4th order model in a stationary cartesian coordinate system (α, β) which also fully describes the dynamics of the motor. This coordinate system and its relationship to the physical displacement of the windings a , b , c is depicted in Figure 3.2. The α, β axes then represent the real and complex parts of the sinusoidal stator current phasor. Further simplification can be achieved by

Variable/coefficient	Description	Value (normalised units)
i_a, i_b, i_c	Stator phase currents	
V_a, V_b, V_c	Stator phase voltages	
θ_e	Rotor electrical position	
θ_m	Rotor mechanical position	
ω_m	Rotor speed (mechanical)	
ω_e	Rotor speed (electrical)	
ψ_f	Flux linkage	
$load$	Rotor load torque	
T_e	Electromagnetic torque	
P	Number of magnetic pole pairs	5
K_e	Motor torque constant	1
B	Rotor drag coefficient	0.0323
R_s	Stator winding resistance	0.2541
L_s	Stator winding inductance	0.1277
J	Rotor moment of inertia	1.4837

Table 3.1: Key for 3ϕ PMSM model equations

using a rotating coordinate system which tracks the position of the rotor. This cartesian coordinate system defines a ‘direct’ axis (d -axis) in the direction of the rotor magnetic flux and a ‘quadrature’ axis (q -axis) in quadrature to the rotor magnetic flux. In rotor electrical coordinates the d and q axes are perpendicular, but in physical coordinates for a motor with P pole pairs, the mechanical angle between the axes will be $\frac{\pi}{2P}$ radians. Modelling in this rotational co-ordinate system eliminates sinusoidal terms within the model equations, allowing the magnitude and phase of the stator electromagnetic flux to be varied easily with respect to the rotor magnetic flux. This model is still strictly 4th order although for controller design purposes, the rotor position state can usually be ignored. A full description of the d - q axis model is given in Section 3.4.

Conversion between 3ϕ and d - q axis reference frames for the purpose of control is achieved using a rotor position measurement, θ_e , and the nonlinear coordinate transformations of (3.8) and (3.9). Historically, Clarke transformations are used to translate between 3ϕ and α, β reference frames, followed or preceded by Park transformations which are used to translate between α, β and d - q axes [30]. However, as the stationary cartesian reference frame is usually not required, it is more efficient to make the transformation in one operation as shown here.

The signals V_0 and i_0 are referred to as the zero components and have been included here for mathematical precision. However, for a balanced system i.e. one in which the sum of the three phase quantities are always zero, these zero components are always equal to zero and hence can be ignored. The systems considered in this thesis are all balanced.

The form of transformations used in this thesis are referred to as magnitude invariant transformations as the magnitude of voltages and currents are equal in both reference frames. However, power and torque are not equivalent between reference frames and to compensate for this, correction factors are required in the equations defining torque and power. Power invariant transformations such as detailed in [66, 67] also exist but with these, the magnitude of voltages and currents are not equivalent between reference frames.

$$\begin{bmatrix} i_d \\ i_q \\ i_0 \end{bmatrix} = \frac{2}{3} \begin{bmatrix} \cos(\theta_e) & \cos(\theta_e - 2\pi/3) & \cos(\theta_e + 2\pi/3) \\ \sin(\theta_e) & \sin(\theta_e - 2\pi/3) & \sin(\theta_e + 2\pi/3) \\ 0.5 & 0.5 & 0.5 \end{bmatrix} \begin{bmatrix} i_a \\ i_b \\ i_c \end{bmatrix} \quad (3.8)$$

$$\begin{bmatrix} V_a \\ V_b \\ V_c \end{bmatrix} = \begin{bmatrix} \cos(\theta_e) & \sin(\theta_e) & 1 \\ \cos(\theta_e - 2\pi/3) & \sin(\theta_e - 2\pi/3) & 1 \\ \cos(\theta_e + 2\pi/3) & \sin(\theta_e + 2\pi/3) & 1 \end{bmatrix} \begin{bmatrix} V_d \\ V_q \\ V_0 \end{bmatrix} \quad (3.9)$$

3.4 The d - q Axis PMSM Model

The PMSM model in d - q coordinates is given by the following equations:

$$\begin{aligned} \frac{di_d}{dt} &= \frac{1}{L_s} [V_d - R_s i_d + PL_s \omega_m i_q] \\ \frac{di_q}{dt} &= \frac{1}{L_s} \left[V_q - R_s i_q - PL_s \omega_m i_d - \frac{K_e \omega_m}{\sqrt{3}} \right] \\ \frac{d\omega_m}{dt} &= \frac{1}{J} [T_e - load - B\omega_m] \\ T_e &= \frac{\sqrt{3}}{2} K_e i_q \\ \omega_e &= P\omega_m \end{aligned} \quad (3.10)$$

where the new variables are defined as follows:

Variable/coefficient	Description
i_d	The stator current vector component in the d -axis direction
i_q	The stator current vector component in the q -axis direction
V_d	The stator voltage vector component in the d -axis direction
V_q	The stator voltage vector component in the q -axis direction

Table 3.2: Key for d - q axis PMSM model equations

Note that the only non-linearities in this model are the products $\omega_m i_q$ and $\omega_m i_d$ ¹. This means that the model can be linearised easily around a set of equilibrium states $[i_{d,lin} \ i_{q,lin} \ \omega_{m,lin}]^T$ to produce the linear state space

¹This model can be considered therefore as a linear parameter varying (LPV) system where the varying parameter is the motor speed.

model of (3.11). This linear model is ideal for use in model based controller designs using techniques such as \mathcal{H}_∞ mixed sensitivity and \mathcal{H}_∞ loop shaping [103, 81] and linear analysis, and will be put to use in later chapters.

$$\begin{aligned} \begin{bmatrix} \dot{i}_d \\ \dot{i}_q \\ \dot{\omega}_m \end{bmatrix} &= \begin{bmatrix} -\frac{R_s}{L_s} & P\omega_{m,lin} & Pi_{q,lin} \\ -P\omega_{m,lin} & -\frac{R_s}{L_s} & -Pi_{d,lin} - \frac{P\psi_f}{L_s} \\ 0 & \frac{3P\psi_f}{2J} & -\frac{B}{J} \end{bmatrix} \begin{bmatrix} i_d \\ i_q \\ \omega_m \end{bmatrix} + \begin{bmatrix} \frac{1}{L_s} & 0 \\ 0 & \frac{1}{L_s} \\ 0 & 0 \end{bmatrix} \begin{bmatrix} V_d \\ V_q \end{bmatrix} \\ \begin{bmatrix} \omega_m \\ i_d \\ i_q \end{bmatrix} &= \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} i_d \\ i_q \\ \omega_m \end{bmatrix} \end{aligned} \quad (3.11)$$

3.5 Vector Control

In traditional AC machine control, motor speed and torque is controlled in an open loop manner by varying the magnitude and frequency of the three phase voltage signals. This method is suitable for steady-state operation and systems with slow dynamics but does not exploit the high performance capability of these machines. To achieve maximum performance, a closed loop control strategy called vector control is applied. Vector control or ‘field oriented control’ is a control strategy for brushless machines that allows torque to be controlled precisely by manipulating both the magnitude of the stator magnetomotive force (MMF) and its phase relationship with that of the rotor permanent magnet flux. To achieve this, the controller needs to manipulate the stator current components in the rotor d - q axes which produce the stator MMF.

In vector control, measurements of the winding currents are taken and translated into the d - q reference frame using an appropriate measurement of the motor electrical position, θ_e . These d and q axis current samples are compared to the demanded values, generating a current error vector that is used to drive the controller. The control action required to drive the error toward zero is computed, producing a voltage output vector also in the d - q reference frame, $V_{dq} = \begin{bmatrix} V_d & V_q \end{bmatrix}'$. The d - q axis voltage demands are translated into three phase AC waveforms using a measurement of the rotor position and the inverse transformation of (3.9), and applied commonly via a Pulse Width Modulation (PWM) inverter. Since digital control hardware development has enabled the control functions and nonlinear coordinate transformations to be computed at sufficiently fast sample rates, this has become an established method for the control of brushless permanent magnet machines. Material regarding its use can be found in [30], [48], [97] and [11].

The instantaneous control of motor torque allows the speed, acceleration and position of the rotor to be controlled precisely, making it a very capable actuator. A further benefit of vector control is that for steady-state operation, each element of the voltage and current vector is constant, whereas in three phase and α, β coordinates the voltages and currents are sinusoids. This property enables the use of a vast array of linear control solutions and analysis techniques. To understand further the benefit of applying control in the d - q reference frame, it is useful to make a comparison between the PMSM and a conventional DC motor.

In a brushed DC motor the torque produced is linearly proportional to the winding current [16] and hence linear design techniques can be very successful at tuning for both steady state and transient performance. For a PMSM, the torque produced is a nonlinear function of both the magnitude of the stator current and the rotor position (see (3.7)). Fortunately, in the d - q axis reference frame, torque is proportional to the magnitude of the q -axis component of the stator current vector (3.10) since this produces a stator mmf component in quadrature to the rotor mmf². Applying control in d - q axes exploits this linear relationship between q -axis current and torque, reducing the complexity of the control problem toward that of the simple DC motor. However, even in the d - q coordinates, the model is nonlinear, with bilinear interaction between both the d and q axis currents and the motor speed, ω_m . It is possible to obtain a Linear Parameter Varying (LPV) (or more accurately a quasi-LPV model) in which the dynamics vary linearly as a function of motor speed. The control of LPV systems is reasonably established now (see [7] for example) but does tend to lead to complex control strategies not suitable in this application. Similarly the well-known technique of feedback linearisation [72] could be applied directly to this model, although such a technique may suffer from robustness problems. Thus, most controllers are designed typically using simple linear control techniques which are tuned, perhaps in a somewhat ad-hoc manner, for the desired transient performance.

An important point to note is that in order to generate current in the q -axis, voltage may also be required in the d -axis. This follows since the stator windings are inductive and thus the phase of the stator current will lag behind that of the stator voltage. The transfer function from applied voltage to current flowing in an arbitrary stator winding is given by $1/(sL_s + R_s)$ and an example frequency response is plotted in Figure 3.3. Note that both the gain and phase drop significantly with increasing frequency and that at high frequencies the phase lag approaches 90 degrees. In rotor coordinates, this means that if the voltage vector is aligned with the q -axis, the current vector may be rotated by up to 90 degrees away from the q -axis. Noting that the d and q axes are 90 electrical degrees apart, this means that at high frequencies the current induced may be almost entirely in the d -axis. Looking at this from another perspective, to ensure that the current is produced in the q -axis (for torque production), the voltage vector may need to be advanced by up to 90 degrees toward the d -axis. The frequency of the stator winding currents is proportional to the electrical velocity of the motor, $\omega_e = P\omega_m$. Therefore, the amount of voltage advance required is dependent upon motor speed, and with appropriate scaling the x-axis of Figure 3.3 can be interchanged for motor speed.

This effect of stator winding inductance on open loop behaviour is also demonstrated in the simulation response of Figure 3.4 where a sequence of voltage step demands are applied to the q -axis. As each successive step is applied to the q -axis voltage, the motor settles at progressively higher speeds. However, the manner in which the motor accelerates to each step changes. Note that at low speed, the step voltage demand produces a large transient in the q -axis but has very little effect on the d -axis and therefore a large proportion of control energy contributes towards accelerating torque. At higher speeds, the same magnitude step to the q -axis voltage has its most significant effect in the d -axis. This split of control energy between d and q axes reduces the control energy contributing to accelerating torque and the motor speed response is more sluggish as a result.

²Note that the d -axis current does not contribute towards torque




Figure 3.3: Transfer function from stator winding voltage to winding current




Figure 3.4: PMSM open loop step responses

3.6 A Linear Control Structure

Using vector control, linear proportional plus integral (PI) controllers can be arranged as shown in Figure 3.5 to form an inner current control loop, and an outer speed control loop. This is a widely used control structure for PMSM speed regulation, due to its transparency, ease of implementation and general effectiveness. The inner loop controls the d and q axis currents by manipulating the applied voltages, and since torque is proportional to the q -axis current, can be considered as a torque controller. The outer loop controls the speed of the motor by manipulating the q axis current demand i.e. demanding a torque. Under normal operation the d -axis current demand is zero so maximum current is applied to the q -axis and therefore contributes toward the production of torque. The coupling between d and q axis currents, shown in the d - q axis model (3.11), has the effect that in practice, the d -axis current will stray away from zero to an extent as transient q -axis demands are made, leading to a reduction in efficiency. This effect can be reduced by suitable tuning of the controller and can also be eliminated, at least in theory, by the application of feedback linearisation [68].




Figure 3.5: Linear Control Structure

The current and q -axis speed controllers used are described by the following discrete state space models where τ_{spd} and τ_{crt} represent the reciprocal of the sample frequencies used in the speed and current controllers respectively. Normalised versions of the parameters used for the practical applications in later chapters can be found in Appendix A.2.

$$K_{spd} \sim \begin{cases} x(k+1) = [1] x(k) + [\tau_{spd}] \omega_{m,error}(k) \\ i_{q,Dmd}(k) = [k_{i,spd}] x(k) + [k_{p,spd}] \omega_{m,error}(k) \end{cases} \quad (3.12)$$

$$K_{crt} \sim \begin{cases} x(k+1) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} x(k) + \begin{bmatrix} \tau_{crt} & 0 \\ 0 & \tau_{crt} \end{bmatrix} i_{dq,error}(k) \\ V_{dq,Dmd}(k) = \begin{bmatrix} k_{i,crt} & 0 \\ 0 & k_{i,crt} \end{bmatrix} x(k) + \begin{bmatrix} k_{p,crt} & 0 \\ 0 & k_{p,crt} \end{bmatrix} i_{dq,error}(k) \end{cases}$$

Note that with this structure, K_{spd} is a simple single input single output (SISO) controller, whereas K_{crt} is a decentralised multiple input multiple output (MIMO) controller. By decentralised it is meant that there

is no coupling between the channels in the controller i.e. it is a diagonal transfer function matrix. The signal $i_{dq,error} = [i_{d,dmd}(i_{q,dmd}) - i_d \quad i_{q,dmd} - i_q]'$ drives the inner-loop MIMO controller, where $i_{d,dmd}$ is a demand signal constructed in a nonlinear fashion from i_q and will be discussed in more detail later.

3.7 Field Weakening Operation

Under normal operation where the d -axis current demand is zero, the back emf induced by rotation of the rotor reduces the magnitude of the voltage phasor measured across the stator windings. The back emf increases in magnitude in proportion to the motor speed until the point at which the back emf is equal but opposite to the applied voltage vector. At this speed, referred to as the *base speed*, the resultant voltage is zero, hence there is no current induced, and consequently no accelerating torque. This is the theoretical maximum speed at which the motor can operate with zero d -axis current. Since in this mode of operation all the applied current contributes toward torque production, the speed range between rest and base speed is referred to as the ‘maximum torque region’.

Some mathematical clarity can be gained by analysing the q -axis state equation of (3.10). In the maximum torque region, $i_d = 0$ and hence we have the simplified expression of (3.13), the solution to which is given as (3.14). From this we can see that when the back emf, $K_e \omega_m(\tau)/\sqrt{3}$, is equal to the applied voltage, $V_q(\tau)$, the solution is simply the free response given by (3.15), so for zero initial conditions, $i_q = 0$ and no torque is produced.

$$\frac{d}{dt} i_q(t) = \frac{1}{L_s} \left[\underbrace{V_q(t)}_{\text{Applied voltage}} - R_s i_q(t) - \underbrace{\frac{K_e}{\sqrt{3}} \omega_m(t)}_{\text{back emf}} \right] \quad (3.13)$$

$$i_q(t) = e^{-(R_s/L_s)t} i_q(0) + \frac{1}{L_s} \int_0^t e^{-(R_s/L_s)(t-\tau)} \left[V_q(\tau) - \frac{K_e}{\sqrt{3}} \omega_m(\tau) \right] d\tau \quad (3.14)$$

$$i_q(t) = e^{-(R_s/L_s)t} i_q(0) \quad (3.15)$$

In order to achieve speeds beyond base speed, a nonlinear flux weakening algorithm is activated which demands a negative d -axis current. Applying current into the d -axis weakens the flux linkage and consequently reduces the back emf induced, with the effect that greater current can be injected into the q -axis [97]. This allows torque to be produced beyond the base speed, enabling acceleration to higher speeds, and also increases the torque capability for speeds immediately below base speed. During flux weakening, also known as ‘phase advance’ operation, the portion of current applied to the d -axis is essentially lost as it does not contribute to torque. As a result, the benefit of higher speed capability comes at the cost of efficiency. The angle, ϕ , by which the phase of the current vector is to be advanced is pre-determined as a nonlinear function of the motor speed and read from a look-up table within the controller. A graphical representation of a typical phase advance map

used is shown in Figure 3.6 for which the d -axis current demand is calculated as shown in (3.16). Note that in practice, phase advance operation must start below base speed to ensure that sufficient torque is available to accelerate the motor beyond this threshold, particularly in the presence of load disturbances.




Figure 3.6: Phase advance map

The effect of motor current phase advance can also be seen mathematically. As previously, we consider the solution to the q -axis state equation of (3.10). For flux weakening operation, we have $i_d \neq 0$ and we consider the case for which the motor speed is constant, giving the solution as shown in (3.17). By inspection, it is clear to see that the zero torque (zero q -axis current) condition occurs at higher ω_m if $i_d < 0$ than for $i_d = 0$.

$$i_{d,Dmd} = -|i_{q,Dmd}| \tan(\phi) \quad (3.16)$$

$$i_q(t) = e^{-(R_s/L_s)t} i_q(0) + \frac{1}{L_s} \int_0^t e^{-(R_s/L_s)(t-\tau)} \left[V_q(\tau) - PL_s \omega_m i_d(\tau) - \frac{K_e \omega_m}{\sqrt{3}} \right] d\tau \quad (3.17)$$

3.8 Saturation within the Control Loop

Various limitations related to hardware introduce different forms of saturation into the control system. It is important to model these correctly in order to ascertain their effect on the closed loop and also to implement anti-windup compensation to deal effectively with the adverse effects of the associated saturation events. Anti-windup is not the only method of dealing with the saturation problem, as discussed in Section 4.2, but for the EPHS application we consider, anti-windup provides a good blend of performance and simplicity of design with low computational demands.

One constraint imposed is a voltage limit resulting from the power source in the vehicle. The terminal voltage of the battery should be in the region of 13.5V when the alternator is functioning correctly but may drop toward

12V depending upon the status of the alternator and the current being drawn from the supply. To simplify analysis, this limit is considered to be a constant and values of 13.5V and 12V are considered for analysis. With a drive stage voltage of 12V, the PWM inverter is able to produce stator phase voltages which vary between $\pm 6V$. This constraint translates to a hexagonal limit in the stator reference frame as shown in Figure 3.7, where any voltage vector contained within this hexagon is achievable.




Figure 3.7: Physical (hexagonal) and imposed (circular) voltage saturation limits on the stator reference frame

Another constraint is a limitation on the current that can be supplied by the power electronic drive circuits. This corresponds to a static saturation limit on the magnitude of the current flowing in the stator windings, i_a, i_b, i_c and can be modelled in the d - q reference frame as saturation of the norm of $[i_d \ i_q]'$. These saturation constraints change the form of the current equations from

$$\frac{d}{dt} \begin{bmatrix} i_a \\ i_b \\ i_c \end{bmatrix} = [A(\theta_e)] \begin{bmatrix} i_a \\ i_b \\ i_c \end{bmatrix} + [B] \begin{bmatrix} V_a \\ V_b \\ V_c \end{bmatrix}$$

to

$$\frac{d}{dt} \begin{bmatrix} i_a \\ i_b \\ i_c \end{bmatrix} = [A(\theta_e)] \begin{bmatrix} \text{sat}(i_a) \\ \text{sat}(i_b) \\ \text{sat}(i_c) \end{bmatrix} + [B] \begin{bmatrix} \text{sat}(V_a) \\ \text{sat}(V_b) \\ \text{sat}(V_c) \end{bmatrix}$$

where $A(\theta_e)$ and B are the state and input matrices of the nonlinear model's state equation. Thus, not only are the inputs to the plant saturating, but some of the states also saturate. A graphical representation of the physical saturation limits within the system is shown by the block diagram of Figure 3.8.

An example of the effect of saturation is shown in Figure 3.9 where the speed controller saturates in the absence of an anti-windup compensator. When the controller saturates briefly for the small step demand, this does not degrade performance appreciably. However, the extended period of saturation resulting from the larger step demand causes a delayed and sluggish response to the reverse step. In this figure, phase advance operation can

also be observed where the d -axis current demand becomes non-zero beyond a normalised speed of approximately 1 unit. Note that this has a direct effect on the q -axis saturation limit which reduces from unity to around 0.5 units as the d -axis demand is increased.

When implementing a controller for a saturating system, it is commonplace to impose software saturation limits with a threshold slightly lower than the physical limits to ensure that these are not actually reached. Any anti-windup conditioning applied will then function according to these lower software limits. Since we desire to compute control in the d - q reference frame, these ‘software’ saturation limits must also be in d - q coordinates, posing some challenges.

The hexagonal voltage limit is difficult to translate into the d - q reference frame and so instead the limit is approximated by a circle, shown in Figure 3.7. The radius of this circular limit would then be a design parameter determining how close the control signals get to the physical limits. This circular limit translates easily into d - q coordinates as saturation of the norm of $V_{dq} = \begin{bmatrix} V_d & V_q \end{bmatrix}'$. Since the signals V_d and V_q both appear within the controller, these can be limited directly.

Current saturation can also be approximated in the d - q reference frame in a similar way as for voltage saturation. However, we cannot easily impose a similar saturation limit within the controller because the currents are states of the plant rather than directly controlled inputs. An alternative and convenient way of implementing a current saturation function is to limit the norm of $i_{dq,dmd} = \begin{bmatrix} i_{d,dmd} & i_{q,dmd} \end{bmatrix}'$ which is the current demand fed into the inner loop current controller. The idea is that, assuming that the inner loop is stable and reasonably well damped, we can expect the magnitude of the current induced in the motor to be bounded by a value similar to the bounded demand. Thus, provided that the current controller is not tuned too aggressively, this should prevent the magnitude of $i_{dq} = \begin{bmatrix} i_d & i_q \end{bmatrix}'$ reaching the physical limit. This is not an ideal method to limit the current but is simple to implement and transparent. Alternative ideas will be presented in later sections of the thesis.

A further complication is that in both cases, these saturation limits translate into d - q axes as saturation of the norm of a vector signal rather than element-wise saturation of the vector signals. In the case of the voltage limit, saturation of the norm can be imposed by scaling the V_{dq} voltage demand down when the norm exceeds the voltage limit. By scaling both elements simultaneously, this ensures that the phase of the signal is unchanged and only the amplitude is attenuated. In the case of current demand saturation the same method could be applied, however, the optimal compensator designs presented in Chapter 5 require linear controller models for their synthesis and the d -axis current controller is nonlinear. Therefore in order to fit the compensator design framework an alternative method is used: Saturation is only imposed on $i_{q,dmd}$ and the saturation limit is varied as a function of $i_{d,dmd}$ such that the norm of $i_{dq,dmd}$ is constrained to the saturation level. In this way, only the linear q -axis part of the speed controller need be considered for anti-windup design, and the d -axis current demand can be considered as a disturbance signal to the linear control loop. This implementation of software saturation limits is depicted in block diagram form in Figure 3.10. In this figure, $\text{sat}(\|\cdot\|)$ represents saturation of the 2-norm of the input vector signal.




Figure 3.8: Block diagram representation of the physical saturation limits




Figure 3.9: The saturation problem




Figure 3.10: Implemented software saturation limits

3.9 The Effect of Discretisation

To implement vector control a digital controller is needed due to the online computation required. This means that the d and q axis voltage demands are only updated at discrete instances in time. These V_{dq} demands are translated into 3ϕ waveforms based on a sampled measurement of the rotor electrical position, θ_e , so the stator phase voltages also are only updated at discrete instances in time. At the sample instance, this is correct and the d - q axis voltages measured in the motor are the same as the demands. However, in reality, θ_e is varying continuously with the effect that towards the end of the sample period, the rotor magnetic field vector has rotated towards the stator electromagnetic field which has remained static for the sample period. This has the same effect as the applied voltage phasor moving into phase lag during the sample time and hence alters the distribution of voltage between the d and q axes.

At low motor speeds the rate of change of the rotor position is low and hence the error in the phase voltages is also low. At high speed however, θ_e can vary significantly between samples and the reduction in q -axis current incurred can reduce the attainable speed and reduce efficiency. This problem can be solved by increasing the sample rate significantly but this may be undesirable for a commercial product due to the increased hardware cost so a certain amount of error has to be accepted. Therefore an important feature of this controller structure is that there will always be an amount of error involved in controlling the phase relationship between the rotor and stator fields, and that this error is accentuated at high speeds. The size of this error is dependent upon the operational speed range of the motor and is also sample dependent. Because this error would be known for a given application it could be compensated for and in industry, these type of software fixes are commonplace.

Chapter 4

Introduction to Anti-Windup

In this chapter the problem of ‘windup’ is explained with respect to systems with plant input constraints and existing methods to deal with this problem for linear systems are introduced. The issues of closed loop stability and compensator performance are addressed and methods of synthesising optimal compensators are explained. The work from hereon splits quite cleanly into two main strands; application of current techniques to an industrial problem, and research into improving existing techniques. These two topics will be addressed separately in subsequest chapters.

4.1 The Problem of Actuator Saturation

Normally, when designing a linear controller for a certain system, a linear model of the plant is generated and then a controller is designed using linear methods to satisfy some pre-defined stability and performance criteria. A satisfactory design will guarantee stability of the linear closed loop and satisfy given performance objectives which may be defined in terms of rise time, settling time, maximum overshoot and disturbance rejection for example. Provided that the linear model is an accurate representation of the true system and the controller designed is sufficiently robust, these performance criteria will be met in the real system also.

When a system is subjected to plant input constraints such as limits on the magnitude of the control signals, a non-linearity is introduced to the system as shown in Figure 4.1. If, during operation, the controller output exceeds these limits then linearity of the closed loop is lost and the performance criteria achieved by the linear model may not be met. Typical outcomes are failure to meet reference demands, slow rise times and even oscillatory transients or instability, although the effect observed is dependent upon the characteristics of that particular system.

In the simple case of a SISO system that is open loop stable, saturation may simply increase the rise time of the closed loop system to a step demand or prevent a reference demand being met due to the reduction in available control effort. However, the behaviour of the controller during periods of saturation may be more detrimental to performance. The mis-match between controller output and plant input during saturation events allows the




Figure 4.1: Generic closed loop control system with plant input saturation constraints

controller output to evolve in a manner that does not match the system response. In the case of proportional plus integral (PI) control, this phenomenon is referred to as ‘integrator windup’ as the integral state continues to accumulate during saturation, driving the system further into saturation while failing to have any further effect on the performance output. Consider the behaviour of such a system when the plant input constraint causes the reference demand to be infeasible i.e. a reference which in steady-state cannot be tracked. The persistent error signal causes integration to continue for as long as the reference persists in an attempt to eliminate the steady-state error.

If integrator windup is allowed to occur, when the reference demand changes and the control signal is required to drop below the saturation level, the accumulated integral action serves to work against tracking of the new reference and a significant amount of time is required for the energy in the integrator to dissipate before the controller can contribute positively towards attaining the new reference (Figure 4.2). This specific problem was traditionally termed the ‘windup’ problem since any performance degradation is a direct result of integrator windup. During the 1950s and 1960s, a number of ad-hoc methods were devised to prevent or limit the extent of this phenomenon by directly influencing the integrating function. One such method was published by Fertik and Ross in 1967 [15] and R.M. Phelan suggested turning off integrators during saturation in his 1977 book, *Automatic Control Systems* [70]. There are other application-specific schemes available and, due to their development within industry, many are not documented within the wider literature.

Over time, usage of the term ‘windup’ has evolved and it is now used more generically to describe the performance degradation effects associated with plant input saturation. This includes, but is not limited to, SISO systems with integral control. For example, multivariable systems, even with static controllers, can exhibit undesirable behaviour when one or more of the plant inputs saturate as the distribution of control effort is altered -see Figure 4.3. Furthermore, modern \mathcal{H}_∞ type controllers often do not have integrators but the behaviour of their states can still be adversely affected by saturation. Finally, even SISO systems with purely proportional control can exhibit oscillatory or even unstable behaviour when saturation limits are exceeded.

While a system is in saturation, the input to the plant is constant and the system becomes “pseudo open loop”. During this period the dynamics of the system revert, in a sense, to the dynamics of the open loop plant which may be undesirable and/or unstable. Multiple input, multiple output (MIMO) systems often have cross-coupling between channels, so in addition to the effects described above, saturation in one channel can cause detrimental performance in another channel or even instability. Clearly, actuator saturation is a problem that




Figure 4.2: The effect of integrator windup

needs to be addressed in many applications as any real system is likely to have finite limits on plant inputs, be it a mechanical or electrical system or chemical process.

4.2 Solving the Problem of Actuator Saturation

There are a number of approaches to dealing with the problem of actuator saturation. The main methods along with their relative merits and drawbacks are summarised as follows:

- Saturation Avoidance.** Perhaps the simplest approach is to design a controller which will not violate the saturation limits for any input demand, for instance, by using low gain controllers without integral action. Since a successful design will not cause saturation, the associated problems would not be encountered and linear performance will be achieved. Of course, controllers designed in this manner tend to be conservative since the control action will be reduced across the whole range of input demands i.e. small signal performance is forfeited to prevent saturation at the extremes of control behaviour. As such, this approach is not appropriate for many systems. \mathcal{H}_∞ controllers can aid performance improvement with this approach since they tend to distribute control effort in a more efficient manner, particularly for multivariable systems, but performance is limited fundamentally by linearity of the controller since they treat large and small signal behaviour in exactly the same way.
- Saturated Linear Control.** Another approach is described as saturated linear control in which saturation is not avoided, but the control system is designed to ensure that the saturated closed loop remains stable.




Figure 4.3: Saturation inducing control direction change

This has the benefit of simplicity but not all systems can be guaranteed to be stable during saturation and even if stability is guaranteed, performance may not be acceptable. An example can be found in [21]. Another important stream of papers considers the low-and-high gain techniques [51, 52] in which saturated linear control is achieved by simply increasing the gain of a low-gain controller. All open loop systems with poles in the closed left half complex plane can be stabilised with such control laws. However the price paid for this simplicity is a lack of flexibility in the design i.e. no tuning rules and a prescribed architecture.

- Model Predictive Control.** Model predictive control [49, 75] is a strategy in which the control constraints are incorporated into an online finite horizon optimisation procedure, producing controllers which respect the saturation constraints directly. This may well give the best performance potential of all control strategies designed to cope with saturation but it has its drawbacks. Computational demands for implementation are high and thus the approach is not commercially feasible for many applications. In addition, the computation time required to solve the optimisation can preclude application to systems with fast dynamics. At present, the primary application of MPC is in the process industries where time constants are relatively long but improvements in computational power of micro-controllers and the development of more efficient optimisation routines are extending its applicability to ever faster systems. There are also concerns about robustness and tuning since it normally relies on the standard linear plus quadratic (LQ) weighting matrices.
- Linear Conditioning/Anti-Windup** The approach considered in this thesis is referred to as ‘linear conditioning’ or ‘anti-windup compensation’. This involves two separate linear controller elements and a two stage design process. The first stage is the design of a linear controller to meet the performance specification in the absence of saturation. Following this, the system is augmented by an additional linear element referred to as the ‘anti-windup compensator’ which becomes active at the onset of saturation. This compensator is designed to maintain stability of the system during saturation of the nonlinear system and minimise the degradation of performance associated with the saturation event. The benefit of

this approach in terms of performance is that the small signal response of the system is not forfeited and any linear control design methodology can be used to construct the local linear controller. Computational demands can also be low enough for application on high bandwidth and cost sensitive applications. Many variants of anti-windup compensation exist of which a summary will be introduced in the following section.

4.3 Approaches to Anti-Windup Compensation

There are many forms of anti-windup compensator design in existence which function in very different ways. The different designs vary significantly in complexity, structure and performance. Early designs were largely heuristic and as such it could be difficult to design and predict their performance. As tools emerged to analyse these nonlinear systems, more structured designs emerged and tools were generated to synthesise compensators directly which guarantee stability of the saturated system and optimise performance in some sense. This section provides an overview of some of the most well known compensation schemes in existence. Before looking at the specifics of each design let us consider a generic anti-windup control architecture by means of introduction to the concept, and to highlight the areas where differences arise.

4.3.1 A Generic Anti-Windup Architecture

A generic representation of an anti-windup compensated closed loop control system is shown in Figure 4.4. This consists of a controller, K , that responds to reference signals $r \in \mathbb{R}^{n_r}$ and plant outputs $y \in \mathbb{R}^{n_y}$ to generate the control signals $u \in \mathbb{R}^m$. The internal states of the controller are denoted $x_c \in \mathbb{R}^{n_c}$. The plant, G , has control inputs $u_m \in \mathbb{R}^m$, disturbance inputs $d \in \mathbb{R}^{n_d}$ and internal states $x_p \in \mathbb{R}^{n_p}$. The block translating u to $u_m \in \mathbb{R}^m$ - the measured inputs of plant G - represents the saturation function. In multivariable form, the saturation function is a diagonal matrix of scalar nonlinear functions that bound the magnitude of the corresponding input vector elements. This is described mathematically as shown in (4.1) where u_i represents the i th element of the control vector u , $\text{sat}(u_i)$ represents the saturated version of that same signal and \bar{u}_i represent the symmetrical saturation limits applied to the i th channel. The input output relationship for a SISO case is depicted in Figure 4.4. Non-symmetric limits can also be handled but are a trivial extension to the symmetric function shown here and so have been omitted.

$$\text{sat}(u_i) = \begin{cases} \bar{u}_i & \forall u_i > \bar{u}_i \\ u_i & \bar{u}_i \geq u_i \geq -\bar{u}_i \\ -\bar{u}_i & \forall u_i < -\bar{u}_i \end{cases} \quad (4.1)$$

The anti-windup compensator, AW , shown in Figure 4.4 detects when the plant input limits have been exceeded, by monitoring \tilde{u} , and then influences the controller in some way to improve the behaviour of the system. Some compensators only require knowledge that saturation has occurred but more powerful designs




Figure 4.4: Generic anti-windup architecture

also desire to respond to the magnitude and direction of saturation in their anti-windup action. A useful way to capture these different requirements generically is to define a signal $\tilde{u} \in \mathbb{R}^m$ as the difference between saturated and unconstrained control signals. The simple designs that only require to know which signals are saturating can be considered simply to apply a logical test to the signal such as $|\tilde{u}_i| = 0$.

It is possible that the unconstrained control signal u and its constrained counterpart may be real signals that need to be measured and the saturation function to be a physical phenomena. However, in digital control applications, it is most common for the saturation function to be a non-physical constraint implemented in software. Often these software limits represent physical constraints, but are set slightly lower than the real limits. This removes the requirement for a sensor to measure u_m as this is now a signal internal to the controller and by preventing the physical saturation limits from being met, mechanical durability of the system can be improved. It is also possible that physical limits do not exist but software limits are imposed for other reasons such as to minimise power usage or noise.

4.3.2 Conditional Integration

For systems with a stable plant and simple PI control it can be adequate simply to stop the integrator state from accumulating when certain conditions are met. Two common variants are

1. Stop integrating when the controller output is saturated
2. Stop integrating when the integrator state exceeds given bounds

These methods will prevent, or at least reduce, integrator windup and have the benefits of simple tuning and implementation. Condition 1, referred to as ‘intelligent integration’ by Krikelis [47], does not require any tuning whatsoever. However, the work of Krikelis describes an implementation which *can* be tuned in which deadzone feedback is placed around the integrator and the thresholds and slopes of the deadzone function can be altered. Condition 2 requires a bound on the integrator state to be set but this corresponds to only one tuning parameter

per integrator. A disadvantage of these methods are that the plant input(s) are allowed to remain in saturation, meaning that the feedback loop is effectively broken. Therefore, for acceptable performance the plant must have stable and acceptably damped open loop behaviour. Even then, stable open loop dynamics do not imply that the nonlinear closed loop will be stable. No guarantees of stability are afforded by these approaches and their application are restricted typically to PI controllers. In addition, because the integrator state is not altered during saturation, the stored integrator state may restrict the speed of response to a subsequent change in the reference.

4.3.3 Ad-hoc Integrator Reset Methods

In addition to the conditional integration methods there are also simpler methods where the integrator state is reset to a pre-specified value when certain conditions are encountered. Perhaps the simplest implementation resets the integrator state to zero when the controller output reaches the saturation limit, preventing saturation from continuing. Although integrator windup is prevented entirely, when the integrator state is reset a sudden drop in the control signal magnitude is caused i.e. the control is discontinuous and system performance is degraded. An illustration of this is shown in Figure 4.5 where the integrator reset causes a substantial degradation in tracking performance. Note that the model used varies from the norm in that the saturation limits are parameter varying rather than constant.




Figure 4.5: Reset triggered when control signal reaches saturation level

Numerous variations to this scheme are possible including changes to the conditions required to initiate a reset, and variation in the values the integrator is reset to. One example is to trigger the reset of the integrator state only when the control signal exceeds a certain threshold above the saturation limit. Integrator windup is not

prevented completely but an upper bound on windup is imposed. An example of this is shown in Figure 4.6 where a reset will not be triggered until the control signal exceeds the saturation level plus 0.25. Comparing this result to that given in Figure 4.5 shows the improvement in performance.




Figure 4.6: Reset triggered when control signal exceeds saturation + 0.15

Resetting the integrator state to a non-zero value during saturation is discussed in [50]. This can prevent the control signal instantaneously dropping below the saturation level if this non-zero value is chosen appropriately. However, one common problem remains: Whenever an integrator is reset, the balance between proportional and integral action is temporarily changed leading to a period of altered system dynamics. An example of this is discussed in 5.1.1 where momentary saturation at the onset of a step reference diminishes integral action, leading to a much slower rise time than if integration was allowed to continue.

These methods are easy to implement and require very little online computation. They can work very well for simple systems, but none afford any guarantees of performance or stability in the general case. These methods will not work well with systems that are open loop unstable and often do not perform well with high order or MIMO systems.

4.3.4 Back Calculation (and Tracking)

In the Back Calculation method, first proposed by Fertik and Ross [15], the stored value of the integrator state is recomputed (back calculated) when saturation occurs such that the controller output at the next sample lies approximately at the saturation limit. In this way, the controller makes full use of the available control energy but its proximity to the saturation limit enables it to effect a quick response to a change in the reference demand. A simple implementation of this for a digital PI controller is shown in (4.2) where K_i and K_p represent the

proportional and integral gains, τ represents the sample period, e represents the tracking error and u and u_{max} represent the control signal and limit respectively.

$$x(k+1) = \begin{cases} x(k) + K_i \tau e(k), & |u(k-1)| \leq u_{max} \\ u_{max} \text{sign}(K_p e(k)), & |u(k-1)| > u_{max} \end{cases}$$

$$u(k) = x(k) + K_p e(k) \quad (4.2)$$

It can be beneficial not to reset the integrator in one sample but rather to reset the integrator dynamically with a given time constant. This special case is referred to as ‘Back Calculation And Tracking’ (BCAT) since the anti-windup compensator causes the controller output to ‘track’ back to the saturation level. The latter approach can be represented very neatly in block diagram form as shown in Figure 4.7 where T_t represents the time constant with which the reset takes place. This method is described in more detail in [3] and is also referred to as ‘classical’ anti-windup in [14].




Figure 4.7: Back calculation and tracking block diagram

Aström and Hägglund suggest that the back calculation time constant, T_t , should be chosen according to (4.3) where T_d is the derivative time constant (if used) and T_i is the integral time constant. These time constants can be related to proportional, integral and derivative gains, K_p, K_i, K_d , as shown in (4.4) and (4.5). If we prefer to think in terms of gain rather than time constants, the suggested range for the anti-windup feedback gain $1/T_t$ to be tuned within is given by (4.6) provided that $K_d, K_p \neq 0$.

$$T_d < T_t < T_i \quad (4.3)$$

$$T_d = K_d/K_p \quad (4.4)$$

$$T_i = K_p/K_i \quad (4.5)$$

$$\frac{K_i}{K_p} < \frac{1}{T_t} < \frac{K_p}{K_d} \quad (4.6)$$

4.3.5 Conventional Anti-Windup

Another simple technique, found in [14] and [42], involves comparing the constrained and unconstrained control signals and feeding this back through a high gain to the controller input (Figure 4.8). This anti-windup feedback artificially reduces the controller input, y_k , during saturation to drive the controller output toward the saturation level. The magnitude of the feedback gain determines how close the controller output gets to the saturation limit, and if a suitably high gain is used the control signal can be constrained to the saturation limit, eliminating windup.




Figure 4.8: High gain anti-windup scheme block diagram

Mathematical Analysis:

An approximate analysis of the anti-windup behaviour is as follows. For a simple PI controller with state-space realisation $K(s) \sim (0, 1, k_i, k_p)$, when the system is in positive saturation, i.e. $u_m = \bar{u}$, the controller output, u , is given by (4.7) where $e = r - y_k$

$$u(s) = \frac{k_p s + k_i}{s + AW(k_p s + k_i)}(e + AW\bar{u}) \quad (4.7)$$

Steady-state analysis:

Using the final value theorem, the controller output to a unit step input as time tends to infinity is given by

$$\begin{aligned} u_{ss} = u(t) &= \lim_{t \rightarrow \infty} \frac{k_p s + k_i}{s + AW(k_p s + k_i)}(e + AW\bar{u}) \\ &= \frac{1}{AW}(e + AW\bar{u}) \end{aligned}$$

and provided the AW gain is large enough, $u \approx \bar{u}$. Thus, the controller output will be forced down toward the saturation level, \bar{u} , and the integrator will not wind up. Note that this analysis assumes that a steady-state condition is reached and so for the analysis to hold, the closed loop (nonlinear) system must be stable.

There are no stability guarantees with this technique but stability is usually observed, in simple, low order stable systems, provided that the anti-windup gain is not too large. Tuning of the anti-windup gain is typically a simple iterative process using experimental tests and/or simulations. Good performance can be obtained for simple systems but responses for higher order and/or MIMO systems can be undesirable and unstable. For discrete-time systems, the magnitude of the anti-windup feedback gain is restricted by the sample rate used and so the achievable performance can also be limited by such implementation issues.

4.3.6 The Observer Approach

The back calculation and tracking method for PI(D) controllers prevents windup by causing the controller output to track down to the saturated output during saturation by manipulating the controller state. A similar approach can also be applied to higher order controllers such as \mathcal{H}_∞ designs by means of the observer approach [1, 2, 76]. In this method the saturation nonlinearity is considered to cause a mis-match between the controller states and the control signal applied to the plant, and the ‘error’ between the controller output and plant input is interpreted as an observation error.

Consider that the nominal controller is given by the state-space realisation (4.8) and that we wish to include an anti-windup feedback to the controller states that is proportional to the difference between the controller output and plant input. This can be represented by the state-space model of (4.9) where L is the anti-windup feedback gain.

$$K(s) \sim \begin{cases} \dot{x} = Ax + B_r r + By \\ u = Cx + D_r r + Dy \end{cases} \quad (4.8)$$

$$\begin{aligned} \dot{\hat{x}} &= A\hat{x} + B_r r + By + L(u_m - u) \\ u &= C\hat{x} + D_r r + Dy \\ u_m &= \text{sat}(u) \end{aligned} \quad (4.9)$$

As the state-space realisation of the nominal controller is known, this can be interpreted in an observer structure by substituting the definition of u from (4.8) into the state equation of (4.9), yielding the state-space description of the observer-based controller in (4.10). Note that in this approach, there is no need to have a separate nominal controller and observer; a single state-space representation is sufficient. A block diagram of the corresponding practical implementation is shown in Figure 4.9.

$$\hat{K}(s) \sim \begin{cases} \dot{\hat{x}} = (A - LC)\hat{x} + (B_r - LD_r)r + (B - LD)y + Lu_m \\ u = C\hat{x} + D_r r + Dy \\ u_m = \text{sat}(u) \end{cases} \quad (4.10)$$

From (4.9) it is evident that in the absence of saturation i.e. when $u_m = u$, the observer-based controller will behave exactly as the nominal controller of (4.8). During saturation, the controller behaviour is governed by (4.10) which can be influenced by the choice of observer gain, L . For systems with integral control, the integral pole can be moved from the origin in the complex plan along the negative real axis such that the controller is open loop stable during saturation. In this case, rather than the integrator ramping up in response to a persistent input (u_m, r, y) , the controller output u will settle to a constant value at steady-state, limiting the extent of integrator windup. For controller types that are already open loop stable, the observer gain can be tuned to speed up convergence of the controller states to an equilibrium condition during saturation. Moreover, if the




Figure 4.9: The observer approach

observer gain is large enough it will cause the observation error $u_m - u$ to converge to zero, thereby minimising the extent of integrator windup.

The extra “information” about the saturated controller output can assist the controller in maintaining desirable closed-loop behaviour. In fact, by appropriate choices of the observer gain, L , it is easy to see that the CAW scheme (and also the Hanus scheme introduced in Section 4.3.8) can be considered special cases. Furthermore, the observer scheme is the basis of Kothare’s [43] unifying scheme for anti-windup compensators. However, despite these appealing features, it is important to remember that the observer scheme by itself will not *guarantee* nonlinear stability of the arising closed-loop system - extra conditions must be imposed for this to be the case.

4.3.7 The Observer Technique

Another method referred to as the ‘observer technique’ is presented in the work of Peter Hippe [37]. Not to be confused with the observer-based anti-windup methods of Aström and Glattfelder; here an observer is employed to recreate state estimates for the plant, allowing a state feedback control law to be applied as shown in Figure 4.10. With state feedback control, there are no controller states to windup and so in one sense, the windup problem is solved. However, the presence of control constraints still has an impact on closed loop performance. This approach does not fit into the generic framework of Section 4.3.1 as there is not a clear demarcation between nominal controller and anti-windup compensator and so strictly speaking would be best described as a state-feedback control law with an anti-windup feature rather than an anti-windup method in its own right. It has been included here to make the distinction between this and the observer-based anti-windup approach of 4.3.6.

Any states that can be observed in the measured outputs, y_m , are reconstructed via the matrix Ψ , and the remaining un-observable states are reconstructed using the state observer. Note that the measured plant input is used as the input to the observer rather than the unconstrained control signal, avoiding observation errors. A significant benefit of this approach is that anti-windup is essentially provided for free and does not require any tuning. Two obvious disadvantages are (i) that a higher computational overhead is required for implementing the observer compared to PI control implementations with add-on anti-windup compensators; and (ii) the method is




Figure 4.10: The observer technique of Peter Hippe

restricted to state-feedback control laws. Other drawbacks associated with conventional state-feedback control also apply. For instance, performance is dependent upon the accuracy of any reconstructed and/or observed states.

4.3.8 Hanus Conditioning

Hanus was one of the first practitioners of anti-windup to seek a generic anti-windup controller design, having identified that the designs of the day were all dependent upon the controller structure, controller type, or the nature of the constraints on the manipulated variable. He proposed that a good anti-windup system should possess the following properties [27]:

- Independence of the control structure. Later work determined that this is typically achievable if one allows the anti-windup compensator to be of sufficiently high order [100]
- Independence of the controller type. Later work determined that again, this is achievable if one allows the anti-windup compensator to be of sufficiently high order [100]
- Independence of the controlled system. Later work determined that in general this is not achievable, since the plant determines the fundamental behaviour of a system during saturation.
- Independence of the nature of the constraints on the manipulated variable. Later work determined that is, in general, not the case.
- Simplicity

He used the term windup to represent the phenomena observed when a mis-match between the controller output and plant input occurred and attributed the windup phenomenon with the information storage properties of the controller. In addition, he considered that the nominal control behaviour should be designed for the system

in the absence of a windup cause and that controller behaviour should revert to the nominal design when the windup cause disappears.

Hanus considered two main causes for the mis-match between controller output and plant input; actuator nonlinearities, for which actuator saturation and rate limits are examples, and the switching between two control modes, for which the example of manual and automatic control was given. However, in the scope of this thesis, we focus on magnitude saturation only.




Figure 4.11: Hanus conditioning

The essence of Hanus' *conditioning technique* involves 'piloting' the controller state variables in order to reduce the windup effect. This is achieved by implementing the feedback structure shown in Figure 4.11 where the direct feed-through term, D_c , of the nominal controller is separated to allow the anti-windup feedback to effect the controller states only. The anti-windup feedback gain, D_c^{-1} , is a partial inversion of the controller at $s \rightarrow \infty$ and by manipulating the error signal, has the effect of reducing the reference signal to a *realisable* level i.e. one for which no saturation occurs and hence $\tilde{u} = 0$.

It should be noted that for the case where the nominal controller is of PI design, Hanus conditioning can be interpreted as a special case of back calculation and tracking. In this case D_c is equal to the proportional gain K_p and the Hanus anti-windup gain, $1/K_p$, entering the controller just prior to the integral gain is equivalent to the BCAT gain $1/T_t = K_i/K_p$ which enters the controller just prior to the integrator (Figure 4.7). This choice of BCAT gain corresponds to the minimum value suggested by Aström and Hägglund in (4.6).

The Hanus' scheme can also be represented in the form of Figure 4.12 which is more convenient for analysis. By inspection of Figure 4.12, the transfer function from \tilde{u} to u is seen to be

$$-K(s)D_c^{-1} + 1 \quad (4.11)$$

and thus for all frequencies which satisfy $|K(s)| < D_c$, the gain is negative, causing the anti-windup action to reduce the controller output toward the saturation level. This condition should be achieved by all sensibly designed controllers as the dominant control effort should lie in the low frequency region. This analysis also




Figure 4.13: Internal Model Control (IMC) anti-windup design

operation when in positive saturation i.e. $u \geq (\bar{u})$. In this condition, the signal y_k is defined as shown in (4.13).

$$y_k = y + \tilde{y} \quad (4.12)$$

$$= G\bar{u} + G_m(u - \bar{u}) \quad (4.13)$$

By defining the direct model dynamics equal to that of the plant, $G_m = G$, the equation simplifies to $y_k = Gu$ which is exactly as if saturation had not occurred. Therefore, the controller will behave as if saturation was not present and in one sense the problem of windup is solved. Provided that the plant is open loop stable, the nonlinear closed loop will also be stable. However, saturation will still take place and when it does, the nominal controller is unaware of its occurrence. In this mode of operation, the controller essentially seeks to cause y_k to track the reference demand regardless of any difference between y and y_k . As a result, undesirable characteristics of the open loop plant may become evident in the plant output, y , [14]. If the plant contains slow or underdamped dynamics this can lead to undesirable closed-loop behaviour characterised by oscillatory behaviour and/or slow modes - even though y_k smoothly tracks the reference. The limited control of the plant output afforded whilst the system is saturated opens up the way for extensions to the IMC scheme such as found in [104] that aim to improve performance of IMC methods.

More design flexibility is given if the direct model is allowed to differ from the linear model of the plant dynamics, $G(s)$, but this leaves the question of how to design an appropriate filter and also how to characterise performance of the system. Many of the approaches mentioned previously are simple enough to be tuned in an entirely ad-hoc fashion using simulation analysis and successive iterations to the tuneable parameters. However, for this approach and some more complex designs yet to be introduced, some analysis and synthesis tools are required. The basis for the tools to be presented are coprime factor representations, which will be introduced.

4.4 Optimal Dynamic and Static Anti-Windup

The designs presented so far all have their merits and may well be desirable choices for certain applications but they are also not without their limitations. The introduction of a saturation nonlinearity into the closed loop system can result in very complex behaviour and de-stabilise an otherwise stable system. As a result, there is often no guarantee that a simple anti-windup approach will result in a stable saturated system, and even when

stability *is* provided, acceptable performance may not be provided unless certain additional conditions are met; for instance, in the case of IMC anti-windup, performance may be poor if the open loop dynamics are slow and/or poorly damped. In addition, many lack generality in that they are applicable only to specific control structures such as PI control or state feedback.

A more general criticism levelled at the simple ad-hoc methods are that they tend to focus on the behaviour of the controller and ignore the effect of saturation on the behaviour of the plant. As identified earlier in the chapter, the effects of saturation are not limited to windup of the controller states. During saturation, we can consider the system to be pseudo open loop where at least one plant input is constant, dictated by the constraint. Until the point at which this control signal comes out of saturation the system response is governed, at least partially, by the dynamics of the open loop plant which could be slow, oscillatory, or even unstable. As such, although the simple approaches may well be able to constrain integral control states and/or bring control signals out of saturation, overall performance may still be poor.

It is desirable to be able to produce compensator designs that are independent of the type of nominal controller employed, provide a guarantee of nonlinear stability and consider explicitly the performance of the resulting nonlinear closed loop system. In recent years modern control design methods have been applied to the field of anti-windup that enable these objectives to be met and a number of these approaches will be introduced here. The resulting designs can have many degrees of freedom and so ad-hoc tuning methods are considered to be insufficient. As a result, the concept of optimising performance in some sense and providing a limited number of tuning parameters to the designer is of central importance.

4.4.1 The Weston and Postlethwaite Framework

In order to provide the potential of high performance for as many systems as possible, a framework is adopted in which anti-windup signals are injected into the system at both the controller input and output as shown in Figure 4.14. This framework was originally introduced in [99] and has been widely used since. Anti-windup feedback via θ_1 allows the compensator to bypass the controller dynamics and feedback via θ_2 allows the compensator to influence the internal states of the controller, giving a great deal of flexibility without requiring direct access to the controller states. The compensator, Θ , may be parameterised by a static gain or as a dynamic filter and in the multivariable case, a gain or transfer function matrix. Given this framework, our attention is turned to the problems of how to assess stability of the nonlinear system for a given choice of Θ , and how to characterise performance.

Weston and Postlethwaite [100] proposed that the compensator be parameterised using a linear transfer function, M , and a replica of the plant dynamics from control inputs to performance outputs, G_2 , as shown in Figure 4.15. Many anti-windup compensators can be represented in the framework of Weston and Postlethwaite, corresponding to different choices of M . This is useful since it allows the behaviour of very different designs to be compared using a common framework. Similar work in existence is the unified framework of Kothare et al.




Figure 4.14: Anti-windup compensated closed loop

[43]. However, in this work, the compensator is parameterised by two transfer functions to be chosen, U and V , rather than one and so for analysis and synthesis purposes it is less appealing. Miyamoto and Vinnicombe [62] also explored Kothare’s scheme but their work can also be translated to the framework of [100]. The most obvious parameterisation of M that corresponds to a known anti-windup design is $M = I$, giving the IMC design of Section 4.3.9.

$$\begin{bmatrix} \Theta_1 \\ \Theta_2 \end{bmatrix} = \begin{bmatrix} M - I \\ N \end{bmatrix} \quad (4.14)$$




Figure 4.15: Compensation with M

Note that in the Weston and Postlethwaite framework, the anti-windup signals θ_1 and θ_2 are re-labelled as u_d and y_d , and the signal names u_{lin} and y_{lin} are also defined (Figure 4.15). The reason for this can be seen by the application of some simple algebra as follows.

The signal labelled y_{lin} is defined as shown in (4.15). The signal y_d is defined as shown in (4.16). By making the substitution $u_m = u - \tilde{u}$ and working back, the signal y can be defined as shown in (4.17). Finally, by substituting (4.16) and (4.17) into (4.15), we arrive at the result of (4.18). This shows that for anti-windup conditioning with M , the nominal controller, K , ‘sees’ the nominal plant, G , in spite of the presence of saturation and anti-windup compensation. As a result, saturation and the design of the compensator do not affect the behaviour of the nominal controller, and the signals u_{lin} and y_{lin} can be considered to represent the control and output responses of the ideal linear system. The signals u_d and y_d , represent the deviation of the control signals and performance outputs of the compensated system from that of the linear system.

$$y_{lin} = y + y_d \tag{4.15}$$

$$y_d = G_2 M \tilde{u} \tag{4.16}$$

$$y = G_2 u_{lin} - G_2 M \tilde{u} \tag{4.17}$$

$$y_{lin} = G_2 u_{lin} \tag{4.18}$$

Figure 4.15 can also be represented equivalently as shown in Figure 4.16, referred to as the decoupled representation. The structure of the lower portion of the diagram - the nominal linear system - is validated by (4.15). The disturbance filter is a direct carry over from Figure 4.15, and the nonlinear loop simply makes use of the deadzone operator in the identity $Dz(u) = u - \text{sat}(u)$ where $u_m = \text{sat}(u)$. This de-coupled representation proves to be very useful for the analysis of both stability and performance as will be explained.




Figure 4.16: Decoupled representation of compensation with M

The nominal linear system with reference input r and performance output y_{lin} is considered to represent the ideal system behaviour and is designed to stabilise the plant and achieve the nominal performance objectives while the system remains in the linear region (not saturated). For stability of the saturated system we require that the linear transfer function $G_2 M$, referred to as the disturbance filter, be stable and also that the nonlinear

loop be stable. Provided that the plant is open loop stable and M is chosen as a stable transfer function, the problem of guaranteeing stability of the saturated system is reduced to that of ensuring that the nonlinear loop is stable. Thus for open loop stable systems, the stability condition is simplified greatly to that of a single feedback loop with dynamic element $M - I$ and an isolated nonlinearity, which in this case is the deadzone function. In terms of performance, the signal $y_d = y_{lin} - y$ represents directly the deviation of the plant output from the desired behaviour and so can be very useful in defining measures of performance.

4.4.2 Performance of Anti-Windup Compensators

Systems incorporating anti-windup compensation are expected to operate in the linear region for the majority of the time. Thus, linear performance should equate to good performance provided that the linear controller has been designed appropriately. When the system does saturate, the response will degrade and deviate from linear performance. Hence, the notion of good performance of an anti-windup compensator is perhaps expressed best as *minimising the deviation from linear performance* during and immediately after saturation. The difficulty here is to find an appropriate way of expressing this notion of performance mathematically which allows for a tractable performance optimisation. The work of Weston and Postlethwaite in [100] (and also related work by [84] and [62]) gives a reasonably intuitive way of optimising the deviation from linear performance.

Considering Figure 4.16, it follows from the previous section that y_d represents the deviation of the nonlinear (saturated) systems behaviour from the intended linear behaviour, due to the linear control signal u_{lin} . Hence the nonlinear map $\mathcal{T} : u_{lin} \mapsto y_d$ gives important information about the nature of the deviation which can occur. It would therefore be desirable to minimise the “size” of this map, to ensure that the deviation in y_d is small when u_{lin} causes saturation. A computationally tractable way of doing this is by minimising the induced \mathcal{L}_2 norm, or ‘ \mathcal{L}_2 gain’. That is we would like to synthesise an anti-windup compensator (or equivalently choose $M(s)$) such that we have

$$\|y_d\|_2 < \gamma \|u_{lin}\|_2 \quad \gamma > 0 \quad (4.19)$$

where γ is to be minimised. Note also that ensuring $y_d \in \mathcal{L}_2$ also ensures that $\lim_{t \rightarrow \infty} y_d(t) = 0$ and hence that $\lim_{t \rightarrow \infty} y(t) = y_{lin}(t)$, thereby ensuring that asymptotically at least linear performance is preserved.

Note that because the map $\mathcal{T} : u_{lin} \mapsto y_d$ is inherently nonlinear, to obtain rigorous results, we are forced to turn to nonlinear control tools which are noticeably less “sharp” than linear tools. In particular the \mathcal{L}_2 gain of the system is, by definition, the bound on the “energy gain” of the system for *any* signal in the space \mathcal{L}_2 . In reality a practical system will be subject to only a small subset of these signals. Thus there can sometimes be some difference between an anti-windup compensator which is deemed “optimal” in terms of (4.19) and one that provides the best performance in nonlinear simulation. However, there is often a rough agreement between the two.

There are several approaches to synthesising anti-windup compensators which attempt to minimise the \mathcal{L}_2 gain of the map $\mathcal{T} : u_{lin} \mapsto y_d$. Below, the ones which are used most in the thesis are reviewed briefly.

4.4.3 Full Order Anti-Windup

A so-called “full-order” anti-windup compensator is one in which the number of anti-windup states equals the number of plant states. While, in applications, it is desirable to keep the number of states in the anti-windup compensator to a minimum, full-order anti-windup compensators are of interest because of some of their remarkable features [22, 24].

In the work here, we follow the full order compensator design of [88, 94, 100], where a particularly simple construction is proposed: let $M(s)$ be part of a right coprime factorisation of the plant according to $G_2 = NM^{-1}$. This means then that the disturbance filter $G_2M = N$, and thus by appropriate choice of $M(s)$, to some extent, the undesirable dynamics of the disturbance filter can be cancelled (subject to ensuring stability of the nonlinear loop of course). An appealing method for specifying a coprime factorisation of the plant is the state-space realisation of (4.20) in which F is a free variable akin to a state-feedback gain. This assumes that the plant is given by the minimal state-space realisation of (4.21). This approach is adopted in [88, 94].

$$\begin{bmatrix} M \\ N \end{bmatrix} \sim \left[\begin{array}{c|c} A_p + B_{pu}F & B_{pu} \\ \hline F & I \\ \hline C_p + D_{pu}F & D_{pu} \end{array} \right] \quad (4.20)$$

$$G \sim \left[\begin{array}{c|cc} A_p & B_{pd} & B_{pu} \\ \hline C_p & D_{pd} & D_{pu} \end{array} \right] \quad (4.21)$$

In this case, the search for an optimal compensator is reduced from finding $M(s)$ which minimises $\|\mathcal{T}\|_2$ to finding the state-feedback gain F which minimises $\|\mathcal{T}\|_2$, assuming the state-space realisation of (4.20). This problem was solved in [88, 94] using a Lyapunov formulation of the Circle Criterion and reducing the \mathcal{L}_2 gain synthesis problem into an LMI framework. In this work, additional robustness considerations were incorporated that enable the design to be tuned to improve stability robustness of the nonlinear closed loop system to a class of additive uncertainty.

Let us define the uncertain plant model as $\tilde{G}(s) = [G_1(s) \ G_2(s) + \Delta_G]$. From a small gain argument, it can be shown that the nominal linear system (in the absence of saturation) is robustly stable to an uncertainty model $\|\Delta_G\|_\infty < 1/\gamma$ where $\gamma = \|(I - K_2G_2)^{-1}K_2\|_\infty$. In order to assess stability robustness of the nonlinear system, the de-coupled representation of the uncertain anti-windup closed loop system in Figure 4.17 is considered. Applying a small gain argument it can be shown that the nonlinear closed loop system is robustly stable to the uncertainty model Δ_G if $\|\Delta_G[I - M\mathcal{F}(u_{lin})]\|_2 < 1/\gamma$ where $\mathcal{F}(u_{lin})$ is the nonlinear map $u_{lin} \mapsto \tilde{u}$. Therefore, in order to maximise the size of the plant uncertainty model Δ_G that the nonlinear system is robust to, we desire to minimise $\|[I - M\mathcal{F}(u_{lin})]\|_2$. This can be achieved by solving

$$\|z\|^2 \leq \gamma^2 \|u_{lin}\|^2 \quad (4.22)$$

for minimal γ where $z = M\tilde{u}$. This robustness condition is incorporated together with the \mathcal{L}_2 performance condition to form a weighted combination as shown in (4.23). The matrices $W_p \in \mathbb{R}^{n_y}$ and $W_r \in \mathbb{R}^m$




Figure 4.17: De-coupled AW block diagram with uncertainty model

are positive definite (and normally diagonal) weighting matrices chosen by the designer. For emphasis on performance W_p is chosen large and W_r small; for emphasis on robustness, the reverse is done. Note that when the ratio $\bar{\sigma}[W_p]/\bar{\sigma}[W_r]$ is sufficiently small, the design becomes more “IMC-like” - in which case performance is sacrificed for robustness.

$$\left\| \begin{array}{c} \sqrt{W_p} y_d \\ \sqrt{W_r} z \end{array} \right\|^2 \leq \gamma^2 \|u_{lin}\|^2 \quad (4.23)$$

In [88, 94] it was proved that there exists an anti-windup compensator with $F = LQ^{-1}$ such that (4.23) is satisfied if there exist matrices $Q > 0 \in \mathbb{R}^{n_p}$, diagonal $U > 0 \in \mathbb{R}^m$, an unstructured matrix $L \in \mathbb{R}^{m \times n_p}$ and a scalar $\gamma > 0$ such that the following LMI holds.

$$\begin{bmatrix} QA'_p + A_pQ + L'B'_p + B_pL & B_pU - L' & 0 & QC'_p + L'D'_p & L' \\ * & -2U & I & UD'_p & U \\ * & * & -\gamma I & 0 & -I \\ * & * & * & -\gamma W_p^{-1} & 0 \\ * & * & * & * & -\gamma W_r^{-1} \end{bmatrix} < 0 \quad (4.24)$$

The above LMI provides a computationally convenient manner for synthesising anti-windup compensators and has been used in a number of applications [32, 44, 98] successfully. An alternative method solving the same problem was proposed in [79] in which a Riccati equation was the central computational tool; this is computationally more efficient and also identified strong links between the stability multiplier ($W = U^{-1}$) and the robustness of the saturated system. In this thesis the original LMI approach is used.

Although they require more computational power to implement in practice, the full order AW designs have the following benefits:

- Provided that $G(s) \in \mathcal{RH}_\infty$ there will always exist a globally stabilising full order compensator [24, 100]. This follows since the IMC design, for which closed loop stability is guaranteed, is a special case of the full order design.
- When constructed as part of a right coprime factorisation of the plant $G_2 = NM^{-1}$, the full order designs are *completely independent* of the controller. This is not to say that saturated performance will not be influenced by the nominal controller, but stability of the saturated system and optimality of the compensator design in terms of performance are not affected. Hence, once the compensator is designed, the nominal controller could be retuned without having to revisit the anti-windup design process.
- A full order compensator will achieve the lowest \mathcal{L}_2 gain of any linear compensator. That is, by this performance measure, it yields optimal performance.

The full-order anti-windup technique can suffer from high computational demands since it introduces the same number of states to the control system as the linear plant model used for design. Thus, even if the nominal controller is simple, having only one integral state, the compensator may introduce many more. In the interest of minimising computational demands, there is an interest in designs for which Θ is simply a gain matrix, so-called static anti-windup compensation, and designs for which Θ is transfer function matrix of order less than the plant, so called low-order dynamic anti-windup compensation. Approaches to the static and low-order dynamic anti-windup synthesis problem are presented in the following subsections.

4.4.4 Static Anti-Windup

Static anti-windup compensators are so called because they contain no dynamics; they are activated upon saturation being detected and deactivated immediately upon saturation ceasing. In common with full-order anti-windup compensators, it transpires that computationally attractive algorithms exist for their synthesis, but unlike the full-order compensators they are also computationally attractive for implementation since they are just gains - they require no additional states. The method we follow for the synthesis of optimal static anti-windup compensators was introduced in [91] (see also [65]). The compensator is described by the static relationship

$$\begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} = \Theta \tilde{u} = \begin{bmatrix} \Theta_1 \\ \Theta_2 \end{bmatrix} \tilde{u}$$

where $\Theta_1 \in \mathbb{R}^{m \times m}$ and $\Theta_2 \in \mathbb{R}^{n_y \times m}$. In the Weston and Postlethwaite Framework, the static compensator described above corresponds to the choice of M as shown in (4.25).

$$M = (I - K_2 G_2)^{-1} (K_2 \Theta_2 + \Theta_1 + I) \quad (4.25)$$

A minimal state space realisation of the static compensator in the Weston and Postlethwaite framework is given as (4.26) for which the matrices are defined in Appendix B.1.

$$\begin{bmatrix} M \\ N \end{bmatrix} \sim \left[\begin{array}{c|c} \bar{A} & B_0 + \bar{B}\Theta \\ \hline \bar{C}_1 & D_{01} + \bar{D}_1\Theta \\ \bar{C}_2 & D_{02} + \bar{D}_2\Theta \end{array} \right] \quad (4.26)$$

Given this state-space realisation, it was proved in [91], that a static anti-windup compensator, $\Theta = LU^{-1}$ which ensures that $\|\mathcal{T}\|_2 < \gamma$, can be computed if there exists a positive definite matrix $Q > 0 \in \mathbb{R}^{n_p}$, diagonal $U > 0 \in \mathbb{R}^m$, an unstructured matrix $L \in \mathbb{R}^{m \times n_p}$ and a scalar $\gamma > 0$ which satisfy the following LMI:

$$\begin{bmatrix} Q\bar{A}' + \bar{A}Q & B_0U + \bar{B}L - Q\bar{C}'_1 & 0 & QC'_2 \\ \star & -2U - UD'_{01} - D_{01}U - L'\bar{D}'_1 - \bar{D}_1L & I & UD'_{02} + L'\bar{D}'_2 \\ \star & \star & -\gamma I & 0 \\ \star & \star & \star & -\gamma I \end{bmatrix} < 0 \quad (4.27)$$

Compensator synthesis is achieved using essentially the same LMI tools as in the full order case except that there are some important differences to note

- A static anti-windup compensator ensuring $\|\mathcal{T}\|_{i,2} < \gamma$ is not always guaranteed to exist. Recall, that in order for a compensator to exist, it must satisfy the multivariable Circle Criterion; this is not always possible when Θ is stipulated to be static.
- A static compensator, when it exists, may not achieve “optimal” \mathcal{L}_2 performance; that is $\gamma_{full} \leq \gamma_{static}$. This is essentially because the static compensator has far fewer degrees of freedom with which to tailor the \mathcal{L}_2 performance.
- The barred matrices in the state-space realisation (4.26) are functions of the controller parameters, meaning that a controller re-tune would typically require an anti-windup compensator re-design.

Nevertheless, the simplicity with which static anti-windup can be implemented is the main selling point and it may be considered acceptable to compromise on performance. Furthermore, as mentioned previously, the \mathcal{L}_2 gain is only an approximation of the performance problem and it may well be that even if a static compensator yields a nonlinear closed-loop with a much greater \mathcal{L}_2 gain, in typical simulations it may perform sufficiently well to be considered suitable for implementation.

4.4.5 Low Order Anti-Windup

It transpires that full-order and static anti-windup synthesis can be cast directly as a convex optimisation problem using LMIs. For compensators of intermediate order, the optimisation problem is in general non-convex

and thus difficult to solve. From a practical perspective however, it is often desirable to obtain low order compensators: normally the complete dynamics are not required to furnish satisfactory performance on the saturated closed loop system and thus low-order approximations can be used to obtain very similar performance with much lower computational overheads.

In [91], a method for designing low order anti-windup compensators was proposed. In that work, the anti-windup compensator is partitioned into a set of linear filters and some static gains according to (4.28).

$$\Theta = \begin{bmatrix} \Theta_1(s) \\ \Theta_2(s) \end{bmatrix} = \begin{bmatrix} F_1(s)\tilde{\Theta}_1 \\ F_2(s)\tilde{\Theta}_2 \end{bmatrix} \quad (4.28)$$

The synthesis routine is partly manual and partly automatic in that the designer chooses the filter dynamics $F_1(s)$, $F_2(s)$ and then an LMI optimisation process (almost identical to the static one) is carried out to determine the gain matrices $\tilde{\Theta}_1$ and $\tilde{\Theta}_2$ which minimise the \mathcal{L}_2 gain of \mathcal{T} as before. Clearly, stability will not be guaranteed for an arbitrary choice of filter dynamics and so some care needs to be taken in specifying appropriate filters. A good starting point is to synthesise a full order design and choose $F_1(s)$ and $F_2(s)$ as low order approximations of the resulting $M(s) - I$ and $G_2(s)M(s)$.

For the low order design, the resulting M is again given by (4.25), except that now Θ_1 and Θ_2 are dynamic as defined in (4.28). A minimal state-space description of the compensator is given as (4.29) and adopts a similar form to that of the static design. The state space matrices are given in Appendix B.1.

$$\begin{bmatrix} M \\ N \end{bmatrix} \sim \left[\begin{array}{c|c} \bar{A} & B_0 + \bar{B}\tilde{\Theta} \\ \hline \bar{C}_1 & D_{01} + \bar{D}_1\tilde{\Theta} \\ \bar{C}_2 & D_{02} + \bar{D}_2\tilde{\Theta} \end{array} \right] \quad (4.29)$$

The gain matrix $\tilde{\Theta} = [\tilde{\Theta}_1 \quad \tilde{\Theta}_2]$ is then constructed as $\tilde{\Theta} = LU^{-1}$, where again L and U are solutions of the LMI (4.27) (with the state-space realisation (4.29)). This gain matrix is then combined with the filters $F_1(s)$ and $F_2(s)$ in order to produce the anti-windup compensator. Note that the LMI optimisation process itself has no more degrees of freedom than the static case, but the flexibility afforded by allowing the designer to specify the filter dynamics does lead to a greater range of possible solutions. As a result, although there is no guarantee that a stabilising solution will exist, the likelihood is improved significantly compared to the static case.

The main drawback with the low-order approach is that the designer has to choose the compensator dynamics in order to tune the system. However, the control that this affords the designer seems often to allow the resulting designs to outperform that of the full order approach (from time domain simulations), even though $\gamma_{full} \leq \gamma_{low}$. Further work is required to identify any guidelines or tuning rules that could be used to guide the designer in this choice. As with the static design, a retune of the nominal controller will require a redesign of the anti-windup compensator. However, provided that the change is not dramatic, simply re-running the optimisation using the same filter dynamics and tuning parameters is likely to be sufficient to yield another successful design. Another

similar design to this is proposed in [46] where in addition to the gain matrices being optimised, the zeros of the compensator are optimised as well.

4.5 Summary

A summary of some of the main points of comparison between the various anti-windup approaches discussed on this chapter is shown in Table 4.1. The conditional integration and back calculation are grouped with the ad-hoc integrator reset methods as they share the same properties.

One further point that helps to distinguish between the different approaches concerns the concepts of controller windup and plant windup proposed by Peter Hippe [37]. The term controller windup is used to describe the degradation of performance associated with the controller states evolving effectively in an open loop manner when plant input saturation effectively breaks the feedback loop. The term plant windup is used to describe the degradation of performance associated with saturation that cannot be attributed to the adverse behaviour of controller states during saturation. It is most clearly understood when considering the behaviour of a closed loop system with a pure proportional controller. In this case, no controller windup can take place, but depending upon the dynamics of the open loop plant, saturation may cause oscillatory transient behaviour or even instability of the closed loop. Peter Hippe attributes this plant windup phenomena to the closed loop dynamics being chosen to be too fast and that with the plant input constraints, the plant states cannot be brought to equilibrium quickly enough.

It is the opinion of the author that the overall windup phenomenon cannot be separated clearly into the individual phenomena of controller and plant windup except in the case when one is avoided as in the proportional controller case. However, the concepts are useful to bear in mind. An important point concerning plant windup is that it is a function of the nominal controller and the saturation constraint, so in the two stage anti-windup design process we do not seek to eliminate plant windup.

Most of the *classic* anti-windup designs seek only to solve the controller windup phenomenon. These include all the integrator reset strategies, back calculation and tracking, conventional anti-windup, the observer approach, the observer technique, Hanus conditioning and internal model control. In contrast, the optimal designs of [24, 91], by virtue of their performance optimisation, also take into account the effects of plant windup and seek to limit its effect. This property makes the optimal designs much more appealing for application to generic systems.

AW Design	Simplicity of Design	Simplicity of implementation	Nonlinear Stability Guarantee	Controller Restriction	Plant Restriction	Performance	Existence of stabilising solution
Ad-hoc integrator reset	Very simple	Very simple	None	PI(D) only	None	System Dependent	Not guaranteed
BCAT	Simple	Simple	A priori	PI(D) only	None	System Dependent	Not guaranteed
CAW	Simple	Simple	A priori	None	None	System Dependent	Not guaranteed
Observer Approach	Quite Simple	Simple	A priori	None	None	System Dependent	Not Guaranteed
Hippe Observer Technique	Quite Simple	Simple	A priori	State feedback only	Observable Stable G(s)	System Dependent	Guaranteed
Hanus conditioning	Simple	Simple	A priori	Invertible at $\omega \rightarrow \infty$ Minimum phase		System Dependent	Not guaranteed
IMC	Simple	Additional states as $\deg(G(s))$	Inherent in design	None	Stable G(s)	Requires G(s) to be well damped	Guaranteed
Turner Full-order Dynamic	Quite simple	Additional states as $\deg(G(s))$	Inherent in design	None	Stable G(s)	Optimised	Guaranteed
Optimal Static	Quite simple	Simple	Inherent in design	None	Stable G(s)	Optimised	Not guaranteed
Optimal Low-Order Dynamic	Requires choice of filter dynamics	Few additional states	Inherent in design	None	Stable G(s)	Optimised	Not guaranteed

Table 4.1: Comparison of anti-windup design characteristics

Chapter 5

Anti-Windup Design for EPHS Motor Speed Control

In the control of the PMSM constituent of the EPHS system, a restriction on the current flowing in the motor is required. One way of attempting to enforce this restriction is by constraining the current which is demanded by the PMSM speed controller. In this chapter various anti-windup techniques are applied to the motor speed control system to compensate for violation of this constraint on the current demand. The constraint is a nonlinear function of the elements of the current demand vector and so does not naturally fit into the linear anti-windup frameworks described in Chapter 4. A novel method of transforming this multivariable nonlinear constraint into a time-varying SISO constraint compatible with linear anti-windup synthesis techniques is presented. This method is employed to allow optimal anti-windup designs to be constructed for the EPHS system. Simulation analysis then follows on a nonlinear discrete-time model in order for comparisons between the optimal designs and some more traditional approaches to be made. The most appealing traditional and modern designs are then subjected to practical testing with the real system, yielding encouraging results. This is believed to be the first application of modern optimal anti-windup compensation to PMSM control and the main results of this work were published in [59] and are soon to appear in [60].

Before the application of anti-windup compensation to this complex system is tackled, attention is given to the application of anti-windup to a low order linear continuous time “single axis” model of the system for which both the controller and plant model are SISO systems. This allows some of the features of each design to be observed more transparently prior to application to the complex nonlinear model. Following this, the vector saturation constraint is described, a mathematical model is developed and a method by which it can be incorporated into a linear anti-windup framework is developed. Linearised models of the open loop and closed loop dynamics are developed for the system in the absence of saturation and finally, anti-windup compensators are designed and tested on the complete system model, both in simulation and experimentally.

5.1 Single Axis Speed Controller Anti-Windup Design

We first consider a simplified speed control system in which the “plant” under consideration represents an inner loop motor current control system and the “controller” is a PI speed regulator in the outer loop. The inner loop

current control system has been substantially simplified and is modelled as a first order continuous time system with transfer function $G(s)$ given in (5.1). The dynamics of the d -axis are ignored entirely and the single input represents the q -axis current demand, with the output of the system being the rotational velocity of the motor. As a result of considering only the q -axis, the effects of phase advance and cross-coupling between the d and q axes are omitted. The normalised parameters k_t , j_m and b_m represent the torque constant and inertia of the motor and the mechanical damping of the system. Due to variation in the viscosity of the hydraulic fluid with temperature, b_m will vary from an initial value as high as 0.2 at startup, but then quickly converge toward a nominal value of approximately 0.05. The mechanical damping is time varying in nature but in the analysis it is considered to be a constant as it is understood to vary with significantly slower dynamics than the rest of the model. A discrete PI motor speed controller, K_{dt} , is given by the transfer function of (5.2) and normalised parameters for this and the plant model can be found in Table 5.1. All anti-windup designs will be tested in discrete time but to allow compensator synthesis in continuous time, the equivalent continuous time controller is given as (5.3).

$$G(s) = \frac{k_t}{j_m s + b_m} \quad (5.1)$$

$$K_{dt} \sim \begin{cases} x(k+1) = x(k) + k_i \tau e(k) \\ u(k) = x(k) + k_p e(k) \end{cases} \quad (5.2)$$

$$K_{ct} \sim \begin{cases} \frac{d}{dt} x = k_i e(t) \\ u(t) = x(t) + k_p e(t) \end{cases} \quad (5.3)$$

Parameter	Value	Description
k_t	0.83	Normalised motor torque constant
j_m	0.8	Normalised motor inertia
b_m	[0.05,0.2]	Normalised motor mechanical damping
k_p	0.45	Normalised proportional gain
k_i	0.05	Normalised integral Gain
τ	0.15	Normalised discrete iteration rate

Table 5.1: Single axis speed control model parameters

A bode plot of the loop transfer function, $L = GK$, for each extreme of the mechanical damping variation (Figure 5.1) reveals that the nominal closed loop model has an infinite gain margin and a phase margin in excess of 80 degrees, and therefore is robustly stable. The reduction in loop gain with the higher damping factor suggests that performance of the linear system will be degraded in the cold start condition, but coupled with the reduced phase lag, robust stability actually improves. In addition to the stable closed loop dynamics, the plant is also open loop stable and heavily damped.

Because the model is so simple and the open loop dynamics are relatively benign, one might expect the system to tolerate saturation very well, with the only negative effect being the well known problem of integrator windup. As such, this application is not expected to be a difficult test of the various anti-windup designs, but




Figure 5.1: Bode plot of single axis motor control system

its simplicity will be useful to show the various structures and methods of operation. To mimic variation of the q -axis current limit in the real system, a saturation limit is imposed on the plant input that varies with motor speed according to Table 5.2. Between breakpoints, linear interpolation is used.

motor speed breakpoint (normalised)	0	1	3.8	5
plant input limit, \bar{u} (normalised)	1	1	0.33	0.33

Table 5.2: Plant input saturation limit lookup table

5.1.1 Anti-Windup Designs

Prior to the application of any anti-windup designs it is important to gain an understanding of the behaviour of the unconstrained linear system (without saturation) and the behaviour of the constrained system (saturation limits included) without anti-windup compensation, for which a block diagram is shown in Figure 5.2. The latter describes the expected behaviour of the system if no anti-windup conditioning is applied, and the former is generally considered to represent the *desired* behaviour of the system i.e. that which the conditioned system should preserve as closely as possible. Step responses for both models are plotted in Figure 5.3 for the nominal mechanical damping value of 0.05 and in Figure 5.4 for the cold start value of 0.2. Note that the control system should be tuned for the nominal condition with $b_m = 0.05$ but performance in the cold start condition is also required to be acceptable. The model is scaled such that a normalised speed of 3.5 units represents the maximum operational speed of the motor.




Figure 5.2: Single axis constrained model block diagram

- Unconstrained performance:** With the nominal damping value, the linear unconstrained system exhibits a small overshoot of each speed set point and zero steady-state error. With the higher mechanical damping value the system is more sluggish and although steady-state error is still zero, tracking performance is degraded and the rise time is increased by a factor of 5.
- Constrained performance:** The constrained closed loop system exhibits a very similar response with the nominal mechanical damping value (Figure 5.3) as saturation only occurs for a short period of time at the onset of each step reference. Note that the saturation constraint becomes tighter when the motor speed is higher and so the severity and duration of saturation is increased for the deceleration step. The restriction in available control energy during tracking of the deceleration demand slows tracking initially. However, whilst saturated, the integrator state winds up resulting in a greater overshoot of the set-point. With the high mechanical damping value (Figure 5.4), the plant input level required to achieve the high speed set-point is beyond the saturation limit and so the reference is infeasible. Acceleration performance is similar until the point at which the saturation limit is reached and then the motor speed quickly plateaus. While the infeasible reference persists the integrator continues to accumulate in an attempt to eliminate the steady-state error, driving the controller output further and further into saturation. When the deceleration step demand is applied, there is an instantaneous drop in the controller output as the proportional term in the controller responds to the step change in the error signal. However, the control signal remains above the saturation limit for some time and so the change in reference does not have an immediate effect on the plant. As time progresses, integral action in the controller drives the control signal back into the linear range, at which point the plant output starts to track toward the new reference. However, the rise time is large as the response is dominated by integral action.

Selection of anti-windup designs

A selection of the anti-windup designs presented in Chapter 4 are chosen for application to the model. The back calculation method is chosen as a performance benchmark as it represents one of the most widely used anti-windup approaches used in industry. The BCAT method is also selected as it provides improvements over




Figure 5.3: Single axis model response without anti-windup and nominal mechanical damping




Figure 5.4: Single axis model response without anti-windup and increased mechanical damping

the back calculation approach with minimal additional complexity for design and implementation. The conventional anti-windup (CAW) design is also applied as this also provides the potential of improved performance with no more complexity than the BCAT approach. These are designs that typically may be adopted in industry if warranted by a suitable performance improvement, as their tuning process is simple, intuitive and quick. The other designs chosen for testing are the optimal dynamic and static designs of [91] and [94] as the benefits they provide in terms of guaranteed stability for the nonlinear system and optimised performance may warrant the additional complexity for design and implementation. As the plant in this simple model is only first order, the low order dynamic anti-windup approach does not provide any reduction in the compensator order compared to the full order case and so is ignored. In addition, static compensation was not possible as the LMI solver failed to find a feasible solution, and so only the full order optimal design is included.

Back calculation

Figures 5.5 and 5.6 show the responses of the system with back calculation to the same doublet demand and compares performance against the desired linear behaviour. The back calculation approach requires no tuning and simply recomputes the integral state of the controller when saturation is encountered such that the controller output at the next sample instance lies approximately at the saturation limit. This is observed most clearly at the onset of the deceleration step demand for the model with nominal b_m (Figure 5.5) but it also occurs at the onset of the acceleration step. When a step demand is made, the proportional action of the controller causes a spike in the control signal that may exceed the saturation limit. When this spike does cause a transient violation of the limit, back calculation resulting from this event causes an instantaneous drop in the integrator state. As the motor accelerates and the proportional component of the control signal drops, the controller output drops below the saturation limit and the system fails to make full use of the available control energy. As a consequence, rise times are extended, particularly for the deceleration step since saturation is more severe. It should be noted that this problem can be avoided if the proportional gain is reduced enough to prevent the initial violation of the saturation limit. However, it is an undesirable feature of an anti-windup design to place such a limitation on the tuning of the nominal controller.

With the higher mechanical damping value, back calculation impairs the tracking of the acceleration step demand as the controller output drops below the saturation limit for approximately 8 units of time. While the infeasible reference persists, the integrator state is not allowed to accumulate and the controller output is held very close to the saturation limit. When the deceleration demand is made, tracking is significantly better than for the acceleration demand and even the nominal linear system. This is helped in part by the fact that in the constrained simulation, the motor speed is already closer to the low speed set-point. Also, since the speed error is smaller, the step change in reference does not cause the controller output to exceed the limit and back calculation does not occur. However, another factor affects the tracking of the deceleration step that could be considered a negative characteristic of the anti-windup design.

In a conventional PI regulatory system, at steady-state the tracking error tends to zero and the output of the




Figure 5.5: Single axis model response with Back Calculation and nominal mechanical damping




Figure 5.6: Single axis model response with Back Calculation and increased mechanical damping

controller is influenced predominantly by the integral state since the proportional term tends to zero. When a step change to the reference is applied, a step change in the control signal is also observed that is proportional to the error between the output level and the new reference. This error signal, initially, is also equal to the magnitude of the step change in the reference. When the reference is infeasible at steady-state as in the case we consider here, back calculation constrains the integral state but the proportional term of the controller is unaffected and remains non-zero. As a result, when the reference demand is reduced in a step-wise manner, a step change is observed in the control signal that is proportionate to the step change in reference *demand* even though the change in motor speed required to meet the new set-point may be significantly less. In this application, the effect only serves to quicken the tracking performance but in other applications, this may result in undesirable overshoots.

Back Calculation and Tracking (BCAT)

Back calculation and tracking is applied to the single axis model with discrete control as shown in Figure 5.7 and simulation responses to the same doublet demand are shown in Figures 5.8 and 5.9. The crucial difference between this approach and back calculation *only* is that the integrator state is reset progressively over time through the BCAT time constant, $1/T_t$, rather than instantaneously at the next sample. Note that a factor τ is included which represents the controller discrete sample period to facilitate correct operation in discrete time (as standard with discrete-time integration). Due to the progressive reset of the integrator, when saturation occurs for brief periods, for instance due to proportional kick, the effect on the integral state is only minor and the subsequent loss of performance observed with the back calculation only approach may be avoided. In the example shown here, the compensator is tuned by progressively increasing the magnitude of $1/T_t$ such that when an infeasible reference is applied, the anti-windup feedback is large enough to drive the controller output down close to the saturation limit at steady-state. As a result, the design will drive the integrator state toward similar values as achieved by the Back Calculation only approach and so the design is expected to exhibit similarities with this method.

The simulation results using $b_m = 0.05$ (Figure 5.8) exhibit significantly better tracking performance than the back calculation only approach as expected, as the integral state is not depleted significantly during the brief saturation events. Tracking performance does lag slightly behind that of the linear system but this is expected to some extent as the saturation limits restrict the available control energy. However, part of this loss of performance can be attributed to the anti-windup design and improved performance could be observed with a lower anti-windup gain. In the simulations using $b_m = 0.2$ (Figure 5.9), performance is very similar to that of the back calculation *only* approach with the exception that the rise time to the acceleration demand is improved. As with the back calculation only method, the response to the reverse step demand following an infeasible reference is faster than the linear response. If this is undesirable, the anti-windup gain can be reduced toward the lower limit of K_i/K_p in (4.6) which is equivalent to Hanus conditioning. For comparison, a compensator tune equivalent to Hanus conditioning is included in Figures 5.10 and 5.11. The Hanus method gives a better

match with the linear responses in both simulation conditions but the response to a reverse step following an infeasible reference is slower (Compare Figure 5.11 with 5.9). By fine tuning the anti-windup feedback gain a chosen compromise between the behaviour of each design can be achieved.




Figure 5.7: Single axis model back calculation and tracking discrete-time implementation

Conventional Anti-Windup

Conventional Anti-Windup is quite similar to the back calculation and tracking approach except for the fact that the anti-windup feedback path influences the *whole* controller via the error signal, as shown in Figure 5.12, not just the integrator. Responses of the CAW system are shown in Figures 5.13 and 5.14.

The ability to constrain the integrator state during extended periods of saturation is dependent upon the anti-windup feedback gain being quite high, hence its common description as *High Gain* anti-windup. With PI controllers, the presence of a direct feed-through (proportional) term can result in instability if the anti-windup feedback gain is set too high, particularly for discrete-time implementations. This is made more likely still since a unit delay required in the anti-windup feedback path to prevent an algebraic loop introduces additional phase lag. In this application, the anti-windup gain is tuned as aggressively as possible without causing instability in order to extract maximum performance. Note that with the given anti-windup tune, some chattering is observed in the control response during saturation (See Figure 5.13) although deviation from linear performance for the simulation with nominal mechanical damping is very small. This design is not susceptible to performance degradation when large but brief saturation events occur as the proportional and integral behaviour of the controller are treated together. However, it can be sensitive to conditions in which saturation continues for extended periods.

Because of the restriction on the anti-windup gain to retain stability, an amount of integrator windup is unavoidable when an infeasible reference is demanded as shown in Figure 5.14. The integrator windup that results leads to a sluggish response to the reverse step which would be worse still were the infeasible reference to persist for a longer period. However, some simple analysis that follows shows that the integrator state would actually be bounded.




Figure 5.8: Single axis model response with Back Calculation & Tracking and nominal mechanical damping




Figure 5.9: Single axis model response with Back Calculation & Tracking and increased mechanical damping




Figure 5.10: Single axis model response with Hanus conditioning




Figure 5.11: Single axis model response with Hanus conditioning




Figure 5.13: Single axis model response with Conventional Anti-Windup and nominal mechanical damping




Figure 5.14: Single axis model response with Conventional Anti-Windup and increased mechanical damping

magnitude of the robustness weight W_r (also a scalar) with respect to W_p . Starting from a value of 1, the magnitude of the robustness weight is increased to bias the optimisation toward maximising performance and a value of 100 was found to provide good results. The resulting design achieves an \mathcal{L}_2 gain bound $\gamma = 5.25$, and the compensator state space model (4.20) is parameterised by (5.8) giving the continuous time state space model of (5.9). The frequency domain behaviour of this compensator is shown in Figure 5.16. Note that the block diagram and state space models are given in continuous time as the design is synthesised in this domain, but conversion to discrete-time in both cases is trivial.




Figure 5.15: Single axis full order anti-windup block diagram

Single axis full order design 1:

$$F = -0.1094 \quad (5.8)$$

$$\begin{bmatrix} M(s) - I \\ G(s)M(s) \end{bmatrix} \sim \left[\begin{array}{c|c} -0.1719 & 1 \\ -0.1094 & 0 \\ \hline 1.037 & 0 \end{array} \right] \quad (5.9)$$

With the first design, the response to acceleration demands is very good with rise times that are among the fastest and very comparable to the linear system response. Whilst the infeasible reference is present (Figure 5.18), the compensator acts to constrain the controller output, preventing the integrator from accumulating excessive energy. Tracking of the deceleration demand is much improved over the case without anti-windup but cannot match that of the back calculation methods. A crucial difference in the design of this compensator is that it does not seek to minimise the extent of saturation as the other designs do. Instead, the design brief for this compensator type is to minimise the deviation from *linear performance*, i.e. the performance of the system without saturation. This is dependent upon the linear model used in synthesis and so for any system with varying plant dynamics, the trim point chosen can have a significant effect on the behaviour of the resulting design. This is illustrated by comparing this design with a second design achieving an \mathcal{L}_2 gain performance bound of $\gamma = 3.16$ defined by (5.10) and (5.11), based on the linear plant model with $b_m = 0.2$, $W_p = 1$ and $W_r = 100$. The frequency domain behaviour of this compensator design is also included in Figure 5.16. Note that with this design, channel one of the compensator has very little effect and the anti-windup action is

affected almost entirely by feedback to the controller input as with Conventional Anti-Windup. In spite of this, performance exceeds that of the CAW design. Part of the reason for this is that the frequency shaping present in the second compensator channel allows the compensator to constrain the integrator state more effectively at steady-state without causing stability problems at higher frequencies.

Single axis full order design 2:

$$F = -3.12e - 6 \quad (5.10)$$

$$\begin{bmatrix} M(s) - I \\ G(s)M(s) \end{bmatrix} \sim \left[\begin{array}{c|c} -0.25 & 1 \\ \hline -3.12e - 6 & 0 \\ 1.037 & 0 \end{array} \right] \quad (5.11)$$

Simulation results for this second compensator design are shown in Figures 5.19 and 5.20. When comparing the performance of the two designs on the simulation model with $b_m = 0.05$ (Figures 5.17 and 5.19), little difference is observed because the period of saturation is quite short. However, for the simulation model with $b_m = 0.2$, there is a marked difference: Design 2, which is designed around this condition, exhibits a much faster return to linear behaviour, taking approximately 10 units of time (Figure 5.20) compared to 40 for Design 1 (Figure 5.18). The explanation for this is simply that since Design 1 is based on the model dynamics with $b_m = 0.05$, it aims to provide performance akin to the linear case with $b_m = 0.05$ even when the system properties deviate from those of the design model.

As the system is expected to operate with nominal mechanical damping for the majority of the time, Design 1 would be the obvious choice provided that it is sufficiently robust, in terms of stability and performance, to the known variation in plant dynamics. If this was not the case, for instance if the performance with $b_m = 0.2$ was deemed to be undesirable, it may be appropriate to alter the dynamics of the design model to compensate. As an example, a better performance compromise between the two simulation cases may be obtained if the design model was parameterised with $b_m = 0.1$ i.e. somewhere between nominal and cold. Another observation regarding the two designs is that tracking of the reverse step demand for the case where $b_m = 0.2$ is faster with Design 1 and so in one sense, overall performance of the system may be perceived to be better even though it deviates more from linear performance.

5.1.2 Performance Comparison

The output tracking responses for each of the anti-windup designs described are included together in Figures 5.21 and 5.22 to aid comparisons between the different approaches.

Simulations with $b_m = 0.05$

With the nominal mechanical damping value it is clear that the back calculation approach is the worst performing design. The remaining designs all perform quite similarly, partly because saturation is not severe and is quite short-lived in the simulation. The system with conventional anti-windup exhibits the smallest deviation from linear performance, arguably followed by the simulation without anti-windup and the full-order dynamic




Figure 5.16: Bode plots of full-order AW designs 1 and 2




Figure 5.17: Single axis model response with Full Order Dynamic Anti-Windup and nominal mechanical damping. Anti-windup design 1: Based on linear plant model with $b_m = 0.05$




Figure 5.18: Single axis model response with Full Order Dynamic Anti-Windup and increased mechanical damping. Anti-windup design 1: Based on linear plant model with $b_m = 0.05$




Figure 5.19: Single axis model response with Full Order Dynamic Anti-Windup and nominal mechanical damping. Anti-windup design 2: Based on linear plant model with $b_m = 0.2$




Figure 5.20: Single axis model response with Full Order Dynamic Anti-Windup and increased mechanical damping. Anti-windup design 2: Based on linear plant model with $b_m = 0.2$

design; the former overshooting a little and the latter undershooting a little. The Back Calculation and Tracking design suffers a larger undershoot but as shown previously, this could be re-tuned to improve performance.

Simulations with $b_m = 0.2$

For the simulations replicating a cold start with higher mechanical damping, the extended period of saturation introduces a much greater spread of performance between the acceleration and deceleration tests. Only the back calculation approach shows a noticeably different response for the acceleration step and this is because the integrator is reset unnecessarily.

In response to the deceleration step, the system without anti-windup performs worst followed by that with Conventional Anti-Windup as the tracking is slower than the linear simulation. All of the other designs aid the system to meet the next setpoint quicker than the linear case. If we are to consider deviation from linear performance, CAW and the full order dynamic design would be considered best, followed by the two back calculation methods. However, if overall system performance is considered, we might consider that the two back calculation designs perform best, followed by full order dynamic AW and CAW. Note that for other systems, this deviation from linear performance may not be considered to be beneficial.

Summary

This section has studied a simple model of the PMSM system under consideration. Although the model is simple, it enables a rough illustration of some of the saturation problems and also a rough idea of the suitability of the various anti-windup architectures. Both the CAW and Back Calculation anti-windup designs have been seen to perform relatively poorly with this model, the former due to its inability to constrain the integrator state adequately during saturation and the latter due to its excessive manipulation of the integrator state when saturation is caused very briefly at the onset of a large change in the reference. Consequently both of these techniques will be discounted for use on the more complex model. The Back Calculation and Tracking, and the optimal dynamic anti-windup designs all performed well and are consequently chosen as candidates for application to the more complex model.

The complex model, described in Chapter 3, is of significantly greater complexity than the simple model, with more states and nonlinearities in both plant and controller, making it a more challenging design problem. The remaining sections of the chapter discuss how the more promising anti-windup techniques can be adapted to this model and reports extensive simulation and experimental results.

5.2 Current Demand Constraint Implementation

One of the crucial differences between the simplified model discussed above and the real system is that, in the simplified model, the d -axis is simply ignored and the constraints on the d - q -axis current demand (control signal) are severely approximate. A more realistic model of the PMSM speed control system is depicted in Figure 5.23, where the constraint on the motor current demand is a limit on the norm of the vector according to (5.12). This can also be expressed as (5.13) and is represented in the block diagram by $\text{sat}(\|\cdot\|)$.




Figure 5.21: Performance comparison of anti-windup designs for single axis model with nominal mechanical damping




Figure 5.22: Performance comparison of anti-windup designs for single axis model with high mechanical damping




Figure 5.23: Closed loop control system

$$\left\| \begin{bmatrix} i_{d,dmd} \\ i_{q,dmd} \end{bmatrix} \right\| \leq i_{max} \quad (5.12)$$

$$\sqrt{i_{d,dmd}^2 + i_{q,dmd}^2} \leq i_{max} \quad (5.13)$$

It is obvious that this multivariable constraint is not consistent with standard anti-windup design, both because it is not decentralised and because the constraint is not applied at the plant input. In order to fit an anti-windup framework where the saturating signals are inputs to the plant, we consider the inner loop current control system to be the effective plant, and the outer loop PI regulator and phase advance algorithm to constitute the controller. Thus, the current demand saturation constraint is considered to be a plant input limitation. The complexity of this saturation constraint is that the limit is a nonlinear function of the plant inputs. It may appear simple to apply independent constraints on each element of the current demand vector. In principle this is possible and modern anti-windup methods such as found in [24, 84] are certainly capable of dealing with multivariable saturation constraints. However, for this application there are a number of complications with this. Firstly, there is not a unique solution to (5.13) for a given limit i_{max} .

Let us define limits on the d and q axis components of the current vector $i_{d,lim}, i_{q,lim}$ such that

$$\sqrt{i_{d,lim}^2 + i_{q,lim}^2} = i_{max} \quad (5.14)$$

holds. Any current demand vector $[i_{d,dmd} \ i_{q,dmd}]'$ for which $|i_{d,dmd}| \leq i_{d,lim}$ and $|i_{q,dmd}| \leq i_{q,lim}$ will also satisfy (5.13). However, there are a continuum of possible solutions $i_{d,lim}, i_{q,lim}$ satisfying (5.14) and it is not trivial to select a set of constraints that are appropriate. This is shown in Figure 5.24 where we consider a possible current demand vector $i_{dq,dmd}$ with norm at the limit, i_{max} .

The set of constraints $[i_{d,lim} \ i_{q,lim}]' = [i_{d2} \ i_{q2}]'$ satisfy (5.14) but are inappropriate for the given current angle since the limit on $i_{q,dmd}$ would already be violated even though $|i_{dq,dmd}| \not\geq i_{max}$. Such a set of constraints would cause the compensator to activate prematurely and seek to drive the motor current vector toward $\tilde{i}_{dq,dmd}$. An appropriate set of constraints for the given current demand vector is $[i_{d,lim} \ i_{q,lim}]' = [i_{d1} \ i_{q1}]'$. Given

this example it is simple to see that the limits on each element of $i_{dq,dmd}$ would be required to vary in real time in accordance with the current angle. This is entirely possible but would require significant online computation.




Figure 5.24: Possible elementwise current limit models

The second complication relates to the action of the controller being asymmetric in that its behaviour during deceleration is significantly different to that during acceleration. This will be explained with reference to Figure 5.25 which shows the structure of the speed regulator and phase advance function a little more explicitly. During normal operation, the q -axis current demand can adopt both positive and negative values, producing accelerating or decelerating torque. However, the d -axis demand is always negative or zero whether accelerating or decelerating as $i_{d,dmd}$ is defined to be proportional to $-|i_{q,dmd}|$. When the speed controller is linearised the nonlinear modulus term is lost and thus when $i_{q,dmd}$ becomes negative, the linear model produces $i_{d,dmd} \geq 0$ which is incorrect. This discrepancy is highlighted by considering the expected behaviour of the nonlinear and linear models in acceleration and deceleration conditions as summarised in Table 5.3.




Figure 5.25: Nonlinear speed controller

The error identified in the linear model would prevent any synthesis routine based on linear systems theory from generating a compensator design that would deal effectively with both d and q axes of the system. This can be

Operating condition	Linear model	Nonlinear model
Acceleration	$\dot{\omega}_m > 0$	$\dot{\omega}_m > 0$
	$\omega_m > 0$	$\omega_m > 0$
	$i_{q,dmd} > 0$	$i_{q,dmd} > 0$
	$i_{d,dmd} < 0$	$i_{d,dmd} < 0$
Deceleration	$\dot{\omega}_m < 0$	$\dot{\omega}_m < 0$
	$\omega_m > 0$	$\omega_m > 0$
	$i_{q,dmd} < 0$	$i_{q,dmd} < 0$
	$i_{d,dmd} > 0$	$i_{d,dmd} < 0$

Table 5.3: Possible operating conditions and their representation by the linear and nonlinear models

seen by considering how a successful anti-windup compensator would act during acceleration and deceleration.


In the acceleration condition the linear and nonlinear models broadly agree, and a successful anti-windup compensator would seek to make $i_{q,dmd}$ less positive and drive $i_{d,dmd}$ in the positive direction in order to reduce the magnitude of the current demand vector. In the deceleration condition, the linear and nonlinear models fundamentally disagree and the linear model predicts $i_{q,dmd} < 0$ and $i_{d,dmd} > 0$. A successful anti-windup compensator for the linear model would seek to make $i_{q,dmd}$ less negative and also drive $i_{d,dmd}$ in the negative direction in order to reduce the magnitude of the current demand vector. When such a compensator is applied to the nonlinear system, it will operate as desired during acceleration, but during deceleration it will drive the d -axis current demand in the negative direction. Since, this is already negative, this will actually increase the magnitude of the d -axis element of the vector and will either slow the departure from saturation or even drive the system further into saturation.

SISO treatment of norm saturation

To overcome the problems described in the previous section with respect to the d -axis portion of the model we convert the norm limit to a time-varying constraint exclusively on the q -axis current. The threshold of this constraint, $i_{q,lim}$, is determined as a function of the d -axis current demand according to (5.15) where i_{max} is the maximum allowable magnitude of the current vector and $i_{q,lim}$ solves (5.16). A graphical depiction of this limit is shown in Figure 5.26.

$$i_{q,lim} = \pm \sqrt{i_{max}^2 - i_{d,dmd}^2} \quad (5.15)$$

$$\left\| \begin{bmatrix} i_{d,dmd} \\ i_{q,lim} \end{bmatrix} \right\| = i_{max} \quad (5.16)$$


 Figure 5.26: q -axis current limit model

This saturation limit, although parameter varying, can still be described as a sector bounded nonlinearity to fit in with the compensator synthesis frameworks of [91, 88, 24]. This is shown in Figure 5.27 where the following two saturation functions with different saturation limits are shown to be enclosed by the sector bound $[0,1]$.

$$sat_1(u) = \begin{cases} \bar{u}_1 & \forall u \geq \bar{u}_1 \\ u & \forall -\bar{u}_1 < u < \bar{u}_1 \\ -\bar{u}_1 & \forall u \leq -\bar{u}_1 \end{cases} \quad sat_2(u) = \begin{cases} \bar{u}_2 & \forall u \geq \bar{u}_2 \\ u & \forall -\bar{u}_2 < u < \bar{u}_2 \\ -\bar{u}_2 & \forall u \leq -\bar{u}_2 \end{cases}$$




Figure 5.27: Sector bounded nonlinearities

A limitation with this approach is that $i_{q,dmd} \leq i_{q,lim}$ will only imply $\|i_{dq,dmd}\| \leq i_{max}$ if $|i_{d,dmd}| \leq i_{max}$. For $|i_{d,dmd}| > i_{max}$, the solutions to (5.15) are complex, which is of course not possible in practice, so a more appropriate description of the constraint is given in (5.17).

$$i_{q,lim} = \begin{cases} \pm \sqrt{i_{max}^2 - i_{d,dmd}^2}, & |i_d| \leq i_{max} \\ 0, & |i_d| > i_{max} \end{cases} \quad (5.17)$$

One may have the concern that a compensator driven by violation of this scalar saturation limit will affect the q -axis current magnitude only, leaving the d -axis current unchanged, and altering the current angle from that of i_{dq} to that of \tilde{i}_{dq} (Figure 5.26). However, for the controller we consider, the current angle is determined exclusively by the phase advance map and so any reduction made to the q -axis current demand will also have a direct and proportionate effect on the d -axis demand. Thus, we are able to handle a nonlinear multivariable constraint effectively using a single time-varying scalar constraint.

5.3 Models for Anti-Windup Design

All of the optimal anti-windup design techniques are model-based; that is they require linear models of the plant, and also the controller in the low-order and static cases, in order to be synthesised. Furthermore such linear models enable other anti-windup designs, such as the BCAT techniques, to be analysed in a more rigorous way. Unfortunately, both the controller and the “plant” (inner loop current control system) depicted in Figure 5.25 are both nonlinear and hence to apply the anti-windup analysis and synthesis techniques discussed, need to be linearised.

Due to the modelling of the current constraints described above, we concentrate only on the q -axis behaviour of the speed controller. The behaviour of the nonlinear d -axis portion of the controller is far less important to consider for anti-windup design since it does not contain any dynamic elements and its main function is simply to alter the trim point at which the motor current control loop operates. In addition, the phase advance angle is scheduled with motor speed which varies with significantly slower dynamics than the electrical system in which saturation occurs. Therefore, we choose to consider only the q -axis portion of the controller in the model that will subsequently be used for compensator synthesis. The d -axis current demand, along with the mechanical load torque are considered to be disturbance inputs to the linear plant model.

To generate a linear plant model, the d – q axis motor model equations (3.10) are linearised using the equilibrium (trim) conditions $\omega_{m,lin}$, $i_{d,lin}$, $i_{q,lin}$, with corresponding equilibrium inputs $V_{d,lin}$, $V_{q,lin}$ and $load_{lin}$, yielding the linear state-space model of (5.18). The dynamics of the linear model are dependent on the trim conditions, most notably the motor speed. Thus, the speed at which the model is linearised is an important consideration in developing a representative linear model. It should also be noted that, as the motor has several inputs, there is not a unique trim point at a given speed; a wide range (a continuum) of voltage vectors could result in the same speed at steady-state. The trim condition is dependent upon the current angle and the applied load as well as the chosen speed. To ensure that the linear models produced represent the true behaviour of the model at a given speed and load, the trim conditions were determined from closed loop simulation results at steady-state, as for a given motor speed and applied load, the phase advance map in place enforces a unique voltage vector.

$$\begin{aligned}
\begin{bmatrix} \dot{i}_d \\ \dot{i}_q \\ \dot{\omega}_m \end{bmatrix} &= \begin{bmatrix} -\frac{R_s}{L_s} & P\omega_{m,lin} & Pi_{q,lin} \\ -P\omega_{m,lin} & -\frac{R_s}{L_s} & -Pi_{d,lin} - \frac{P\psi_f}{L_s} \\ 0 & \frac{3P\psi_f}{2J} & -\frac{B}{J} \end{bmatrix} \begin{bmatrix} i_d \\ i_q \\ \omega_m \end{bmatrix} + \begin{bmatrix} \frac{1}{L_s} & 0 \\ 0 & \frac{1}{L_s} \\ 0 & 0 \end{bmatrix} \begin{bmatrix} V_d \\ V_q \end{bmatrix} \\
\begin{bmatrix} \omega_m \\ i_d \\ i_q \end{bmatrix} &= \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} i_d \\ i_q \\ \omega_m \end{bmatrix}
\end{aligned} \tag{5.18}$$

The linear model that results is then combined with the PI regulators of (5.19) to form a linear model of the closed loop current control system with d and q axis current demand inputs. As the d -axis demand is considered to be a disturbance signal, this input is ignored, and by considering only the motor speed output, ω_m , we are left with the SISO model from $i_{q,dmd}$ to ω_m of (5.20). Here, states one and two represent the d and q -axis currents respectively, state three represents the motor speed and states four and five correspond to the d and q -axis integrator states in the current loop.

$$\begin{aligned}
\dot{x}_{crt} &= \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_d \\ x_q \end{bmatrix} + \begin{bmatrix} k_{i,crt} & 0 & -k_{i,crt} & 0 \\ 0 & k_{i,crt} & 0 & -k_{i,crt} \end{bmatrix} \begin{bmatrix} i_{d,dmd} \\ i_{q,dmd} \\ i_d \\ i_q \end{bmatrix} \\
\begin{bmatrix} V_d \\ V_q \end{bmatrix} &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x_d \\ x_q \end{bmatrix} + \begin{bmatrix} k_{p,crt} & 0 & -k_{p,crt} & 0 \\ 0 & k_{p,crt} & 0 & -k_{p,crt} \end{bmatrix} \begin{bmatrix} i_{d,dmd} \\ i_{q,dmd} \\ i_d \\ i_q \end{bmatrix}
\end{aligned} \tag{5.19}$$

$$\begin{aligned}
\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \\ \dot{x}_4 \\ \dot{x}_5 \end{bmatrix} &= \begin{bmatrix} -\frac{P+k_{p,crt}}{L} & P\omega_{m,lin} & Pi_{q,lin} & \frac{1}{L} & 0 \\ -P\omega_{m,lin} & -\frac{R+k_{p,crt}}{L} & -Pi_{d,lin} - \frac{P\psi_f}{L} & 0 & \frac{1}{L} \\ 0 & \frac{3P\psi_f}{2J} & -\frac{B}{J} & 0 & 0 \\ -k_{i,crt} & 0 & 0 & 0 & 0 \\ 0 & -k_{i,crt} & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix} + \begin{bmatrix} 0 \\ \frac{k_{p,crt}}{L} \\ 0 \\ 0 \\ k_{i,crt} \end{bmatrix} i_{q,dmd} \\
\omega_m &= \begin{bmatrix} 0 & 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \end{bmatrix}
\end{aligned} \tag{5.20}$$

It was noted in (3.10) of Chapter 3.4 that the motor model is a Linear Parameter Varying (LPV) system. This potentially opens the door to the use of LPV anti-windup techniques such as found in [4, 17] in which the anti-windup compensator is scheduled with the varying dynamics of the system. Such an approach may offer

improved performance and greater robust stability properties. However, the possibility of achieving this in the PMSM speed control loop is complicated somewhat by the fact that the LPV model is present within the feedback loop of the current control system, and additional parameter variation is brought about by the phase advance function. For these reasons, LPV anti-windup techniques were not tested although in principle it may be possible to derive an LPV approximation of the plant model dynamics in order to apply these techniques.

5.3.1 Linear Model Dynamics

A linear model produced at a speed of 1/3 normalised units, given by the parameterised state space model of (5.21) was found to be a good choice for compensator design. The choice of linearisation speed was made based on the conditions for which current demand saturation was considered to be most problematic. To understand this better let us consider the required operation of the control system and the conditions for saturation to occur.

$$\begin{aligned} \dot{x} &= \begin{bmatrix} -1463 & 1047 & 24.15 & 22730 & 0 \\ -1047 & -1463 & -477.1 & 0 & 22730 \\ 0 & 168.3 & -3.882 & 0 & 0 \\ -9.535 & 0 & 0 & 0 & 0 \\ 0 & -9.535 & 0 & 0 & 0 \end{bmatrix} x + \begin{bmatrix} 0 \\ 1109 \\ 0 \\ 0 \\ 9.535 \end{bmatrix} \dot{i}_{q,dmd} \\ \omega_m &= \begin{bmatrix} 0 & 0 & 1 & 0 & 0 \end{bmatrix} x \end{aligned} \quad (5.21)$$

The speed control system is required to operate at an idle speed and accelerate quickly to step inputs. The magnitude of these steps is variable but can extend to the maximum speed capability of the motor. For current demand saturation to occur, quite a large change in the speed demand is required, for instance, accelerating from an idle speed of 1/6 units to 2/3 units or more. In such an example, during the time the compensator is active, the motor will accelerate toward the speed reference and so the dynamics of the plant will vary. The linear model dynamics should be a good representation of the model throughout this speed range. The mid-point between idle speed of 1/6 units and maximum speed of 1 may be a good compromise. However, when load disturbances are applied the speed range will be limited and saturation will occur at lower speeds. In addition, saturation is likely to be most severe at the onset of the largest step demand i.e. when the speed is low. For these reasons a linearisation speed in the lower third of the operational speed range is considered to be a good compromise. Should the compensator remain active for a significant period following the high speed set-points being reached, it may be necessary to increase the linearisation speed to enable the compensator behaviour to better match the system in this region.

Figure 5.28 shows the variation in dynamics of the linear current loop model from q -axis current demand to motor speed. Initially as the speed increases from 0 there is no change in dynamics, then at 1/6 units where phase advance comes into operation variation in the dynamics start to appear. As the linearisation speed is increased further, gain in the 1 to 100 Rad/s region increases and beyond 100 Rad/s, the gain drops off more

abruptly. In terms of phase, increased motor speed provides reduced phase lag below 2 Rad/s but significantly increased phase lag in the 200 to 500 Rad/s region. Based on this analysis, if a linearisation point were to be chosen to best represent the dynamics over the whole speed range, approximately 2/3 units would be a good choice.




Figure 5.28: Variation of effective plant dynamics with motor speed under no load conditions

A time domain comparison of the nonlinear model and the linear model in response to step inputs at the q -axis current demand is shown in Figure 5.29. This comparison shows that the design model dynamics are a good representation of the nonlinear system behaviour up to a speed of at least 0.6 units.




Figure 5.29: Time domain comparison of linear and nonlinear models without saturation

5.4 Anti-Windup Designs for the Complex Model

In this section we present a number of different anti-windup compensators designed to deal with saturation of the current demand vector. Tuning of the model based designs is aided by continuous time simulations based on the linear model descriptions presented in Section 5.3. Following a successful first pass, the designs are discretised and tested with a multi-rate discrete time model incorporating the nonlinear phase advance controller and a full three phase description of the motor and drive system. The designs were iterated upon and re-simulated on the nonlinear model in order to maximise performance and provide sufficient robustness to cope with the variation in dynamics that is present. For the ad-hoc methods, tuning was performed directly on the nonlinear model.

Before the anti-windup designs are presented we consider first the behaviour of the nonlinear model without anti-windup and also that of the model in the absence of the saturation constraints. These are important first steps to gain understanding of the intended performance and also the extent of performance degradation caused by saturation.

5.4.1 Nonlinear Model Behaviour without Saturation

The response of the nonlinear model to a series of step demands under no load conditions without any saturation constraints applied is shown in Figure 5.30. The plots are normalised such that the maximum operational speed of the motor, the maximum available torque, and the maximum magnitude of the current demand vector are




Figure 5.30: Nonlinear model step responses - no saturation

equal to one. The magnitude of each step demand is $1/6$ units. Due to the changing dynamics of the system with motor speed and increased coupling between d and q axes at higher speeds, damping is observed to deteriorate with increasing speed. In addition, the existing tune of the control system causes very large current demand transients for such step demands in the high speed range; in the region of 2.5 times that of the current limit. When larger step references are applied such as $2/6$ or more in normalised units, the nonlinear model exhibits instability as shown in Figure 5.31. In addition larger steps at lower speeds such as from $1/6$ units to $5/6$ units also introduce instability. There are a number of possible causes of / contributors to this instability but for a full consideration this discussion is deferred to a later section. For now, the key points to note are that the system dynamics vary with motor speed and also with the size of the step reference. Of course, being a nonlinear system, the type of input also has a significant effect on the response of the system but for the purposes of this thesis, abrupt repeatable inputs such as steps are required to generate the saturation phenomenon we desire to investigate. To gain a better understanding of the system, it is useful to employ some linear analysis tools to assess the properties of the nonlinear model linearised at various equilibrium points within the operating range of the system.

5.4.2 Linear Stability Analysis

In order to perform any linear analysis, the d - q axis model is used rather than the three phase model as this can be linearised readily. The limitation of this approach is that it does not consider the operation of the PWM inverter which introduces a phase lag for discrete implementation as described in Chapter 3.9 and a contributor




Figure 5.31: Nonlinear model unstable step response - no saturation

of phase lag is effectively ignored. Based on the PWM inverter model employed in the EPHS three phase model, the phase lag is proportional to motor speed. Therefore, when assessing the linear analysis results, this must be taken into consideration.

The EPHS model was linearised as described in Section 5.3 under a range of load conditions and operational speeds to allow analysis of the stability properties of the speed loop. The speed and load range of interest was quantised and stability properties were tested at each feasible combination. In Figure 5.32 the Nyquist plots for each linearisation condition are overlaid. From visual inspection it is apparent that the gain and phase margins are of a healthy size. Figures 5.33 and 5.34 show the gain and phase margins as a function of the motor speed and applied load. For motor speeds below $1/6$ units where field weakening operation is inactive, there is very little variation in the closed loop dynamics as a function of motor speed and applied load, and the gain and phase margins are large. The application of a load torque during flux-weakening operation increases the gain margin but decreases the phase margin. The minimum gain margin occurs at maximum speed under no load and is large at 14.3 in linear units. The conditions in which the phase margin is at a minimum are in the speed region of $1/4$ under maximum load and also at maximum speed, giving values of approximately 57 degrees.

Based on this analysis, there are no conditions in which the closed loop d - q axis model would be expected to be unstable in the absence of saturation. It was thought initially that the instability observed in the simulation of Figure 5.31 may be due to the increased phase lag present in the three phase model. However, this was ruled out as simulation tests using the $d - q$ axis plant model also revealed the same instability problem.

Nyquist plots for the speed control system OLTF (SISO approximation) across the speed and load range




Figure 5.32: Nyquist plot for $d-q$ axis model speed loop




Figure 5.33: Discrete $d-q$ axis model speed loop gain margin variation




Figure 5.34: Discrete $d-q$ axis model speed loop phase margin variation

For the $d - q$ axis model without saturation constraints although instability could be induced, it was found for a smaller set of conditions than with the three phase model. Possible explanations for this difference between the predictions of stability and observed instability in simulation are as follows:

- When the motor accelerates very quickly, the system operates away from any of the linear model trim conditions and therefore the linear models fail to represent the system behaviour accurately.
- The phase advance map is linearised as two gains that are dependent upon the trim condition and so in reality these gains are time-varying. This time variation is not taken into consideration in the linear analysis and as such, we are effectively making use of the Aizermann Conjecture [45]. Aizermann hypothesised that for a Lure type nonlinear system with nonlinear gain bound by $[\alpha, \beta]$, the nonlinear closed loop would be stable if the linear closed loop formed by replacing the nonlinear gain by a static gain γ is stable for all $\alpha \leq \gamma \leq \beta$. This was shown to be untrue and hence could be a reason why the linear analysis is optimistic.

5.4.3 Nonlinear Model Behaviour without Anti-Windup

A simulation response for the nonlinear model with saturation constraints applied is given in Figure 5.35. The exogenous inputs used for this simulation are the same as applied in the simulation response of Figure 5.30 which shows the response of the system without saturation. Thus, comparison of these two figures allows

the effect of saturation on the system to be seen. The bottom plot in Figure 5.35 shows the boolean result of $\|\tilde{u}\| > \epsilon$ where \tilde{u} is the difference between the constrained and unconstrained q -axis current demand and ϵ is a small positive number. For the first three step demands, the saturation limit is not exceeded and the response is identical to the case without the saturation constraints in place. For the subsequent step demands, the rate at which the motor accelerates reduces as the current the control system is able to inject into the q -axis reduces. In spite of the saturation event, overshoot of the reference is of a similar level to that in the simulation without the saturation constraints and the oscillations that follow are of lower magnitude. Ordinarily one may expect the saturated behaviour to be worse, but in this case the reduced motor speed gradient as the set-point is reached, caused by the restriction on control energy, makes the system less susceptible to such oscillations. In addition, the effective plant is stable and well damped (Section 5.3) and so the temporary break in the feedback loop caused by saturation is not catastrophic to performance. Additional responses with higher magnitude step demands are shown in Figures 5.36, 5.37 and 5.38. For these cases the effects of saturation are more problematic and generally lead to more pronounced overshoots and longer settling times. With the acceleration demand in Figure 5.38 the motor is accelerated from its idle speed up to the maximum operational speed of the motor. Although this represents the largest step we may expect to observe in the real system, the response initially is very good. It is only after the initial rise that the system exhibits strange behaviour and windup causes the motor speed to exceed the set point. A critical point to note is that the presence of the saturation constraints actually stabilises the system as many of these inputs caused instability in the simulations without saturation. Some theoretical justifications for this phenomena have been given by Teel [83] but the complexity of the model used here prevents an analysis based on those results.

The effect of static loads on the system response is shown in Figures 5.39 and 5.40. The load is calibrated such that when the maximum speed is demanded as in Figure 5.40, the speed attained at steady-state is reduced by approximately 10%. The presence of this load increases the rise time to the feasible 4/6 unit reference of Figure 5.39 compared to the no load case in 5.37 but the percentage overshoot and settling time are very similar. When the load renders the reference demand infeasible as in Figure 5.40, the system is locked into saturation and following some oscillations in the speed response, the response to the reverse step demand is delayed and sluggish.

5.4.4 Anti-Windup Test Conditions

Analysis of the nonlinear system responses without anti-windup reveals that in order to test the anti-windup designs we wish to apply to the system, a number of reference demands and load conditions must be considered. The following are a small selection of conditions that capture the main characteristics of the nonlinear system and shall be used for subsequent anti-windup tests in both simulation and experiment. These focus on applying step changes to the speed reference that are large enough to induce problematic saturation and consider the worst case applied loads. Note that the behaviour during deceleration is not considered. This is because for practical application, deceleration demands are rate limited to prevent the motor acting as a generator and




Figure 5.35: Nonlinear model 1/6 unit step responses - no AW




Figure 5.36: Nonlinear model 3/6 unit step responses - no AW




Figure 5.37: Nonlinear model idle to 4/6 unit doublet responses under no load - no AW




Figure 5.38: Nonlinear model idle to max. speed doublet - no AW




Figure 5.39: Nonlinear model idle to 4/6 unit doublet response under static load - no AW




Figure 5.40: Nonlinear model idle to max speed doublet response under static load - no AW

damaging the power electronic drive circuit. This alleviates the problems of saturation during deceleration as will be shown in the experimental tests later in the Chapter. The presence of a static load improves the damping of the system and so the no load condition is generally considered to be the condition of greater interest for anti-windup testing. However, the behaviour of the system when an applied load renders the reference infeasible also leads to undesirable behaviour and so is also considered. The following three simulation conditions are selected to allow performance comparisons between the varying anti-windup strategies.

- **Condition 1:** Step demand from idle speed to maximum speed under no load
- **Condition 2:** Step demand from idle speed to maximum speed with a static load applied such that the speed attained is reduced by 10% from the reference level
- **Condition 3:** Step demand from idle speed to 4/6 units under no load

5.4.5 Anti-Windup Designs

In this section the compensator designs developed for the nonlinear model are presented and the specifics of their tuning detailed.

Back Calculation and Tracking

The back calculation and tracking design method described in Chapter 4.3.4 is followed, with the anti-windup feedback gain increased progressively, starting at the reciprocal of the speed controller proportional gain (equivalent to Hanus' Conditioning). The gain was tuned to obtain the desired trade-off between constraining the integrator state during prolonged saturation, and limiting the loss of performance when severe but brief saturation is caused. The resulting design for discrete time implementation as in Figure 5.7 is defined as

$$1/T_t = 0.38 \quad (5.22)$$

Full order dynamic anti-windup

The current control system model of (5.21) is selected for anti-windup design for the reasons given in Section 5.3, along with the controller model of (3.12). Since the nominal linear closed loop system is SISO, tuning of the full order dynamic anti-windup compensator is quite simple, requiring only that a robustness weight be tuned, and the linearisation point of the current control system be chosen. In terms of linearising the plant, the key parameters that define the trim point are the motor speed, the phase advance angle, and the applied load. As the phase advance angle is fixed as a function of motor speed, this is a given. Any applied load only serves to increase the damping of the plant model, leading to less problematic behaviour and so the no load condition is

selected for anti-windup design, leaving only the motor speed trim point to be chosen. Although decreasing the magnitude of the robustness weight leads to higher levels of \mathcal{L}_2 performance, we wish to maximise robustness in order for the design to be capable of maintaining closed loop stability for operation away from the linearisation point and so the robustness weight is increased as much as possible without sacrificing performance. The resulting design derived using the formulae of Turner et al. [88, 94] was achieved using $W_r = 20$. This was discretised with a time step of 0.001 normalised units using the Tustin approximation, yielding the following discrete model.

$$x(k+1) = \begin{bmatrix} 0.1433 & 0.0265 & -0.3818 & 0.6207 & -8.3890 \\ -0.0390 & 0.0411 & -5.9200 & -0.0140 & -130.10 \\ -0.0047 & 0.0496 & 0.5808 & 0.0005 & -9.1240 \\ 0.0001 & 0.0000 & 0.0000 & 0.9999 & 0.0007 \\ 0.0000 & -0.0001 & -0.0007 & 0.0000 & 0.9809 \end{bmatrix} x(k) + \begin{bmatrix} 0.0472 \\ 0.7315 \\ 0.0513 \\ 0.0000 \\ 0.0001 \end{bmatrix} \tilde{u}(k)$$

$$\theta(k) = \begin{bmatrix} -0.0558 & 0.5006 & -7.5870 & 0.0674 & -178.80 \\ 0.0000 & 0.0000 & 1.0000 & 0.0000 & 0.0000 \end{bmatrix} x(k) + \begin{bmatrix} 0.0000 \\ 0.0000 \end{bmatrix} \tilde{u}(k)$$

$$\|\mathcal{T}\|_2 = 43.3595$$

Low order dynamic anti-windup

Due to the high order of the current control system model, the full order dynamic compensator produced has five states. This is undesirable, particularly considering that the whole of the nominal control system has only three states. In addition, high order implementations on the target microprocessor are particularly undesirable as the limitation to fixed point arithmetic requires that each input, state, multiplicand and output signal be assigned scalings to ensure that quantisation errors do not cause a significant deviation from the intended behaviour using floating point arithmetic. This is trivial for first order filters but quite complex beyond second order filters. As a result, the low order dynamic anti-windup compensation scheme of [91] described in Section 4.4.5 is also employed. Static compensation would be even more desirable in terms of computational simplicity but in this case the LMI constraints for its design were found to be infeasible. In the low order method, providing that the choice of filter dynamics allow the stability and performance conditions in the LMI to be satisfied, a dynamic design consisting of two independent first order filters can be produced.

The same plant and controller model used for the full order dynamic anti-windup synthesis were selected for design. The choice of filter dynamics was guided initially by approximating the dominant characteristics of the successful full order design. Thus, given a plant linearisation condition and performance weight equal to 1 as in the full order case, the only other tuning parameter required was the robustness weight. Initially this was set to a small value so as not to restrict the optimisation and the attention was given to the design of filter

dynamics that provide a feasible solution to the LMI optimisation. A successful design was found in which both filters were chosen as first order low pass designs. From this feasible solution, the bandwidths of these filters were varied and simulation tests performed in order to observe the effect on system behaviour. It was found that the bandwidth of filter $F_1(s)$ had to be restricted as for stability of the nonlinear loop, the gain at high frequencies had to be limited. If the bandwidth was too low, the compensator had little effect via θ_1 and the design behaved very much like Conventional Anti-Windup, except less effectively. A higher bandwidth was permissible for feedback via θ_2 and allowed better control over the integral state. Note that the Nyquist frequency of the speed loop must be considered during the design stage as the speed of the filter dynamics are bound by this. However, in this case the design was not restricted due to the controller sample rate. Fine tuning the compensator bandwidths and the robustness weight resulted in the following design which was discretised at a time step of 0.001 normalised units using the tustin approximation.

$$\Theta_1(z) \sim \begin{cases} x(k+1) = 0.4274x(k) + 0.7018\tilde{u}(k) \\ \theta_1(k) = -0.6224x(k) - 0.3112\tilde{u}(k) \end{cases} \quad (5.23)$$

$$\Theta_2(z) \sim \begin{cases} x(k+1) = 0.6667x(k) + 0.8333\tilde{u}(k) \\ \theta_2(k) = 0.2206x(k) + 0.1103\tilde{u}(k) \end{cases} \quad (5.24)$$

$$\|T\|_2 = 43.3597$$

5.4.6 Simulation Analysis

The three compensator designs above, tuned for use with the nonlinear model, are compared against the Back Calculation only approach in Figures 5.41-5.49 where the three test conditions defined in Section 5.4.4 are applied. All three designs provide improved performance over the simulations without anti-windup and simulations with back calculation only.

The plots for Simulation Condition 1 (Figures 5.41-5.43) reveal that as soon as saturation ceases the BCAT design ceases to act and nominal control behaviour is resumed. However, the dynamic compensators remain active for a period of time as the energy stored in their states dissipates, and the full-order design is observed to remain active for a significantly longer period than the low-order design. All designs are very successful and allow the speed response to converge quickly to the set-point when saturation ceases. A greater performance difference is observed in the response to the reverse step, and although this condition will not occur in practice due to rate limiting of deceleration demands, it does suggest that the BCAT and low-order designs may be more desirable than the full-order compensator.

The plots for Simulation Condition 2 (Figures 5.44-5.46) reveal that saturation persists for the entire duration of the high speed reference demand. Therefore, performance is equivalent until the reference or loading changes. Whilst the control signal is saturated the plant output evolves in a manner determined mainly by the plant

dynamics and no difference between the three systems is observed. Whilst in saturation, the behaviour of the compensators may not have a direct influence on the plant output, but by piloting the states of the controller and the controller output, the time at which saturation ends can be influenced and the behaviour of the control signal immediately following the escape from saturation can be influenced. When saturation ceases, the differences between the differing anti-windup designs are observed.

The plots for Simulation Condition 3 (Figures 5.47-5.49) reveal a more useful comparison. Saturation ceases quite quickly after the onset of the step reference, allowing the anti-windup compensators to have a greater influence on system performance. In terms of overshoot of the high speed reference, the BCAT design performs best, followed by the low-order and full-order designs. The three candidates provide very similar levels of performance and so the differences are subtle. A comparison of the output responses for the three designs and simulation conditions is given in Figure 5.50 allows the differences in overall performance to be seen more clearly. The main points to note between the different designs are as follows:

- Performance levels are very similar for the three candidates
- The low-order design consistently matches or exceeds the performance of the other designs




Figure 5.41: Condition 1 nonlinear simulation with BCAT




Figure 5.42: Condition 1 nonlinear simulation with Full-order AW




Figure 5.43: Condition 1 nonlinear simulation with Low-order AW




Figure 5.44: Condition 2 nonlinear simulation with BCAT




Figure 5.45: Condition 2 nonlinear simulation with Full-order AW




Figure 5.46: Condition 2 nonlinear simulation with Low-order AW




Figure 5.47: Condition 3 nonlinear simulation with BCAT




Figure 5.48: Condition 3 nonlinear simulation with Full-order AW




Figure 5.49: Condition 3 nonlinear simulation with Low-order AW




Figure 5.50: Speed response comparison: Nonlinear model simulations with BCAT, low-order AW and Full-order AW

5.4.7 Perturbed Anti-Windup Designs

The designs presented here represent the best performance attainable by each of the three anti-windup approaches on the nonlinear model. Due to inevitable differences between models and real systems there is a further stage of fine-tuning when the designs are applied to the real system for experimental tests. With simple control designs such as PI controllers, or simple anti-windup designs such as back calculation and tracking, minor modifications can be applied ‘online’ simply by altering gains, allowing the effects of perturbing the tune from the nominal design to be observed directly. With more complex controllers such as \mathcal{H}^∞ designs, or the optimal compensators we consider, the design cannot be fine-tuned manually. Instead, a useful approach that can be employed is to generate a number of ‘perturbed’ designs by altering tuning parameters a small amount in one direction and another. Then, the effect of such perturbations on performance in both simulation and practical implementation can be seen. If the model is accurate enough to be useful for tuning, the best performing design on the simulation model should also be the best performing design on the real system. In addition, further confidence can be gained if the performance of the perturbed design differs from that of the nominal design in the same manner in both simulation and experiment. Perturbed designs for both the BCAT and low-order designs are presented in the following sub-sections. Note that the full-order design is considered to be too computationally demanding for implementation and so is not chosen for experimental tests. Also note that because the agreement between the model and experimental responses is not high, quite coarse perturbations are applied.

Perturbed BCAT designs

The following back calculation & tracking designs will be tested on the practical system, for which design 1 is the nominal design and the anti-windup feedback gain is decreased and increased from nominal in designs 2 and 3 respectively.

$$\text{BCAT design 1: } 1/T_t = 0.38\tilde{u}$$

$$\text{BCAT design 2: } 1/T_t = 0.28\tilde{u}$$

$$\text{BCAT design 3: } 1/T_t = 0.48\tilde{u}$$

Perturbed low-order designs

The following low-order anti-windup designs will be tested on the practical system, for which design 1 is the nominal design introduced in Section 5.4.5, and in designs 2 and 3, the bandwidth of the filter in channel 1 was increased and reduced respectively. The amount by which this bandwidth was altered was selected to cause a noticeable difference in simulation performance in order for a correlation with experimental behaviour to be made. Bode plots of the resulting compensator transfer functions are shown in Figure 5.54. Note that the

low frequency gain of channel one is maintained at approximately one but that altering the bandwidth in this channel significantly influences the gain in channel two returned by the optimisation.

Low-order design 1:

$$\begin{bmatrix} \Theta_1(z) \\ \Theta_2(z) \end{bmatrix} \sim \begin{cases} x(k+1) = \begin{bmatrix} 0.4035 & 0 \\ 0 & 0.6667 \end{bmatrix} x(k) + \begin{bmatrix} 0.7018 \\ 0.8333 \end{bmatrix} \tilde{u}(k) \\ \begin{bmatrix} \theta_1(k) \\ \theta_2(k) \end{bmatrix} = \begin{bmatrix} -0.6224 & 0 \\ 0 & 0.2206 \end{bmatrix} x(k) + \begin{bmatrix} -0.3112 \\ 0.1103 \end{bmatrix} \tilde{u}(k) \end{cases}$$

$$\|\mathcal{T}\|_2 \leq 43.3597$$

Low-order design 2:

$$\begin{bmatrix} \Theta_1(z) \\ \Theta_2(z) \end{bmatrix} \sim \begin{cases} x(k+1) = \begin{bmatrix} 0.8182 & 0 \\ 0 & 0.6667 \end{bmatrix} x(k) + \begin{bmatrix} 0.9091 \\ 0.8333 \end{bmatrix} \tilde{u}(k) \\ \begin{bmatrix} \theta_1(k) \\ \theta_2(k) \end{bmatrix} = \begin{bmatrix} -0.1989 & 0 \\ 0 & 0.3711 \end{bmatrix} x(k) + \begin{bmatrix} -0.0995 \\ 0.1856 \end{bmatrix} \tilde{u}(k) \end{cases}$$

$$\|\mathcal{T}\|_2 \leq 43.3669$$

Low-order design 3:

$$\begin{bmatrix} \Theta_1(z) \\ \Theta_2(z) \end{bmatrix} \sim \begin{cases} x(k+1) = \begin{bmatrix} 0.25 & 0 \\ 0 & 0.6667 \end{bmatrix} x(k) + \begin{bmatrix} 0.625 \\ 0.8333 \end{bmatrix} \tilde{u}(k) \\ \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} = \begin{bmatrix} -0.7593 & 0 \\ 0 & 0.0797 \end{bmatrix} x(k) + \begin{bmatrix} -0.3796 \\ 0.0398 \end{bmatrix} \tilde{u}(k) \end{cases}$$

$$\|\mathcal{T}\|_2 \leq 43.3588$$

Perturbed design performance analysis

Nonlinear simulation results for the perturbed designs are shown in Figures 5.52 and 5.53 for Simulation Conditions 1 and 3. As found previously, Simulation Condition 1 does not enable the subtle differences between the different designs to be observed and the responses from Simulation Condition 3 found in Figure 5.53 are much more revealing. Here we see that the more aggressive BCAT design 3 exhibits less pronounced overshoot and design 2, that is less aggressive than the nominal design, exhibits more overshoot. As expected, the effect of this single gain on anti-windup behaviour is quite transparent. For the low-order compensator, designs 2 and 3 both exhibit degraded performance compared to the nominal design. Noting that the tuning parameter selected was shifted in opposite directions from nominal to generate these two designs, it shows that tuning is less transparent.




Figure 5.51: Frequency response plots of low-order compensator designs 1-3




Figure 5.52: Nonlinear simulation responses of perturbed anti-windup compensator designs for Simulation Condition 1




Figure 5.53: Nonlinear simulation responses of perturbed anti-windup compensator designs for Simulation Condition 3

5.5 Experimental Validation

The low-order anti-windup and back calculation and tracking anti-windup designs were selected for experimental testing on an EPHS test rig. Each design was coded in C using TargetLink rapid prototyping software and incorporated into a prototype EPHS software release in order for it to run on the target 16 bit fixed-point processor. The following sections give details of the tests and comparisons of anti-windup performance.

5.5.1 Experimental Setup

The PMSM shaft is attached directly to a rotary pump used to direct hydraulic fluid through a sequence of hoses. The main flow path splits into two parallel hoses, one with a rotary valve flow restriction and one with a stop tap. These valves are used to mimic the behaviour of the rotary valve on the steering column of a vehicle equipped with EPHS. Down stream of these obstructions the fluid flow meets again and enters the return path at the rotary pump. In the pump, as with all hydraulic power steering setups there is a pressure relief valve that allows hydraulic fluid to bypass the main hydraulic circuit when the pressure within the pump exceeds a given limit.

The electronic control unit (ECU) contains a 16 bit fixed point arithmetic processor on which the control software runs, and the power electronic drive stage for the PMSM. CANape by Vector Software running on a local PC is used to apply motor speed reference demands to the ECU and also to record test signals. Communication between CANape and the ECU is achieved using the CAN protocol [5]. This is a serial communications protocol using a two wire bus that has become widespread in automotive systems, partly due to its low cost. Bandwidth limitations of the CAN protocol mean that a limited amount of data can be transmitted across the bus. As a result, the frequency at which signals are logged by CANape in the following experimental results are generally lower than the frequency at which they are updated within the ECU and the number of signals to be logged simultaneously is also restricted. For commercial applications this is not a restriction as the requirements for communication over CAN are quite limited. As an example, a vehicle stability control system may only require to receive measurements of the vehicle speed and a steering angle measurement. However, for the purposes of these tests, the bandwidth limitation restricted the number of signals that could be measured at a fast sample rate to four. For the simple designs such as BCAT, $i_{d,dmd}$, $i_{q,dmd}$, ω_m , $\omega_{m,dmd}$ and \tilde{u} were measured of which $\omega_{m,dmd}$ is a reference signal generated within CANape and the anti-windup signal is reconstructed as the BCAT gain multiplied by \tilde{u} with a single speed loop sample delay applied.

For the low order dynamic compensator, the two anti-windup signals also needed to be logged. Reducing the sample rate allowed the extra signals to be logged but the lower rate was not adequate to represent the fast electrical dynamics. As a “work around”, signals were logged at the higher rate in two separate experimental runs and the presence of the speed demand signal in both sets of measurements allowed them to be combined afterwards. In run 1, the speed and current signals were recorded, and in run 2, the anti-windup signals were recorded along with the speed reference. Although for the plots produced the anti-windup signals do not match

the speed response exactly, the repeatability of the experiments was very good so the character of the responses are genuine. This method is undesirable but unavoidable.

A further complication arises due to the serial nature of the CAN interface. Only one message can be sent at any one point in time and this includes reference demands sent to the ECU, measurement signals requested from the ECU and any additional diagnostic communications that may exist between the ECU and the vehicle computer, or in this case, the PC-based simulator. CAN communication is arbitrated by priority and precedence and the reference signals applied and measurement signals requested may not be of the highest priority. Therefore they may also conflict with each other and this is evident in some of the plots presented here. For instance, when the results for different anti-windup designs are overlaid and aligned by the speed reference, some systems respond a little later than the others, indicating that reception of the reference by the ECU was delayed in some cases.

5.5.2 Experimental Results - Perturbed Designs

First let us assess the behaviour of the perturbed anti-windup designs. To allow comparisons with the simulation results, experimental tests were performed with a reference demand sequence equal to that of Simulation Condition 1, and with both hydraulic valves fully open.

BCAT perturbed designs

Time histories for the speed response of the system with BCAT designs 1-3 are shown in Figure 5.54. The first observation is that unlike the simulation results for this condition in Figure 5.52, the motor speed accelerates almost linearly up to and beyond the reference demand. The behaviour of the real system is better represented by the behaviour of the model in Simulation Condition 3, shown in Figure 5.53.

Comparing the behaviour of the three designs, it is clear that design 2 exhibits increased overshoot and design 3 exhibits reduced overshoot compared to the nominal design. This trend matches that observed in the simulation analysis. The discrepancy between Simulation Condition 1 and the equivalent experimental test is thought to be a limitation of the model and will be discussed in a later section.

Low-order perturbed designs

Time histories for the speed response of the system with Low-order designs 1-3 are shown in Figure 5.55. Correlation with the simulation responses for condition 1 shown in Figure 5.52 is poor. However, when compared to the model behaviour in Simulation Condition 3, the same characteristics are observed. Design 1 performs best, and designs 2 and 3 exhibit greater overshoots and worse damping. Furthermore, as with the simulation predictions, design 3 performs the worst and the nominal design performs similarly to the nominal BCAT design.




Figure 5.54: BCAT designs 1-3: Experimental responses under no load




Figure 5.55: Low order designs 1-3 (based on linearisation speed 1/3 units): Experimental responses under no load

The correlation exercise reveals a deficiency in the nonlinear model toward the upper end of the speed range. However, the match between experimental results and the model behaviour at slightly lower speeds suggests that the fundamental dynamics are correct. This is supported by the fact that the nominal designs perform best both in simulation and in practice, and also that the manner in which the performance of each perturbed design degrades from the nominal is equivalent between simulation and experiment.

5.5.3 Experimental Performance Comparisons

Comparisons between the BCAT and low-order anti-windup designs with the experimental equivalents of Simulation Conditions 1-3 are shown in Figures 5.56-5.58. Also included in these figures are experimental results for an in-house form of Back Calculation and the constrained system without anti-windup conditioning.

- **Condition 1:** Without anti-windup compensation, the high speed reference of Simulation Condition 1 is overshoot by more than 20 percent (Figure 5.56). Each anti-windup design successfully improves performance and reduces the overshoot to less than 10 percent. There is not a significant difference in performance between each design and even the in-house Back Calculation method performs well.
- **Condition 2:** Figure 5.57 shows the system responses equivalent to Simulation Condition 2 where a flow restriction in the hydraulic hose restricts the maximum speed of the motor to approximately 10% less than the magnitude of the high speed demand. As in the simulation analysis, saturation persists for the whole duration of the high speed reference. During this time, the anti-windup designs have no direct impact on the output response as the control signal remains in saturation. However, the compensation can influence the length of time spent in saturation and such a difference can be seen between the compensated responses and that without anti-windup.

When the reference demand drops, saturation ceases at the same point in time for all anti-windup compensated responses and the deceleration ramp is properly tracked, indicating a swift return to linear behaviour. For the system without anti-windup, windup causes saturation to cease later and the response to the deceleration ramp is delayed. The delay is minor in this case because the controller output was only lightly saturated at the onset of the deceleration demand, indicated by \tilde{u} being small. Had the deceleration demand been applied at 0.9 units of time just prior to an overflow of the integrator state, the effect would be significantly worse.

- **Condition 3:** Figure 5.58 shows the system responses to the acceleration demand of Condition 3. Here the back calculation approach offers no improvement at all over the case without anti-windup and the set-point is overshoot by 24%. The BCAT and low-order designs perform similarly although damping is marginally better with low-order dynamic compensation, and overshoot is reduced to 14%.

In Figure 5.59 the response of the system to the acceleration demand of Condition 3 with a hydraulic load applied is shown. In this case, the test without anti-windup exhibited a similar overshoot of 23%.

Interestingly, the back calculation approach does improve performance in this case and overshoot is reduced to 18%. Again, the low-order and BCAT designs perform best and provide similar levels of performance. The low-order design does appear to provide better damping than the BCAT design and overshoot of the reference is slightly less at just under 11% rather than 13.5%.

5.6 Discussion

The practical implementation exercises presented in the previous section reveal that the low-order dynamic anti-windup and back calculation and tracking approaches are both successful and appealing methods of conditioning an EPHS speed control system to cope with current constraints. The back calculation and tracking approach does not afford any guarantees of stability for the nonlinear system yet for the system it was applied to, yielded very good performance with very little complexity in both design and implementation. Only one tuning parameter is required and due to its simplicity, tuning is quite transparent, allowing the design to be tuned online.

The low-order dynamic anti-windup design has some significant benefits over the more traditional designs commonly used. One of these is that provided that the designer chooses the filter dynamics appropriately, a guarantee of global stability for the nonlinear system is provided at the synthesis stage. Furthermore, to aid its tuning, the design approach is partly automatic as the synthesis routine seeks out an optimal gain matrix to place in series with the chosen filter dynamics such that performance is optimised in an \mathcal{L}_2 sense. In this application the performance of the low-order design is not a large improvement over that of the simpler BCAT design but for more complex systems, the benefits can be more significant. The design also proves to be robust as good performance and stability are observed away from the trim condition used for its design.




Figure 5.56: Experimental responses - Condition 1




Figure 5.57: Experimental responses - Condition 2




Figure 5.58: Experimental responses - Condition 3




Figure 5.59: Experimental responses - As Condition 3 but with static load applied

5.6.1 Robustness

Low-order compensator stability and performance robustness

Stability of the nonlinear system with low-order anti-windup for operation away from the trim point used for design can be assessed by using an LMI optimisation similar to that used for compensator synthesis. For low-order compensator synthesis, recall that the plant, controller and compensator models are given by (5.25), (5.26) and (5.27) respectively. For compensator synthesis we sought to satisfy (5.28) for minimal \mathcal{L}_2 gain, γ . The matrix inequality of (5.29) follows by substituting in for $\dot{\bar{x}}$, y_d and $u = u_{lin} - u_d$, where $\bar{D} = D_{02} + \bar{D}_2\tilde{\Theta}$.

To check stability and \mathcal{L}_2 performance of a pre-designed low-order compensator, we desire to find a solution to this inequality for a given perturbed plant model and compensator design parameterised by matrix $\tilde{\Theta}$ and continuous time filters $F_1(s) \sim (A_1, B_1, C_1, D_1)$ and $F_2(s) \sim (A_2, B_2, C_2, D_2)$. To this end, the LMI of (5.30) in variables $Q > 0$, diagonal matrix $U > 0$ and scalar $\gamma > 0$ is solved for minimal γ . This LMI is obtained by pre and post-multiplying (5.29) by $[Q \ U \ I]' = [P^{-1} \ W^{-1} \ I]'$ and applying the Schur Complement to remove quadratic and bilinear terms.

$$G \sim \left[\begin{array}{c|cc} A_p & B_{pd} & B_{pu} \\ \hline C_p & D_{pd} & D_{pu} \end{array} \right] \quad (5.25)$$

$$K \sim \left[\begin{array}{c|cc} A_c & B_{cr} & B_c \\ \hline C_c & D_{cr} & D_c \end{array} \right] \quad (5.26)$$

$$\begin{bmatrix} M \\ N \end{bmatrix} \sim \left[\begin{array}{c|cc} \bar{A} & B_0 + \bar{B}\tilde{\Theta} \\ \hline \bar{C}_1 & D_{01} + \bar{D}_1\tilde{\Theta} \\ \bar{C}_2 & D_{02} + \bar{D}_2\tilde{\Theta} \end{array} \right] \quad (5.27)$$

$$\dot{\bar{x}}'P\bar{x} + x'P\dot{\bar{x}} + y_d'y_d - \gamma^2 u_{lin}'u_{lin} + 2\tilde{u}'W(u - \tilde{u}) < 0 \quad (5.28)$$

$$\begin{bmatrix} \bar{x} \\ \tilde{u} \\ u_{lin} \end{bmatrix}' \begin{bmatrix} \bar{A}'P + P\bar{A} + \bar{C}_2'\bar{C}_2 & P(B_0 + \bar{B}\tilde{\Theta}) - W\bar{C}_1' + \bar{C}_2'\bar{D} & 0 \\ * & -2W + \bar{D}'\bar{D} - W(D_{01} + \bar{D}_1\tilde{\Theta}) - (D_{01} + \bar{D}_1\tilde{\Theta})'W & W \\ * & * & -\gamma^2 I \end{bmatrix} \begin{bmatrix} \bar{x} \\ \tilde{u} \\ u_{lin} \end{bmatrix} < 0 \quad (5.29)$$

$$\begin{bmatrix} Q\bar{A}' + \bar{A}Q & (B_0 + \bar{B}\tilde{\Theta})U - Q\bar{C}_1' & 0 & Q\bar{C}_2' \\ * & -2U - (D_{01} + \bar{D}_1\tilde{\Theta})U - U(D_{01} + \bar{D}_1\tilde{\Theta})' & I & U(D_{02} + \bar{D}_2\tilde{\Theta})' \\ * & * & -\gamma I & 0 \\ * & * & * & -\gamma \end{bmatrix} < 0 \quad (5.30)$$

The LMI of (5.30) is employed for the full range of operating conditions as used in the linear stability analysis of 5.4.2 with low-order compensator design 1. The existence of a solution to the LMI indicates that the nonlinear anti-windup compensated system is globally finite gain \mathcal{L}_2 stable at the chosen trim condition. Furthermore, the \mathcal{L}_2 gain bound, γ achieved is an indicator of performance at that condition. Figure 5.60 depicts the feasible operating points for the system and identifies the conditions for which the LMI is satisfied by the intersection of dark horizontal and vertical lines in the mesh. This result reveals that a guarantee of nonlinear stability is given for the system compensated by low-order design 1 for the majority of the region of operation, that the conditions for which stability cannot be guaranteed by the Circle Criterion are beyond a motor speed of 0.8 units and that stability can be guaranteed at slightly higher speeds for reduced loads. Although the Circle Criterion does not ensure stability for operation close to maximum speed, because this region of uncertain system behaviour is very small, it is not impractical for empirical stability tests to be employed.

In Figure 5.61 the \mathcal{L}_2 gain bound achieved is plotted against speed and load. The \mathcal{L}_2 gain bound appears constant for the majority of guaranteed stable conditions but adjacent to the high speed conditions for which the LMI was infeasible (and, hence, for which the \mathcal{L}_2 gain bound is infinity), the \mathcal{L}_2 gain was significantly larger. Although only isolated peaks are observed here, if the number of points within the matrix of speed and load conditions was increased sufficiently, a gradual rise in \mathcal{L}_2 gain would be observed as speed and load increased toward the conditions for which the LMI became infeasible. Away from these boundary conditions, the \mathcal{L}_2 gain is essentially static. This is a positive result as it means that \mathcal{L}_2 performance is largely independent of the operating condition, and hence the low-order compensator exhibits both robust stability *and* performance.

BCAT stability robustness

For the purposes of analysis, the back calculation and tracking design can be considered as a special case of the static AW approach of [91]. To achieve this, the nominal controller and plant models used for analysis are altered from that of the standard structure to (5.31) and (5.32). Note that the measured plant outputs are duplicated, allowing one to be used for the proportional term of the controller, and one to drive the integrator, that can be modified by the anti-windup feedback.

$$K(s) \sim \left[\begin{array}{c|c} A_c & B_{cr} \\ \hline C_c & D_{cr} \end{array} \left[\begin{array}{cc} 0 & B_c \\ \hline D_c & 0 \end{array} \right] \right] \quad (5.31)$$

$$G(s) \sim \left[\begin{array}{c|c} A_p & B_p \\ \hline C_p & D_p \\ \hline C_p & D_p \end{array} \right] \quad (5.32)$$

A static compensator for this system, $\Theta = [\Theta'_1 \ \Theta'_2]'$ is equivalent to BCAT if $\Theta_1 = 0$ and $\Theta_2 = [0 \ 1/(k_{i,spd} \times T_t)]'$ where $k_{i,spd}$ is the controller integral gain. The integral gain term is required because for this implementation, the anti-windup signal enters at the input to the controller rather than directly into the integrator state as




Figure 5.60: Feasible operating conditions (mesh vertices) and guaranteed stable conditions for the nonlinear system with Low-order Design 1 (blue intersections)




Figure 5.61: \mathcal{L}_2 gain variation for all guaranteed stable conditions with low-order design 1

with the conventional implementation of BCAT. Using the same tools as for the low-order anti-windup robustness analysis, the LMI of (5.33) can be derived, which if feasible, provides a guarantee of stability for the BCAT compensated closed loop system. Note that due to the modified controller and plant model implementation, the constant matrices in the LMI have a different parameterisation, given in Appendix B.4.

$$\begin{bmatrix} Q\bar{A}' + \bar{A}Q & (B_0 + \bar{B}\Theta)U - Q\bar{C}'_1 & 0 & Q\bar{C}'_2 \\ * & -2U - (D_{01} + \bar{D}_1\Theta)U - U(D_{01} + \bar{D}_1\Theta)' & I & U(D_{02} + \bar{D}_2\Theta)' \\ * & * & -\gamma I & 0 \\ * & * & * & -\gamma I \end{bmatrix} < 0 \quad (5.33)$$

By requiring two elements of the anti-windup matrix, Θ , to be zero, the structure of the static anti-windup compensator is restricted. In order for quadratic stability to be guaranteed the 2×2 upper left hand block of (5.33) must be negative definite and if this is satisfied, the remaining part of the inequality allows the \mathcal{L}_2 gain to be calculated. Thus, if there exists a $Q > 0$ and a diagonal matrix $U > 0$ such that this upper left block is negative definite, the BCAT compensated system at the given trim point is guaranteed to be stable.

However, for this system and BCAT compensator, the upper left block of the LMI was always found to be infeasible for all operating conditions and so no guarantee of stability for the BCAT compensated system could be given, and consequently, no \mathcal{L}_2 gain could be given either.

5.6.2 Model accuracy

With Simulation Condition 1 the model failed to exhibit overshoot (Figure 5.43) and the rise time was significantly longer than for the equivalent experimental test. This characteristic was only observed at the upper end of the speed range when a large step change in the reference demand was applied. There are three main areas of the model that could be in error to cause this discrepancy; the PWM inverter model, modelling of the hydraulic load, and the parameters of the mechanical and electrical system model. As the PWM inverter model is a known over simplification, this was an obvious first place to look. In reality, this is a complex nonlinear system which if implemented fully within a Simulink environment would lead to a significant increase in simulation time. Alternative software such as Saber is typically used for such models as this is designed specifically to represent the complex behaviour of such electronic and control circuits. An accurate model of the inverter was not available during the research, although some useful information about the effect of the simple model used can be obtained by comparing the behaviour of the $d-q$ axis and 3Φ models.

- **PWM inverter model:** The $d-q$ axis model is equivalent to the three phase model with an ideal inverter in place. In comparison, with the discrete 3Φ model, as the phase voltages are only updated at discrete intervals in time, motion of the motor causes the voltage vector to effectively move into phase lag in between samples compared to its demand. This effect is accentuated at higher speeds. Comparing the




Figure 5.62: Simulation condition 1, comparison between d - q axis and 3Φ model speed responses




Figure 5.63: Simulation condition 2, comparison between d - q axis and 3Φ model speed responses

$d - q$ axis and 3Φ models in open and closed loop provides some insight about the appropriateness of the PWM inverter in the 3Φ model.

Closed loop simulation responses for Simulation Conditions 1 and 2 are shown in Figures 5.62 and 5.63 respectively. The oscillations present in the 3Φ model response, particularly for Simulation Condition 2, give a much better agreement with the true system behaviour (Figure 5.57) and the frequency of the oscillations at approximately 12 cycles per unit time is a good match.

To compare the open loop behaviour, a sequence of step voltage demands were applied to the q -axis via the input to the PWM inverter. The resulting speed and current responses of the 3Φ model and the real system are shown back to back in Figure 5.64. The fundamental difference between the two sets of time responses is that the voltage demand applied to the q -axis induces significantly more current in the d -axis with the three phase model compared to that of the real system. This is more and more pronounced as the speed of the motor increases, and the smaller proportion of current induced in the q -axis leads to reduced torque and therefore lower speeds for the same input. The same input sequence applied to the $d - q$ axis model generated the responses shown in Figure 5.65. In this case, the agreement is significantly better, suggesting that the real PWM inverter operates much more like the ideal case than its implementation in the 3Φ model.

This analysis highlights an obvious limitation of the PWM inverter model for open loop behaviour but its effect on the system in closed loop does improve the agreement with experimental data. It is considered that the phase error introduced may not be representative of the real system but when the current loop is closed, the system appears to become robust to this error, allowing for a better match in closed loop. Therefore, although simplistic, the PWM model is considered to be useful.

- Parameter estimation:** To check if modifying model parameters could give rise to a better agreement between the model and experimental data, the parameters that were considered to be most influential were included in a parameter estimation exercise. In this approach, the structure of the model is retained but values of its parameters are altered to minimise the overall deviation from experimental performance. The process can be applied to open or closed loop systems in the time and/or frequency domain. The process can be automated by the application of commercially available optimisation tools or performed manually by altering each parameter in an iterative fashion to best match the measured system behaviour. With parameter estimated models, although the input-output behaviour of the system may better match experimental performance, the internal properties of the model may deviate from that in the real system and it is possible for the estimated parameter values to be non-physical. This is particularly true with the automated tools that are available, where examples of non-physical values may be an estimated resistance that is 50% greater than that of the accurately measured component, or in extreme cases it is possible to have impossible values such as negative inertias or inductances.

To develop a parameter estimated model of the EPHS system, a manual approach was adopted to fit the open loop motor model to experimental responses in the time domain obtained by applying step




Figure 5.64: Open loop response comparison for the nominal 3Φ model with a sequence of step voltage demands applied to the q -axis




Figure 5.65: Open loop response comparison for the nominal d - q axis model with a sequence of step voltage demands applied to the q -axis

voltage inputs to the q -axis via the PWM. Although with a manual approach the model fit may not be quite as good as achieved with an automated approach, there are benefits to a manual approach and it was quite feasible for this system due to the modest number of parameters to be estimated. Only the lumped mechanical inertia, J , lumped mechanical damping, B_m , resistance, R , and inductance, L , were considered. By independently varying each parameter and observing its effect on the system, a better understanding of the system was gained and in principle this can improve confidence about the validity of the parameter estimates derived. Another benefit of the manual approach is that where a trade-off is required in the model fitting, the engineer can make this trade-off intelligently

A comparison of open loop behaviour using an estimated parameter set is shown in Figure 5.66. This model was fitted using the three phase model and so any error understood to be caused by the inverter model is effectively compensated for by altering the model parameters. With the estimated model there is still a shortfall in motor speed compared to the real system but the variation in steady-state speed with applied q -axis voltage is in much better agreement. In addition, the proportion of d and q -axis current induced by the input sequence is in much better agreement throughout the speed range. The only deterioration in agreement is that the parameter estimated model settles at the next steady-state speed condition more abruptly than in the measured response.

These exercises reveal that error in the parameter values of the system or error in the PWM model both appear to be possible candidates for the cause of discrepancy in open loop responses. In reality, it is likely that a certain amount of error is present in both of these aspects of the model and for a more rigorous analysis, a match of frequency domain behaviour should also be sought. However, one barrier to this is that error in the PWM model may alter the frequency response characteristics of the system and lead to an incorrect estimation of system parameters.

Some additional insight into the validity of the estimated parameters is gained by observing closed loop behaviour of the speed loop. A comparison of the three phase model with nominal and estimated parameters against measured performance is shown in Figure 5.67 for Simulation Condition 1. Comparing the speed responses, the estimated model exhibits a similar overshoot of the high speed reference as the measured response. In addition, a significantly better match is observed during the deceleration ramp. However, the rate at which the motor accelerates to the step demand is a little too high with the estimated model, particularly during the initial rise, and the settling time is significantly shorter than in the measured response. A similar comparison for Simulation Condition 2 is shown in Figure 5.68. Here, damping of the nominal model is in better agreement with the measured response but otherwise, ignoring the aforementioned excessive rate of acceleration, the estimated model captures the fundamental characteristics of the measured response well. From observation of these closed loop responses, it is clear that the parameter estimated model is too far removed from the nominal model but in some respects the model agreement is improved. Therefore in combination with an improvement to the PWM inverter model, a more rigorous system identification exercise may be fruitful in fine tuning the model behaviour.




Figure 5.66: Open loop response comparison for the estimated 3Φ model with a sequence of step voltage demands applied to the q -axis




Figure 5.67: Comparison of nominal and estimated parameter models with measured response for Simulation Condition 1 with low-order AW




Figure 5.68: Comparison of nominal and estimated parameter models with measured response for Simulation Condition 2 with low-order AW

5.7 Summary

In this chapter the topic of applying anti-windup conditioning to compensate for current saturation within a PMSM speed control system was tackled. The application of a number of modern optimal designs and more traditional ad-hoc designs were applied to a simple single-axis model of the PMSM speed control system for application in Electrically Powered Hydraulic Steering. This simple model allowed some of the characteristics of the various anti-windup designs to be seen more clearly and a selection were chosen for testing with a complex nonlinear three phase model of the same system with control implemented at multiple rates in discrete time.

In the nonlinear model, the current demand saturation function is a limit on the magnitude of the d - q axis demand vector. In order to apply linear anti-windup conditioning, a novel method of interpreting this nonlinear multivariable constraint as a time-varying magnitude constraint on a scalar signal was devised. This approach allowed the synthesis of modern optimal anti-windup compensator designs for linearised models of the complex model. These designs were tested alongside an appealing classical design in simulation and then validated by application on an EPHS test rig. In spite of some discrepancies between the model behaviour and that of the real system, good performance was achieved by the anti-windup designs both in simulation and in experiment. Furthermore, the effect of perturbations to the designs were observed to be similar in both simulation and experiment. Therefore it was concluded that although the model may not be described as high fidelity, the fundamental dynamics are correct, allowing successful model-based designs to result. A brief foray into understanding the limitation of the model revealed that the most likely cause of the discrepancy is the PWM inverter model, although parameter estimation may also serve to improve the model fit.

The culmination of the chapter is a successful Back Calculation & Tracking design (BCAT) and low-order dynamic design. Both compensators were observed to perform very well in both simulation and experiment and both out-performed an in-house Back Calculation (BC) approach. The BCAT design, benefitted from ease of tuning and exceptionally simple implementation. Unfortunately it did not satisfy any nonlinear stability or performance guarantees and hence tuning was largely simulation based. The optimal low order compensator provided similar, if not slightly better performance and although the two first-order filters were easy to implement, they were obviously of greater complexity than the static BCAT compensator. Tuning was also relatively easy, and, although some iteration is required, the extra degrees of freedom present in the filter tuning makes it easier to achieve a satisfactory design. Perhaps the most significant advantage of the low order compensator was the stability and performance guarantees it provides. The low order compensator tested (synthesised using a trim point at approximately one third of the maximum speed) was able to provide stability guarantees for 95% of the PMSM's operating envelope and yielded an \mathcal{L}_2 gain bound which was constant across most of this. Such characteristics provide reassurance and confidence when implementing such compensators in experimental systems.

Chapter 6

Override Control and its Application to EPHS Motor Control

In this chapter we introduce the topic of override control and demonstrate its application as a method of handling motor current magnitude constraints for the EPHS system of Chapter 3. The anti-windup approach described in Chapter 5 has shown success in simulation and has also been proven by practical industrial application. However, one characteristic of the anti-windup approach is that the desired current constraint is imposed by placing a limit on the magnitude of the current *demand* rather than the *actual* current. This was required to translate the current limit into a plant *input* constraint in order to be amenable to anti-windup compensation. As a result, the anti-windup problem generated is, arguably, somewhat artificial. A crucial implication of this is that in spite of the apparent success of a given anti-windup design, the motor current may still exceed the limit even though the current demand does not. This is a fundamental limitation of the application of anti-windup to this problem.

Override control is a strategy that enables constraints to be applied to states and/or outputs of the plant. Since the motor currents are states of the plant for the PMSM system we consider, the task of limiting current is more naturally cast as an override problem. The following sections give an overview of some existing override control strategies and demonstrate the application of a chosen approach to the current limitation problem. The performance of this approach is then compared to the anti-windup based approach considered in Chapter 5.

6.1 Introduction

Override control has been developed as a strategy which allows limits to be imposed on some of the states and/or outputs of a system. These limits are generally introduced due to safety requirements or to extend component life by ensuring that they are operating within specification. When these limits are approached or exceeded, a secondary controller ‘overrides’ the nominal controller in order to prevent or minimise violation of these limits. The origins of override control lie in the field of process control where limits on system variables such as pressure and temperature are required to be respected [20]. Override control remains a little known branch of control theory but it has been applied successfully to other application areas including flight control.




Figure 6.1: Classic Override Control Problem

Figure 6.1 represents a typical control system for which output limitation is required. The primary control objective is defined as a tracking problem for which the plant output $y_1(t)$ is required to track the reference $r(t)$. The secondary control function is to ensure that an additional plant output, $y_2(t)$, remains within the range (y_{2min}, y_{2max}) .

6.1.1 Classic Override Control Approaches

A classic override control solution to this problem described in [53] is depicted in Figure 6.2. This is a multi-mode control system for which in the primary mode, controller K_1 is active exclusively and governs nominal tracking performance for output $y_1(t)$. The secondary control mode is designed to restrict $y_2(t)$ to the range (y_{2min}, y_{2max}) and a set of logical conditions govern the switching between modes. The secondary controller is essentially a proportional regulator with gain P_2 and an optional dynamic element $K_2(s)$ to aid closed loop stability. This secondary regulator is designed to cause $y_2(t)$ to track toward y_{2max} or y_{2min} via the corresponding control signals $u_M(t)$ or $u_m(t)$ and due to the steady-state error associated with proportional control the constraints will be satisfied at steady state in the absence of disturbances. The proximity of $y_2(t)$ to the limits at steady-state and any overshoot allowed can be tuned by altering the gain P_2 .

To better understand the operation of this system let us consider the behaviour when a ramp reference signal $r(t)$ is applied such that the primary control signal $u_1(t)$ increases linearly with time and the secondary plant output $y_2(t)$ does likewise toward the positive upper limit y_{2max} . Initially, the secondary control signal $u_M(t)$ is positive with a value $P_2 y_{2max}$ and the primary control signal $u_1(t) = 0$. As the reference is applied, $u_1(t)$ departs from zero and becomes progressively more positive, whereas $u_M(t)$ becomes progressively less positive. At a point in time the signals cross over such that $u_1(t) > u_M(t)$. This initiates a mode change and the switching logic passes $u_M(t)$ on to the plant, preventing the ever increasing primary control signal from driving $y_2(t)$ beyond y_{2max} . Equivalent but opposite behaviour is observed as $y_2(t)$ approaches the lower limit y_{2min} .

This override control approach is typical of many classical approaches such as found in [19], sharing the following common characteristics:

- When the output constraints are approached (or violated), an alternative controller is switched in

- The primary controller is always active, thereby helping to maintain tracking performance and disturbance rejection during limit violation
- Output violation cannot be prevented entirely as violation of a limit is required to activate the compensator (although it is possible to artificially lower the limits to produce a similar effect)
- Performance is optimised via a constrained synthesis routine
- Stability of the resulting nonlinear system is sought at the design stage




Figure 6.3: Modern override compensation approach

6.2 Override Compensator Synthesis

In this work the override compensation methodology of Turner and Postlethwaite [86, 85, 90] is adopted due to its powerful approach and intuitive and systematic synthesis. A block diagram showing the generic override compensation framework used in this approach is shown in Figure 6.4. Here we have a reference signal, r , controller, K , and plant model, G . The outputs of the plant are partitioned as $\begin{bmatrix} y'_c & y' \end{bmatrix}'$ where $y \in \mathbb{R}^{n_y}$ represents the plant outputs used for feedback control and $y_c \in \mathbb{R}^q$ represents the outputs we desire to impose limits on, the so-called constrained outputs. The limits are modelled as $y_m = \text{sat}(y_c)$ according to (6.1) and (6.2) where q represents the number of constrained outputs and $\bar{y}_{c,i}$ represents the saturation threshold on the i th channel of y_c . Symmetric limits are defined here for simplicity but asymmetric limits can also be accommodated.

$$\text{sat}(y_{c,i}) = [\text{sat}(y_{c,1}), \dots, \text{sat}(y_{c,q})]' \quad (6.1)$$

$$\text{sat}(y_{c,i}) = \begin{cases} y_{c,i} & \forall y_{c,i} \in (-\bar{y}_{c,i}, \bar{y}_{c,i}) \\ \text{sign}(y_{c,i})\bar{y}_{c,i} & \forall |y_{c,i}| \geq \bar{y}_{c,i} \end{cases} \quad (6.2)$$

The vector signal \tilde{y} , defined as the difference between the measured and saturated versions of y_c , becomes non-zero when the limits are exceeded and is used to drive the override compensator, Φ , to regulate the constrained





Figure 6.4: Override compensation generic framework

Figure 6.5: Simplified override compensation generic framework using $Dz(\cdot)$

outputs to their respective limits via ϕ . The signal \tilde{y} can also be defined using the deadzone function, $Dz(\cdot) = I - \text{sat}(\cdot)$, according to (6.3), producing the simplified block diagram representation of Figure 6.5.

$$\tilde{y} = y_c - \text{sat}(y_c) = Dz(y_c) \quad (6.3)$$

The most powerful form of override compensation is that for which the compensator is given the authority to directly alter all controller states and outputs. Other, more simple, override structures can be considered as special cases of this such as in [61], but are not considered in this thesis. This flexibility is provided by augmenting the state space matrices of the baseline controller. Assuming a two degree of freedom type baseline controller with n_c states, viz.

$$K(s) \sim \left[\begin{array}{c|cc} A_c & B_{cr} & B_c \\ \hline C_c & D_{cr} & D_c \end{array} \right].$$

The augmented controller representation has the form of (6.4) where $\phi_1 \in \mathbb{R}^{n_c}$ and $\phi_2 \in \mathbb{R}^m$ are the elements of the override signal which alter the controller states and outputs respectively. The linear plant model is described by the state-space representation of (6.5).

$$K(s) \sim \begin{cases} \dot{x}_c & = A_c x_c + B_{cr} r + B_c y + \phi_1 \\ u & = C_c x_c + D_{cr} r + D_c y + \phi_2 \end{cases} \quad (6.4)$$

$$G(s) \sim \begin{cases} \dot{x}_p &= A_p x_p + B_{pd}d + B_{pu}u \\ y &= C_p x_p + D_{pd}d + D_{pu}u \\ y_c &= C_{pc}x_p + D_{pdc}d + D_{puc}u \end{cases} \quad (6.5)$$

When the signal y_c lies within the limits, we have $\tilde{y} = 0$ and the override compensator is inactive. Thus, a pre-requisite for stability of the override compensated closed loop is that the nominal closed loop is stable and well-posed, i.e. we have stability when $\Phi \equiv 0$. This is essentially the same requirement as that made on the nominal closed loop system for anti-windup compensation as described in previous chapters except that there is no requirement for A_p to be Hurwitz for global stability.

Due to the fact that the override is activated only when \tilde{y} becomes non-zero, some violation of the limit must be tolerated. However, this can usually be reduced to acceptable levels by appropriate override design. In addition, by setting the limits of the deadzone function a few percent lower than the desired physical limit, an amount of ‘headroom’ can be provided. This can be useful to ensure that the constrained outputs lie below the desired limit at steady-state.

6.2.1 Static Override Compensator Synthesis

Let us consider first the synthesis of a static override compensator for which Φ is a static matrix mapping $\tilde{y} \in \mathbb{R}^q$ to $\phi = [\phi'_1 \ \phi'_2]' \in \mathbb{R}^{n_c+m}$. This is expounded in [86] using essentially the same tools as those used for the synthesis of static anti-windup compensators in [91]. Static override compensation can be desirable as it has no states, making it simple to implement and computationally efficient although there is no guarantee that it will provide the desired performance.

The central objective of override compensation is to bring the signal y_c out of saturation in a timely manner when a limit is exceeded and to minimise the amount by which the limit is violated. This equates to returning \tilde{y} to zero quickly and also minimising the magnitude of \tilde{y} . A useful method of capturing these requirements mathematically is to consider the problem of minimising the \mathcal{L}_2 norm of \tilde{y} . This performance objective is tackled by minimising the induced \mathcal{L}_2 gain between the exogenous input $w = [r' \ d']'$ and \tilde{y} . A simplified representation of the corresponding closed loop system is shown in Figure 6.6 where the state-space representation of G_{cl} is given by (6.6) and the corresponding matrix definitions are deferred to Appendix C. Note that for override synthesis the output y need not be considered.

$$G_{cl}(s) \sim \begin{cases} \dot{x} &= Ax + B_0w + \bar{B}\phi \\ y &= C_yx + D_{y0}w + \bar{D}_y\phi \\ y_c &= Cx + D_0w + \bar{D}\phi \end{cases} \quad (6.6)$$

A secondary performance objective used is to minimise the induced \mathcal{L}_2 gain between $w = [r' \ d']'$ and $\phi = [\phi'_1 \ \phi'_2]'$. This can help to restrict the magnitude of the compensator signal ϕ and therefore allow some control




Figure 6.6: Static override compensation closed loop

over the aggressiveness of the compensator. It also allows a trade-off between adhering to the limits on y_c and performance of the primary control loop. These two performance conditions are combined to form the following cost function for which we seek to minimise γ .

$$\left\| \begin{bmatrix} W_1^{1/2} \tilde{y} \\ W_2^{1/2} \phi \end{bmatrix} \right\|_2 < \gamma \|w\|_2 \quad (6.7)$$

Here the diagonal matrices W_1 and W_2 are used as weights to reflect the relative importance of each performance condition. An initial choice of W_1 is typically an identity matrix, and then the diagonal elements can be varied from 1 to trade-off the relative importance of each channel in the performance optimisation. An initial choice of W_2 is typically a diagonal matrix with diagonals in the order of $1e^{-5}$ such that the second performance objective has minimal effect on the initial design. The gain of this matrix can then be increased to add more importance to the second performance condition in the optimisation and reduce the magnitude of the override signals. In addition, the matrix diagonal elements can be altered individually to favour feedback to either the controller states or controller outputs.

For stability analysis of the override compensated closed loop we apply the Circle Criterion [45]. In order to generate a synthesis routine, the same tools as used for the anti-windup compensator synthesis are applied. Stability of the override compensated closed loop is guaranteed by the Circle Criterion, for which a sector condition describing the deadzone function is combined with a quadratic Lyapunov function to form the inequality of (6.8) in $P > 0 \in \mathbb{R}^{n_c+n_p}$ and diagonal matrix $W > 0 \in \mathbb{R}^q$.

$$\frac{d}{dt}(x'Px) - \tilde{y}'W(y_c - \tilde{y}) < 0 \quad (6.8)$$

The \mathcal{L}_2 performance condition of (6.7) is guaranteed to be satisfied if the inequality of (6.9) in the two norm holds. Using simple matrix arithmetic and by application of the S-procedure, the performance and stability conditions can be combined into the single inequality of (6.10). By substituting in for various signals and applying tools such as the Schur Complement and Congruence Transformations, the LMI of (6.11) is produced. This can be solved using standard LMI solvers, such as present in the MATLAB Robust Control Toolbox, to minimise the scalar $\gamma^2 > 0$ subject to $L \in \mathbb{R}^{(n_c+n_p) \times q}$, $Q > 0 \in \mathbb{R}^{n_c+n_p}$ and a diagonal matrix $U > 0 \in \mathbb{R}^q$.

The static compensator matrix is then recovered as $\Phi = LU^{-1}$. More detail on the derivation of this LMI can be found in [86].

$$\left\| \begin{bmatrix} W_1^{1/2} \tilde{y} \\ W_2^{1/2} \Phi \tilde{y} \end{bmatrix} \right\|^2 < \gamma^2 \|w\|^2 \quad (6.9)$$

$$\frac{d}{dt}(x'Px) + \tilde{y}'(W_1 + \Phi'W_2\Phi)\tilde{y} - \gamma^2 w'w + 2\tilde{y}'W(y_c - \tilde{y}) < 0 \quad (6.10)$$

$$\begin{bmatrix} QA' + AQ & \bar{B}L + QC' & B_0 & 0 & 0 \\ * & -2U + \bar{D}L + L'\bar{D} & D_0 & U & L' \\ * & * & -\gamma^2 I & 0 & 0 \\ * & * & * & -W_1^{-1} & 0 \\ * & * & * & * & -W_2^{-1} \end{bmatrix} < 0 \quad (6.11)$$

6.2.2 Dynamic Override Compensator Synthesis

Static compensation can be applicable to many systems but with some practical systems it is desirable to limit the bandwidth of the override control signal, ϕ . This may be simply to smooth the signal or to better manage the performance and stability trade-off by shaping the frequency content of the signal. To this end, a synthesis routine was designed by Turner and Postlethwaite [86] which allows the compensator to be specified as a transfer function matrix, $\Phi(s)$. In addition to the above, when the constrained output measurements contain noise, it is desirable to filter this out to prevent the noise component of the signal activating the compensator when the true signal level is still within the limits. This is achieved by including a filter, $F_o(s)$, on the constrained outputs of the system.

The dynamic compensator $\Phi(s)$ is chosen to consist of a dynamic element and static element in series according to $\Phi(s) = \tilde{\Phi}(s)K_\phi$. Here $\tilde{\Phi}(s)$ is the dynamic element chosen by the designer - usually a diagonal unity DC gain transfer function matrix - and K_ϕ is a static gain matrix to be synthesised. The positioning of these additional elements in the system is shown in the block diagram of Figure 6.7.




Figure 6.7: Dynamic override block diagram

Splitting the compensator into the components $\tilde{\Phi}(s)$ and K_ϕ allows the additional dynamic elements to be incorporated into the closed loop model $\tilde{G}_{cl}(s)$, simplifying the system representation from that shown in

Figure 6.7 to that of Figure 6.8. The optimisation problem that follows, therefore, is no more complex than for the static compensator since we simply seek a static gain matrix, K_ϕ , which minimises the chosen performance constraint and satisfies the stability condition.




Figure 6.8: Simplified dynamic override block diagram

The output filter, $F_o(s)$, and the dynamic element of the compensator, $\tilde{\Phi}(s)$, are described by (6.12) and (6.13). These dynamics are absorbed into the state-space representation of $\tilde{G}_{cl}(s)$ in (6.14) and definitions of the matrices for $\tilde{G}_{cl}(s)$ are given in Appendix C.

$$F_o(s) \sim \begin{cases} \dot{x}_{F_o} &= A_{F_o}x_{F_o} + B_{F_o}\tilde{y}_c \\ y_c &= C_{F_o}x_{F_o} \end{cases} \quad (6.12)$$

$$\tilde{\Phi}(s) \sim \begin{cases} \dot{x}_f &= A_f x_f + B_f \psi \\ \phi &= C_f x_f + D_f \psi \end{cases} \quad (6.13)$$

$$\tilde{G}_{cl}(s) \sim \begin{cases} \dot{\tilde{x}} &= \tilde{A}\tilde{x} + \tilde{B}_o w + \tilde{B}_1 \psi \\ y_c &= \tilde{C}_1 \tilde{x} + \tilde{D}_{01} w + \tilde{D}_1 \psi \\ \phi &= \tilde{C}_2 \tilde{x} + \tilde{D}_2 \psi \end{cases} \quad (6.14)$$

The performance condition which we seek to minimise is given by (6.7) as in the case of static compensation. However, the signals \tilde{y} and ϕ now have a different representation. In particular, ϕ , is now an internal signal within \tilde{G}_{cl} and so an explicit expression for ϕ is required as given in (6.14). Using this expression for ϕ and substituting in $\psi = K_\phi \tilde{y}$, the performance condition is guaranteed to be satisfied if the inequality of (6.15) holds. This is the equivalent of (6.9) in the case of static synthesis. The derivation continues in the same vein as for static synthesis, resulting in the LMI of (6.16) in variables $L \in \mathbb{R}^{(n_c+m) \times q}$, $Q > 0 \in \mathbb{R}^{n_c+n_p}$, diagonal matrix $U > 0 \in \mathbb{R}^q$ and scalar $\gamma^2 > 0$ for which we seek to minimise γ .

$$\left\| \begin{bmatrix} W_1^{1/2} \tilde{y} \\ W_2^{1/2} (\tilde{C}_2 \tilde{x} + \tilde{D}_2 K_\phi \tilde{y}) \end{bmatrix} \right\|^2 < \gamma^2 \|w\|^2 \quad (6.15)$$

$$\begin{bmatrix} Q\tilde{A}' + \tilde{A}Q & \tilde{B}_1L + Q\tilde{C}'_1 & \tilde{B}_0 & Q\tilde{C}'_2 & 0 \\ * & \tilde{D}_1L + L'\tilde{D}'_1 - 2U & \tilde{D}_{01} & L'\tilde{D}'_2 & U \\ * & * & -\gamma^2I & 0 & 0 \\ * & * & * & -W_2^{-1} & 0 \\ * & * & * & * & -W_1^{-1} \end{bmatrix} < 0 \quad (6.16)$$

Given solutions to (6.16), K_ϕ can be calculated then as $K_\phi = LU^{-1}$.

Notes on tuning:

Tuning of the dynamic compensator follows in a similar method as for static compensation except that there is the additional freedom of choosing filter dynamics. As the presence of filter dynamics increases the complexity of the controller implementation and computational demands, it is desirable to use first order filters where possible.

For the noise filters, $F_o(s)$, first order low-pass filters are usually adequate and the bandwidths can be chosen based on an appropriate noise model. These filters typically have little influence on the optimisation result as their bandwidth is usually significantly higher than those in $\tilde{\Phi}(s)$. The performance trade-off is generally between noise rejection, achieved by reducing the bandwidth, and reducing phase lag in the feedback path, achieved by increasing the bandwidth. To extract higher performance, one might consider higher order filters with damping of less than 0.7 to reduce phase lag in the pass band.

The override input filter, $\tilde{\Phi}(s)$, is chosen usually as a diagonal transfer function matrix of unity gain low pass filters, simplifying the tuning to the choice of appropriate bandwidths. A good method is to start with low bandwidths and gradually increase them whilst observing the effect on performance. It is common that higher bandwidths can be tolerated in the feedback to controller states than to the controller outputs as some low pass filtering usually occurs within the controller.

The tuning of the dynamic override compensator is intrinsically more intuitive than for an anti-windup design of similar complexity for two reasons. First of all, the closed loop system is stable without any compensation applied. Secondly, there is little reason to consider filter types other than low-pass designs as for the output (noise) filter we seek to reduce gain at high frequencies, and the compensator is mainly required to act at low frequencies. However, there is much less work on override control in the wider literature so tuning rules are still relatively crude.

6.3 Application to Current Limitation in PMSM Speed Regulation

As mentioned in the introduction, for the application of anti-windup to the current limitation problem in PMSM speed regulation we made use of the cascade structure of the PMSM speed control system to impose limits on the motor currents via their demands. The success of this approach in limiting the motor current relies on the premise that constraining the magnitude of the current demand to a given limit will cause the magnitude of the

actual motor currents to be constrained to the limit also. Due to integral action present in the current controller this constraint is satisfied at steady-state, but it may be violated transiently due to the dynamic relationship between the motor currents and their demands. An example of this is shown in Figure 6.9 where the magnitudes of the current demand and measured current in the anti-windup simulation model are compared. Note that even though the current demand never exceeds the limit, the magnitude of the *actual* current vector does violate the limit. Also note that at the onset of the large speed step demand, the current demand is limited before the actual current comes close to the limit. Therefore, for a period of time the limit on the current demand proves to be conservative and unnecessarily restricts the control effort applied. This demonstrates the drawbacks involved when using a limit on the current demand to impose a constraint on the actual current.




Figure 6.9: Current saturation in the nonlinear model with anti-windup compensation

Since in the override approach, the limit is applied to the *actual* motor currents rather than the demand, the application of override compensation may afford more confidence in its limitation of the motor current than provided by the anti-windup approach.

6.3.1 Current limit model

The constraint we desire to impose on the motor current is a limit on the norm of the current vector according to

$$\sqrt{i_d^2 + i_q^2} \leq i_{max}, \quad (6.17)$$

hence, the limit is a nonlinear function of the plant states. This is equivalent in structure to the constraint applied to the current demand vector for the purposes of anti-windup in Chapter 5 and so similar options regarding how

to incorporate the constraint and define the structure of the compensator apply. The key requirement for such a constraint model is that it is a tractable linear function of the states.

One option is to define time-varying limits on both the d and q -axis current measurements such that $\tilde{y} = [Dz(i_d) \ Dz(i_q)]$ where the width of each deadzone function is determined by the phase advance angle and the current limit. This would facilitate a multivariable override design that is driven by the vector signal \tilde{y} and feeds back to both elements of the current demand vector and also the motor speed measurement. However, one problem with this is that because the sign of i_d should not change during acceleration and deceleration, violation of the d -axis constraint would cause the compensator to drive the motor speed measurement in the same direction whether the external inputs demand an acceleration or deceleration. Supposing that the limit is violated during acceleration, the compensator could be designed such that violation of the d -axis constraint would drive the speed measurement in the positive direction, thereby reducing the magnitude of the error signal and aiding the system to respect the limits. In the deceleration condition the same compensator would also drive the speed measurement in the positive direction, increasing the magnitude of the error signal and serving to drive the system further into saturation. A neat solution to this problem is to adopt a SISO representation of the constraint similar to that used for anti-windup application.

Let us consider that the current angle is determined exclusively by the phase advance map such that the d -axis current demand is a function of the motor speed and the q -axis demand. Furthermore let us consider that the output of the override compensator $\phi = [\phi_1 \ \phi_2]'$ influences the q -axis current demand directly by ϕ_2 and via the controller dynamics by ϕ_1 . The compensator can now be driven by violation of the q -axis current measurement alone and the threshold of the limit on the q -axis can be defined as (6.18). With this approach the number of compensator input and output signals is reduced by one. This reduces the size of the compensator gain matrix by four elements and may halve the number of states required for the input filter in a dynamic design.

$$i_{q,lim} = \begin{cases} \pm \sqrt{i_{max}^2 - i_d^2}, & |i_d| \leq i_{max} \\ 0, & |i_d| > i_{max} \end{cases} \quad (6.18)$$

As mentioned previously, the form of override compensation adopted does not prevent violation of the constraint entirely because the constraint must be exceeded for the compensator to become active. To reduce the amount by which the limit is exceeded it is common to set the limit in software a few percent lower than the limit we wish to enforce. In this example this would be achieved by calculating $i_{q,lim}$ according to (6.19) with $k < 1$. This can be very effective when the limit is approached quite slowly but when approached quickly, delays in compensation due to filter dynamics, time delays and the dynamic properties of the current control loop can result in the limit still being exceeded. In this case, further reduction to the value of k may be required with the adverse effect that the limit is reduced further in steady-state operation.

$$i_{q,lim} = \begin{cases} \pm \sqrt{(k i_{max})^2 - i_d^2}, & |i_d| \leq i_{max} \\ 0, & |i_d| > i_{max} \end{cases} \quad (6.19)$$

An appealing approach to tackling this problem may be to include derivative action in the compensator. This could increase the aggressiveness of the compensator when the limit is being approached more quickly. However, the ‘aggressiveness’ may still be fundamentally limited by stability considerations and for practical systems, derivative control tends to make a system sensitive to noise. Another option would be to use a lead-lag filter design for $F_0(s)$ or $\tilde{\Phi}(s)$ to provide phase advance in a restricted frequency range although this may also increase noise sensitivity. Ad-hoc methods that temporarily lower the limit when a large transient violation is expected could also be employed.

6.3.2 Linear simulation results

In the same manner as used for the anti-windup designs of Chapter 5, a linear model of the current control system is constructed using a trim speed of 0.5 units under no load. Then the phase advance map is linearised and incorporated into the current control system model to form the linear model, $G(s)$, which we consider to be the plant for the purposes of override design. The derivation of this model is included in Appendix C.2. The resulting state-space model is given in (6.20) for which the single input is the q -axis current demand and the plant output vector is $y = [\omega_m \ i_d \ i_q]'$. Combined in feedback with the linear PI controller of (6.21), for which the input vector is $[r \ \omega_m]'$, the linear closed loop represents the dynamics of the nonlinear speed control system in the vicinity of the trim point and provides the d and q -axis currents as additional outputs for use by the override compensator. A representation of the override compensated closed loop using the linear system model is shown in Figure 6.10.




Figure 6.10: Representation of the override closed loop with linear plant model

$$G(s) \sim \left[\begin{array}{ccccc|c} 0 & 0 & -9.535 & 0 & 0 & 0 \\ 0 & 0 & 0 & -9.535 & 0 & 9.535 \\ 22730 & 0 & -1463 & 522.6 & 12.05 & 0 \\ 0 & 22730 & -522.6 & -1463 & -497.3 & 1109 \\ 0 & 0 & 0 & 168.3 & -3.882 & 0 \\ \hline 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{array} \right] \quad (6.20)$$

$$K_{spd}(s) \sim \left[\begin{array}{c|cc} 0 & 13.3568 & -13.3568 \\ \hline 1 & 0.544 & -0.544 \end{array} \right] \quad (6.21)$$

The d -axis current measurement is used for practical implementation to manipulate the saturation threshold on i_q . However, for the purpose of compensator design, the d -axis current output is ignored as only the q -axis current measurement is passed through the deadzone function and used to drive the compensator. Although the d -axis current does affect the saturation threshold on i_q , this has no bearing on synthesis since the sector bounded nonlinearity used to describe the deadzone encompasses all possible deadzone and saturation functions regardless of the limits used. For dynamic override synthesis where the plant outputs to be constrained are filtered, a SISO output filter is used at the design stage as only the q -axis current is required to be filtered. However, for implementation, the d -axis signal will also influence the signal \tilde{y} via its influence over $i_{q,lim}$. Therefore, in order to reduce noise in the signal \tilde{y} it may also be desirable to filter the d -axis signal also. In the subsequent linear simulation tests, override compensators are designed for the model linearised at a speed of 0.5 units under no load and tested on the same linear model.

Static OR Compensation

A static override compensator was designed using $W_1 = 1$ and $W_2 = \text{diag}(1e-7 \ 1e-7)$. The design produced is defined by the feedback gain matrix of (6.22) and achieves the \mathcal{L}_2 gain performance level specified in (6.23). The choice of W_2 ensures that the secondary optimisation objective has minimal impact on the optimisation for performance.

$$\Phi = [-552.5 \ 108.3]' \quad (6.22)$$

$$\gamma_{static} = 9.23 \quad (6.23)$$

Figure 6.11 compares the nominal linear (unconstrained) performance to that of the constrained system with static override compensation. For best control of the motor current magnitude, $k = 0.98$ is chosen such that the compensator seeks to impose the limit 2% below i_{max} . This enables the compensator to activate just prior to reaching the real limit and avoid its violation completely. Static compensation is very effective at imposing the desired constraint and has very little detrimental impact to the tracking performance. The compensator is able

to consistently hold the motor current at the limit when activated, facilitating maximum use of the available control effort. This allows for minimum deviation from the nominal linear performance which we consider to be the ‘desired’ performance. Although performance of this compensator is essentially ideal, time delays and noise are expected to cause problems for practical implementation and so this performance level may not be achievable in practice. In fact, even in continuous time simulation a small solver time step was required for smooth control behaviour to be observed.




Figure 6.11: Static override compensation - linear model simulation

Dynamic OR Compensation

A dynamic override compensator was designed using $W_1 = 1$ and $W_2 = \text{diag}(1e-7 \ 1e-7)$ and the following filters:

$$F_o(s) = \frac{1}{0.0004s + 1} \quad (6.24)$$

$$\tilde{\Phi}(s) = \frac{1}{0.0159s + 1} \quad (6.25)$$

Constraints on the filter bandwidths that could be implemented on the discrete model were taken into account at this stage such that the designs produced for the linear continuous time system could also be tested directly on the discrete nonlinear system. The bandwidth of the output filter was chosen initially to be as high as

possible in order to impose the minimum constraint on performance. Considering limitations due to sampling, the bandwidth was chosen to equal the Nyquist frequency of the current loop, $\frac{1}{2\tau_{crt}}$, where τ_{crt} is the sampling period of the current loop. For the input filter, the bandwidth was increased progressively until optimum performance was observed. Generally speaking, the greater the bandwidth, the quicker the compensator is able to act, providing the potential for improved performance. However, increasing the filter bandwidth also changes the phase behaviour of the system and it was found that increasing the bandwidth can also result in reduced performance as the gain of the compensator had to be reduced. The filter chosen was deemed to provide a good trade-off between the bandwidth of the override signals and aggressiveness of the resulting compensator. The final stage in tuning was to reduce the bandwidth of the output filter as much as possible to improve robustness and reduce susceptibility to noise whilst having as little impact on performance as possible and the compensator gain matrix produced and associated \mathcal{L}_2 gain performance level are given in (6.26) and (6.27) respectively. Figure 6.12 compares the nominal linear performance to that of the constrained system with dynamic override compensation. Two results are shown, one for which $k = 1$, and one for which $k = 0.92$.

$$K_\phi = [-340.4 \quad -1.606]' \quad (6.26)$$

$$\gamma_{dynamic} = 15.80 \quad (6.27)$$

For the dynamic compensator with $k = 1$, the compensator seeks to impose the limit at i_{max} . The limit is exceeded by approximately 9.5% at the onset of the large step demand compared to $< 1\%$ with the static design. This is because the dynamic compensator has a lower (finite) bandwidth and is unable to act as quickly. This bandwidth limitation also affects the action of the compensator immediately after saturation. The energy accumulated in the compensator filters while the current limit is exceeded is dissipated once saturation ceases, causing the motor current to drop further below the imposed limit. This reduces usage of the available control effort and results in a larger deviation from nominal linear performance. With $k = 0.92$ the dynamic compensator seeks to impose the limit at 8% below i_{max} , causing the compensator to become active sooner and the limit to be respected. Note that the limit has to be reduced by a much greater percentage with the dynamic design to prevent violation of the physical current limit. This lower limit reduces the maximum current at steady-state as well as the magnitude of transient peaks and could therefore be quite detrimental to performance.

Perhaps the most important comparison to be made between the static and dynamic override designs concerns the speed at which they are able to act due to bandwidth constraints imposed by the filter dynamics. As described in the above paragraphs, this accounts for a considerable amount of the performance difference between the two designs. Another related issue is the aggressiveness of each compensator design. Comparing the gain matrices of the static and dynamic compensators (6.22) and (6.26), although the same performance weighting matrices were used for the synthesis of each design, the static compensator is significantly more aggressive, which aids performance. Note also that despite the fact that each channel was given the same weighting in the design via the W_2 matrix, the static compensator makes much greater use of the second channel which directly influences the controller output. This channel bypasses the controller dynamics and therefore also contributes to the faster response.




Figure 6.12: Dynamic override compensation - linear model simulation

In summary, the override compensators perform very well on the linear continuous time model with little performance deterioration and minimal limit violation. This is particularly true for the static design, for which it is difficult to imagine a better performing design.

6.3.3 Nonlinear simulation results

In this section, the override compensator variants introduced for the linear model in the previous section are tested on the nonlinear multi-rate discrete-time model depicted in Figure 6.13. The current controller K_{crt} and output filter of the dynamic compensator operate at the fast sample frequency $1/\tau_{crt}$ and the remainder of the override compensator Φ , speed controller K_{spd} and phase advance controller operate at the slower rate, $1/\tau_{spd}$. A saturation constraint on the voltage is included to ensure that the magnitude of the $d - q$ axis voltage vector is correctly limited by the finite supply voltage and a simple back calculation type of anti-windup strategy is included to prevent windup of the current controller integrator states when the voltage demand exceeds this limit. The inclusion of these additional nonlinearities within the inner current control loop has a significant influence over performance compared to the linear model simulations. In addition, the nonlinearity of the phase advance controller is modelled fully so the current angle will vary through the speed range, giving rise to variations in the current loop dynamics and a time varying saturation limit on the q -axis current.

Two of the simulation conditions defined in Chapter 5 are selected for analysis as they represent two of the most demanding conditions in which override compensation is required to act. For Condition 1 the large step




Figure 6.13: Nonlinear model used for override compensation simulation

change in the speed demand causes large transient violations of the current limit. For Condition 2, the current limit would be violated at steady-state and so the compensator is required to maintain the current at the limit for an extended period. Other input sequences were tested but the two selected were found to be sufficient to capture the most significant characteristics.

- **Condition 1:** Step demand from idle speed to maximum speed under no load
- **Condition 2:** Step demand from idle speed to maximum speed with a static load applied such that the speed attained is reduced by 10% from the reference level

Nonlinear simulations under no load (Condition 1)

- **No OR:** Figure 6.14 shows the performance of the model without override compensation, the speed tracking performance being shown in the upper left plot. An important first point is that the speed controller used in this test was designed for use with anti-windup where a strict saturation limit was imposed on the current demand. The presence of this limit in the anti-windup approach enabled the controller to be tuned more aggressively for better small signal performance without causing large overshoots when large reference demands were made as the control action was naturally limited by this constraint. With this limit removed, the response differs, and an undesirable overshoot is observed for the large step reference. Another consequence of the removal of the current demand limit is that the voltage limit saturates more readily. Voltage saturation is indicated by the voltage magnitude reaching V_{max} in the lower left plot. It can be seen from this plot that voltage saturation starts at the onset of the large step demand and ceases approximately when the overshoot has reached its peak. It is clear that the controller tuning will affect the tracking response, but it is also true that the occurrence of voltage saturation itself is influential

in the response. The bottom right hand plot shows the violation of the current limit with the motor current approaching 1.2 units for a short period of time when the high speed demand is applied. The upper right plot shows the q -axis current against its time-varying limit, $i_{q,lim}$. The key points to note from this simulation are that the current magnitude exceeds the limit by approximately 15%, that the large step reference causes the voltage limit to saturate for at least 0.1 units of time, and that there is a substantial overshoot of the high speed reference.

- **Static OR:** The static compensator presented in the previous section had to be retuned for the nonlinear simulation model as the high gains caused very chattery behaviour in the discrete model. An acceptable design was achieved using $W_2 = [5e - 7 \ 0.1]'$, placing the emphasis on feedback to the controller state rather than the controller output. The compensator matrix produced is given by (6.28) which is much more similar to the DC gain of the dynamic compensator for the linear system (6.26) although still more aggressive. Consequently, the \mathcal{L}_2 gain level achieved by the design increases by two to that given in (6.29).

$$\Phi = [-346 \ -1.32]' \quad (6.28)$$

$$\gamma_{static} = 11.63 \quad (6.29)$$

The behaviour of this static design is shown in Figure 6.15. Tracking of the step reference is much improved compared to the case without override compensation with overshoot avoided completely and the duration of voltage saturation following the large step reference reduced significantly. Violation of the current limit is shown in the bottom right plot and occurs in two phases. Initially, the limit is exceeded by approximately 15% as in the simulation without override compensation. Following this, the compensator acts, bringing the mean magnitude of the motor current down to the limit but also introducing some chatter behaviour. This chatter is most prominent immediately following the initial limit violation, and causes the limit violation to increase toward 22%. Although this is undesirable, the mean magnitude of the motor current is brought down to the saturation level very quickly and the chatter dies down quite quickly.

- **Dynamic OR:** The dynamic compensator was fine-tuned for the nonlinear model by altering the input and output filters to (6.30) and (6.31) respectively. In addition for the purposes of implementation, a copy of the output filter is also required on the d -axis current measurement, by which the q -axis current limit is calculated. Rather than being a direct copy of the output filter, reduced chatter was observed by reducing the bandwidth of this filter. Its transfer function is given as $F_{od}(s)$ in (6.32). The compensator matrix produced is given by (6.33) and achieves the \mathcal{L}_2 gain level given in (6.34). Note that in this case, \mathcal{L}_2 performance is improved over that for the linear design. This is most likely to be because more time was spent optimising the design for the nonlinear system.

$$\tilde{\Phi}(s) = \frac{1}{(1.592e - 4)s + 1} \quad (6.30)$$

$$F_o(s) = \frac{1}{(2.55e - 4)s + 1} \quad (6.31)$$

$$F_{od}(s) = \frac{1}{(1.019e - 3)s + 1} \quad (6.32)$$

$$K_\phi = [-220.1 \quad -1.177]' \quad (6.33)$$

$$\gamma_{dynamic} = 13.38 \quad (6.34)$$

With dynamic override compensation applied and $k = 1$, the performance achieved is shown in Figure 6.16. Violation of the current limit at the onset of the step demand is of a similar magnitude as for the simulations without override and with static compensation. However, the dynamic compensator causes the current magnitude to track down to the limit more smoothly and swiftly than the static design. Tracking of the speed reference is comparable to that with the static design. In fact, the slightly less abrupt limitation of the current results in a small decrease in the rise time to the step reference.

A common feature of the simulation results described above is that the current limit is always exceeded by a similar magnitude regardless of what form of override is applied. This is due to the time delays associated with the discrete-time implementation and fast electrical dynamics of the system. Another important feature of these results is that the quality of the tracking response can be influenced considerably by voltage saturation and this has been seen to cause significant overshoots of the reference. Despite these shortcomings, the system with override engaged performs better than without any override.

Nonlinear simulations with a static load torque applied (Condition 2)

When a static load is applied to the system, the same reference sequence becomes more challenging to track since current is required in the q -axis to reject the load disturbance in addition to that required for the motor to accelerate and track the reference. For the following tests, the magnitude of the applied load is chosen such that the maximum speed reference (1 normalised unit) should be infeasible with the given current and voltage constraints. As a result, the current limit is expected to be reached during steady-state operation rather than just transiently as in the previous results.

- **No OR:** Figure 6.17 shows the performance of the system without override compensation. The first point to make is that the high speed reference is met even though it should be infeasible. This is because the motor currents are allowed to rise above the limit. Another point is that the prolonged period of voltage saturation between 0.2 and 0.6 units of time prevents the phase advance controller from achieving the desired current angle. This allow speeds in excess of the reference to be achieved without exceeding the current limit. This fact suggests that an alternative phase advance map may allow higher speed operation, but this topic is not considered within the scope of the thesis. Violation of the current limit at the onset of the step demand is more severe than for the no load case and the limit is exceeded by approximately 19%. The current limit is also violated later in the simulation as the current angle reverts back to that demanded by the phase advance controller following the escape from voltage saturation.




Figure 6.14: Nonlinear Simulation Condition 1 without override




Figure 6.15: Nonlinear Simulation Condition 1 with static OR




Figure 6.16: Nonlinear Simulation Condition 1 with dynamic OR

- **Static OR:** Figure 6.18 shows the performance of the system with static override compensation. The re-tuned compensator fails to prevent the transient violation of the current limit at the onset of the large step demand, and the limit is exceeded by approximately 22%, partly due to chatter. However, the aggressive action of the static compensator brings the system out of voltage saturation quickly with the effect that correct phase advance operation is maintained for the whole simulation. As a result, there is a smooth rise toward the steady-state motor speed of 0.89 units without overshoot. The obvious drawback to this design is the introduction of chatter to the motor current and voltage signals. It is not clear exactly what causes this but contributing factors are thought to be the use of high gains in a discrete implementation, the interaction of the two discrete feedback loops, and the switching behaviour of the compensator as the q -axis current is repeatedly driven below the limit by the compensator then back above the limit by the nominal controller. Another consideration is the time varying nature of the q -axis current limit which, as shown in the upper right plot, is seen to change at quite a high frequency. This chatter phenomenon also presents itself in the motor speed measurement and so is not a feature that can be overlooked.
- **Dynamic OR:** Figure 6.19 shows the performance of the system with dynamic override compensation. With the dynamic design, violation of the current limit at the onset of the high speed step reference is slightly increased compared to the case without compensation at $\approx 20\%$. This can be attributed to the slower action of the dynamic design, which also presents itself with a slower return to the limit. Chatter type behaviour is also present with this design although the magnitude is significantly lower due to the filtering and although present in the current and voltage signals, does not have a significant effect on the motor speed.

Nonlinear simulations with reduced supply voltage

In a vehicle the supply voltage is held at the required level by the alternator and assuming that the electrical system is functioning correctly, it is appropriate to consider that this voltage varies within a few percent of the specified level. If the alternator malfunctions the supply voltage may drop and this will influence the voltage saturation limit in the EPHS system. In the anti-windup application work of Chapter 5 the effect on EPHS performance was minimal and hence was not reported. However, in the override application the different architecture of the control system allows for a much more significant effect that warrants some attention.

To generate problematic voltage saturation the supply voltage limit, and therefore the voltage saturation limit, is reduced by 10% and Simulation Condition 1 is run, generating the response shown in Figure 6.20. Dynamic override compensation is applied in this case but the response is very similar without override and with static compensation because the current limit is not violated for the vast majority of the simulation.

The simulation result shows that when the large step speed demand is applied, voltage saturation occurs immediately and the initial magnitude of the voltage vector is high, exceeding the limit V_{max} by 2.5 times. The presence of this saturation event causes the integrator in the outer loop to wind up and the current demand




Figure 6.17: Nonlinear Simulation Condition 2 without OR




Figure 6.18: Nonlinear Simulation Condition 2 with static OR




Figure 6.19: Nonlinear Simulation Condition 2 with dynamic OR




Figure 6.20: Override performance with reduced supply voltage, V_{max} , and Simulation Condition 1

magnitude rises to 10 times that of the limit, extending the period of voltage saturation and the speed tracking performance is significantly degraded as a result.

In the simulations with nominal supply voltage, voltage saturation may occur for short periods of time but the current magnitude limit proves to be the more conservative limit i.e. that which saturates more readily. In these cases, the action of the override compensator constrains the integral state of the outer loop controller and thus protects against this form of windup. In the override simulation with reduced supply voltage (Figure 6.20), the motor current magnitude remains below the limit whilst voltage saturation occurs and so the override compensator does not act, and the windup problem ensues. This analysis reveals that for the existing override compensation structure to be fully robust, the current limit should be more conservative than the voltage limit. This could be ensured by a suitable design of the phase advance map as this apportions the voltage and current headroom, or by varying the current limit in accordance with the measured supply voltage. An example response with the latter solution is given in Figure 6.21 where i_{max} is reduced by 15% to coincide with the 10% reduction in supply voltage. This yields promising results as tracking behaviour is not significantly degraded compared to the nominal system despite the reduction in available power.

Voltage saturation is a complex multivariable problem and the effects of violating the limit on the current angle produced is difficult to control with an anti-windup design, particularly with a linear approach. When anti-windup is applied to the outer loop, a hard constraint on the magnitude of the current demand is imposed and this significantly reduces the severity of voltage saturation as the input to the inner loop is bounded. Thus, with voltage saturation being quite mild, directionality problems tend to be less significant. In the override approach, since there is no limit on the magnitude of the current demand, voltage saturation can be quite severe. This leads to potentially greater directionality errors and may be an additional reason why the current magnitude remains below the limit while voltage saturation persists in the simulation of Figure 6.20. Further work in this area could be fruitful and the design of an advanced, possibly nonlinear, anti-windup compensator for the inner loop that better handles the directionality problem may reduce the susceptibility of the override compensated system to such performance degradation.

6.3.4 Comparison with anti-windup compensation

Here we compare the best override compensator design with the best anti-windup compensator design in order to assess the merits of each approach. We choose the dynamic override design with $k = 0.99$ and compare this to low order dynamic anti-windup design 1 of Chapter 5.4.5 under no load (Figure 6.22).

In the upper left plot it can be seen that the anti-windup compensated system exhibits a faster rise-time. However, from the top right hand plot, we can see that this extra performance is achieved by violating the current limit by close to 20% for a period of approximately 1/10 units of time. This highlights the main benefit of override compensation; that it deals directly with the magnitude of the *actual* current vector rather than its demand. Thus, with the exception of some short transient overshoots, the motor current magnitude can be constrained




Figure 6.21: Override performance with i_{max} reduced in accordance with the supply voltage, V_{max} , and Simulation Condition 1

successfully to the limit. Although the override compensator fails to prevent the initial violation of the current limit, it is arguable that this short duration peak would be less problematic to electronic components than a lower peak of longer duration as observed in the anti-windup response. The motor current responses are very similar but perhaps slightly smoother with override compensation. When comparing the voltage limit violation shown in the bottom right hand plot it is clear that the override compensator causes saturation to cease sooner.

Figure 6.23 shows a comparison of the two systems under a static load condition. In this simulation, the tracking performance is much more comparable with the anti-windup approach having initially a faster rise, and the override approach providing lower steady-state error. Although chatter on the motor voltage is not completely eliminated in the override simulation, the motor currents are quite smooth, violation of the current limit is less prominent, and no oscillations are introduced to the motor speed, unlike with the anti-windup approach.

This comparison shows that override compensation is a better conceptual approach to the problem of current limitation than anti-windup within PMSM speed control. The benefits are that violation of the limit can be controlled more reliably, at least in-principle, that violation of the limit is more transparent, and that greater usage of the available current headroom can be afforded. Having said this, the anti-windup approach was also quite successful and if the observed violation of the limit can be tolerated, a slightly faster rise time can be achieved by using this approach. Another appealing feature of the anti-windup approach is that the constraint on the current demand reduces the size of signals that the inner loop is required to respond to and this may reduce the severity and regularity of voltage saturation.

6.3.5 Conclusion

This chapter has examined the override control problem and its application to PMSM current limitation. It has been shown that the override strategy fits this problem more naturally than anti-windup and that improved confidence regarding limitation of the motor currents can be gained with its use. However, there are certain drawbacks yet to receive a robust solution.

Testing on a continuous time linear model revealed very good results with minimal degradation to tracking performance compared to the linear system and very good limiting of the current. This was particularly true for the static design which was able to respond quickly enough to avoid violation of the limit almost entirely. When applied to the discrete time nonlinear model performance dropped to an extent, with the static design proving to be quite sensitive to ‘chatter’ and neither design being able to suppress transient violations of the limit when large inputs were applied to the system. Testing on the nonlinear model has shown some important differences between the two forms of override compensation and given considerable insight into the practicalities of implementing such control systems.

The stability constraints for the design of a static compensator seem to be less restrictive, allowing for higher gains than their dynamic counterparts. These higher gains along with the lack of phase lag from the dynamic




Figure 6.22: Nonlinear simulation comparison between dynamic AW and dynamic OR performance under no load conditions




Figure 6.23: Nonlinear simulation comparison between dynamic AW and dynamic OR performance under static load conditions

compensator filters result in a more aggressive compensator. This provided better performance in continuous time linear simulation but caused chattery behaviour to propagate through the system when applied to the nonlinear discrete-time model. The dynamic compensator cannot be tuned as aggressively and the phase lag introduced by the filters does restrict its ability to respond to very fast transients. However, this allows for much smoother control with significantly less chatter introduced to the system.

One of the difficulties associated with override compensation is that violation of the limits cannot be avoided entirely and when violation is caused transiently, the limit can be exceeded considerably before the compensator manages to effect a response. Even though the duration of such violations may be very short, the amount by which they are exceeded may be unacceptable. The extent to which this is the case is dependent upon a number of factors including the tune of the nominal controller and the dynamics of the plant. The transient current limit violation caused by applying large step references could not be avoided on the nonlinear discrete time model even with the most aggressive static compensator. It was observed that in order for the compensator to have a swift effect, feedback via ϕ_2 is critical. The gain in this channel, and the bandwidth in the case of a dynamic design, is restricted somewhat by the stability condition. However, the behaviour is influenced even more by the latency introduced by the discrete-time implementation.

To reduce this violation, the sampling frequency of the speed controller could be increased to reduce latency in the system and allow the compensator to act more immediately. This would impose higher computational demands and so may not be desirable. An ad-hoc solution could be achieved by reducing the value of k temporarily when conditions are encountered demanding a sharp increase in the current demand. For instance, when the speed tracking error is beyond a given threshold. Such an approach may prove to be very successful if designed appropriately but analysis of its effect on the system would be difficult. However, due to the use of the Circle Criterion for stability analysis of the override design, time varying limits are automatically accommodated. A more appealing theoretical approach may be to employ a dynamic \mathcal{H}_∞ controller in place of the PI regulator. The lower bandwidth of this type of controller may allow these current spikes to be avoided without sacrificing performance significantly. Of course this would come at the cost of increased computational demands.

Simulation analysis with a reduced voltage limit revealed that for the override approach to mitigate against windup in the outer loop, the current limit must be a tighter constraint than the voltage limit. This could be achieved in a number of ways including altering the phase advance map or reducing the current limit in accordance with the supply voltage. A limited analysis also suggested that the influence of voltage saturation on directionality may contribute further toward degraded performance. This area is ripe for further research and possibilities for the application of advanced, possibly nonlinear, anti-windup compensators to this inner loop problem exist. Other possible approaches would be to solve the inner loop anti-windup in an override framework by passing the d and q axis voltage demands through as dummy plant outputs. If successful, this could allow a unified override compensator to be designed to handle both voltage and current constraints.

Overall, override compensation has been shown to be a viable alternative to the anti-windup approach for limit-

ing the magnitude of currents flowing in the motor windings. In concept, it is a more appealing approach as with the exception of the current transients discussed, it allows the chosen constraint to be imposed more reliably and precisely. However, before the design can be considered preferable to the existing anti-windup solution, additional work is required to reduce the susceptibility to transient overshoots and improve the behaviour of the system in low voltage conditions.

Chapter 7

Improving Anti-Windup Synthesis

Anti-windup compensator design is an inherently nonlinear problem. The saturation nonlinearity at the heart of every control system containing an anti-windup compensator means that the performance of the system is strongly dependent upon the magnitude of the control input, or equivalently, the saturation level. The fundamental problem with nonlinear systems is that the results available for their analysis are not as “sharp” as for linear systems. Stability, for example, is either proved approximately (for example with describing functions), or conservatively (for example with Lyapunov functions). The approach taken in this thesis has been to take a cautious approach and to use the Lyapunov approach to design compensators which enforce both stability and performance properties. This has the advantage of giving *guarantees*, but of course, the guarantees given, especially of performance, may be quite vague and potentially of little use.

For the above reasons, it is sometimes the case - particularly for relatively simple systems - that an anti-windup compensator which is designed “by hand”, perhaps using engineering tuition and intimate system knowledge, can out-perform an anti-windup compensator designed using an “optimal method”. While it is difficult to define exactly what “performance” is in such a context, we take it to mean performance yielded in simulation or experiment during common test conditions. In situations where a good *ad hoc* anti-windup compensator can be found, it is difficult therefore to justify the implementation of a more complex optimisation-based compensator on the grounds of theoretical predictions only. By virtue of their greater degrees of freedom, the modern compensator designs offer the potential of significant performance benefits along with their guarantees of stability. However, the difficulties in tuning in order to exploit their full potential are a significant barrier to their widespread adoption within industry. Therefore, the field of anti-windup is ripe for research into improving tuning methods.

In this chapter, a contribution to the design of anti-windup compensators is made. The main aim of the chapter is to introduce modifications to the design algorithms of anti-windup compensators to allow them to bestow improved levels of performance on the closed-loop system during and after saturation. For simplicity, and also for the most freedom, the tuning modifications are demonstrated on the full-order anti-windup algorithms introduced in Section 4.4. Two approaches to improving the tuning of anti-windup compensators are proposed. The

first is based on altering the objectives on the anti-windup compensation design to include a weighting function which incorporates frequency domain information. The second is based on approximating the description of the saturation (actually deadzone) nonlinearity in order to reduce conservatism in the algorithms.

7.1 A Critique of the Existing Algorithms

This section discusses the key problems inherent in the anti-windup design scheme used so far. The discussion is focused on the full-order coprime factor-based scheme introduced in Section 4.4 [100, 94] but the arguments apply more generally to many similar schemes [65, 96, 12], including low order and static schemes [91]. These schemes share two common features: firstly, that performance is guaranteed using the \mathcal{L}_2 gain; and secondly that stability is enforced using, in essence, the Circle Criterion. Both of these features can be criticised as follows:

- *Circle Criterion.* The Circle Criterion is a well-known and tractable stability test for nonlinear systems but it only gives a *sufficient* condition for stability; that is a system which does not satisfy the Circle Criterion may still be stable. In the work of [100, 94] (and elsewhere) the Circle Criterion is equivalent to using a quadratic Lyapunov function and a sector bounded nonlinearity, $Dz(\cdot) \in \text{Sector}[0, I]$ to enforce stability. While quadratic Lyapunov functions are perhaps the most tractable for LMI-based analysis and synthesis, they form only a small class of possible Lyapunov functions. Similarly, while sector-based descriptions of nonlinearities normally lead to tractable computational procedures, for most functions, the sector-based description is not tight and so also introduces conservatism into the design process.
- \mathcal{L}_2 gain. Performance for nonlinear systems is perhaps even more difficult to define than stability and there are few computationally tractable performance measures available to the engineer. One of the most common is the \mathcal{L}_2 gain which has been used extensively in the anti-windup literature. While it has been shown in [24] that there exists a full order anti-windup compensator which can minimise the \mathcal{L}_2 gain of a (nonlinear) map representing the performance of the anti-windup compensator, it is still debatable whether such a map is a natural measure of an anti-windup compensator's performance. In [100, 91], a particular nonlinear map, \mathcal{T}_p was introduced as being particularly pertinent to the anti-windup problem and it was shown that a coprime factor $G_2 = N(s)M(s)^{-1}$ which minimised this map could be chosen by optimisation of an appropriate gain matrix F . However since the optimisation was still \mathcal{L}_2 -based, the performance measure was inevitably coarse. Further attempts on improving the performance of anti-windup compensators has been made in [93, 33, 102], although again these attempts are centred around the \mathcal{L}_2 -gain.

The work described in this chapter will attempt to address both of the above areas and will derive synthesis algorithms for full-order compensators which, in principle at least, should allow less conservative compensators to be produced. Simulation results will demonstrate the extent to which these modifications succeed.

7.2 Frequency Weighted Anti-Windup

This section proposes a modification to the standard full-order coprime factor based anti-windup synthesis algorithms in order to address the coarseness of the \mathcal{L}_2 -gain optimisation identified above. While research has shown that minimisation of the \mathcal{L}_2 gain is a sensible and successful approach at designing anti-windup compensators, such an approach may not capture all the desirable features of an anti-windup design and, as mentioned above, may lead to below-par performance. One reason for these problems is that the \mathcal{L}_2 gain is optimised over the whole space of energy bounded signals, meaning that there is no distinction between commonly encountered signals and those which are unlikely to be encountered in practical situations.

In this section, we take inspiration from mixed sensitivity \mathcal{H}_∞ control and introduce frequency weighting into the synthesis algorithms. In this, way, it is hoped that the designer can use the frequency weight to target the frequencies of interest and optimise the \mathcal{L}_2 gain delivered by the compensator to those important frequencies.

The results derived were published in [58] and show how such an improvement can be achieved using essentially the same tools as devised in [88] and [87]. The full order anti-windup approach is chosen because the associated optimisation has the most degrees of freedom by which the frequency shaping may take effect. However, in principle, the approach could also be applied to static compensation or low order dynamic compensation such as reported in [91].

7.2.1 A Modified Design Procedure

As with the standard full order anti-windup scheme, we consider the stable linear plant $G(s) = [G_1(s) \ G_2(s)] \in \mathcal{RH}_\infty$ and the stabilising two-degree-of-freedom controller $K(s) = [K_1(s) \ K_2(s)]$ which are described by the following state space models.

$$K(s) \sim \left[\begin{array}{c|c|c} A_c & B_{cr} & B_c \\ \hline C_c & D_{cr} & D_c \end{array} \right] \quad (7.1)$$

$$G(s) \sim \left[\begin{array}{c|c|c} A_p & B_{pd} & B_p \\ \hline C_p & D_{pd} & D_p \end{array} \right] \quad (7.2)$$




Figure 7.1: Augmented plant

In frequency weighted anti-windup we re-label the nominal plant \hat{G} and augment it with a frequency weight as shown in Figure 7.1. In this figure, the physical measurement of the plant output, y is augmented with a

frequency weighted version, y_{W_1} , producing the stacked output vector \tilde{y} . W_1 represents the frequency weight - chosen by the designer - which we use to improve performance. $\hat{G} = [\hat{G}_1 \quad \hat{G}_2]$ represents the real, “physical” plant and $G = [G_1 \quad G_2]$ is the augmented plant (see Fig. 7.1). W_1 and \hat{G}_2 are described by the following state-space equations

$$\hat{G}_2 \sim \begin{cases} \dot{\hat{x}}_p = \hat{A}_p \hat{x}_p + \hat{B}_p u_m \\ y = \hat{C}_p \hat{x}_p + \hat{D}_p u_m \end{cases} \quad (7.3)$$

$$W_1 \sim \begin{cases} \dot{x}_1 = A_1 x_1 + B_1 y \\ y_{W_1} = C_1 x_1 + D_1 y \end{cases} \quad (7.4)$$

and hence by making the definitions $x = [\hat{x}'_p \quad x'_1]'$, and $\tilde{y} = [y' \quad y'_{W_1}]'$, we have

$$G_2 \sim \begin{cases} \dot{x}_p = \underbrace{\begin{bmatrix} \hat{A}_p & 0 \\ B_1 \hat{C}_p & A_1 \end{bmatrix}}_{A_p} x + \underbrace{\begin{bmatrix} \hat{B}_p \\ B_1 \hat{D}_p \end{bmatrix}}_{B_p} u_m \\ \tilde{y} = \underbrace{\begin{bmatrix} \hat{C}_p & 0 \\ D_1 \hat{C}_p & C_1 \end{bmatrix}}_{C_p} x + \underbrace{\begin{bmatrix} \hat{D}_p \\ D_1 \hat{D}_p \end{bmatrix}}_{D_p} u_m \end{cases} \quad (7.5)$$

$M(s)$ and $N(s)$ are chosen again as coprime factors of G_2 which is now the “augmented” plant. The decoupled representation of the anti-windup closed loop using the augmented plant is shown in Figure 7.2. Note that y_{lin} and y_d have now become the vector signals $\tilde{y}_{lin} = [y'_{lin} \quad y'_{W_1,lin}]'$ and $\tilde{y}_d = [y'_d \quad y'_{W_1,d}]'$ in which $y_{W_1,lin}$ and $y_{W_1,d}$ represent frequency weighted versions of the linear plant output and disturbance filter output respectively.

Whereas in the standard full-order synthesis [87, 88] the \mathcal{L}_2 norm of the map from u_{lin} to y_d (i.e. $\|\mathcal{T}_p\|_{i,2}$) was minimised, we now propose to minimise the frequency weighted version of this, i.e. the map from u_{lin} to $y_{W_1,d}$, using the state-space realisation given in equation (7.5). Let us call this map \mathcal{T}_{W_1} . This can be achieved by ensuring that (7.6) holds, or equivalently, (7.7) where W_f has the form of (7.8).

$$\left\| W_f^{\frac{1}{2}} \tilde{y}_d \right\|^2 < \gamma^2 \|u_{lin}\|^2 \quad (7.6)$$

$$\tilde{y}'_d W_f \tilde{y}_d - \gamma^2 u_{lin}' u_{lin} < 0 \quad (7.7)$$

$$W_f = \begin{bmatrix} \epsilon I & 0 \\ 0 & \tilde{W}_p \end{bmatrix} \quad (7.8)$$

Typically ϵ would be chosen as a small number to ensure that y_d would have little impact on the design and \tilde{W}_p is a weighting matrix used to trade-off the importance of the various channels in y_{W_1} . Note that such a form

of compensator fits exactly into the standard framework, except the plant has an additional (fictitious) output and the controller has an additional (fictitious) input as illustrated in Figure 7.2. In this figure, the “augmented” controller is denoted $K(s) = [\hat{K}(s) \ 0]$ and has a null input column representing the zero contribution from y_{W1} in the actual system. In this figure the standard anti-windup compensator architecture is used except




Figure 7.2: Augmented plant and controller

now $M(s)$ and $N(s)$ are part of a coprime factorisation of the augmented plant described in (7.5). As with the standard compensator this means that a coprime factorisation can be given by (7.9) where F is a free parameter.

$$\begin{bmatrix} M(s) - I \\ N(s) \end{bmatrix} \sim \begin{cases} \dot{x} = (A_p + B_p F)x + B_p \tilde{u} \\ u_d = Fx \\ y_d = (C_p + D_p F)x \end{cases} \quad (7.9)$$

In a similar manner to standard full-order compensator synthesis, except using the state-space realisation of (7.5), an AW compensator which guarantess stability and \mathcal{L}_2 gain of the map $\tilde{T}_{W1} : u_{lin} \mapsto W_f^{\frac{1}{2}} \tilde{y}_d$ is produced if the inequality of (7.10) holds.

$$\dot{V}(x) + 2\tilde{u}'W(\tilde{u} - u) + \tilde{y}_d'W_f\tilde{y}_d - \gamma^2 u_{lin}'u_{lin} < 0 \quad (7.10)$$

Some algebra then yields the LMI:

$$\begin{bmatrix} QA'_p + L'B'_p + A_pQ + B_pL & B_pV - L' & 0 & QC'_p + L'D'_p \\ * & -2V & I & VD'_p \\ * & * & -\gamma I & 0 \\ * & * & * & -\gamma W_f^{-1} \end{bmatrix} < 0 \quad (7.11)$$

Note here however that because $M(s)$, $N(s)$ are coprime factors of the augmented plant (i.e. that including the weight $W_1(s)$), the order of the anti-windup compensator is greater than that of the real plant, $\hat{G}(s)$. With this in mind, the anti-windup compensator is constructed using $F = LQ^{-1}$ and (7.9).

7.2.2 Case Studies

Here the performance of a number of systems with plant input saturation is compared for standard and frequency weighted designs. Some of the analysis makes use of the description of anti-windup modes [100] where the behaviour of the nonlinear system is divided into three modes:

Mode I - This represents normal linear operation i.e. when saturation limits have not been exceeded.

Mode II - This represents the period for which the system's actuators are saturated. Equivalently it can be thought of as the period of time between \tilde{u} becoming non-zero, and returning to (and remaining at) zero.

Mode III - This represents the final convergence to linear performance, which occurs after saturation has ceased and in which the states of the disturbance filter relax asymptotically to zero.

RC Circuit model

The model of an electrical network taken from [24] and depicted in Figure 7.3 is used. V_i is the plant input voltage, and V_o is the plant output voltage. The state space model of the plant and a robustly stabilising PI controller are given as (7.12) and (7.13) respectively. Note that the matrices of the plant output equation are all negative, hence the plant output is in the opposite direction to that of the applied control signal.

$$\hat{G}_2(s) \sim \left[\begin{array}{ccc|c} -10.6 & -6.09 & -0.9 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ \hline -1 & -11 & -30 & 0 \end{array} \right] \quad (7.12)$$

$$\hat{K}(s) \sim \left[\begin{array}{cc|cc} -80 & 0 & -1 & 1 \\ 1 & 0 & 0 & 0 \\ \hline 20.25 & 1600 & -80 & 80 \end{array} \right] \quad (7.13)$$

For this example, the performance of the anti-windup compensator is limited by slow poles in the open loop plant, which form the basis of the poles of the disturbance filter in the compensator. This causes mode III to be significantly longer than mode II, as shown in Figure 7.4. Inspection of the step response of the standard anti-windup closed loop system (Figure 7.4) reveals that the control performance during mode II cannot be improved upon, as for the duration of mode II, the plant input is saturated in an effort to attain the reference demand. Therefore, any performance gains to be achieved must take effect during mode III. During mode




Figure 7.3: Electrical Network

III, u_{lin} is a low frequency signal and hence an intelligent choice for the frequency weighting filter may be a low pass filter of an appropriately low bandwidth. Figure 7.4 shows a comparison between the standard compensator and that designed using the frequency weight

$$W_1(s) = \frac{0.35}{s + 0.35} \tag{7.14}$$

The application of frequency weighting gives rise to a swifter return to linear performance through its effect during mode III and the performance during mode II is unaffected.




Figure 7.4: RC Circuit Step Response

Missile example

A simplified model of the roll-yaw channel dynamics of a bank to turn missile and LQG/LTR type auto-pilot controller given in [74, 91] is selected for which the linear plant and controller models are given in (7.15) and

(7.16). The two outputs of the simplified plant are roll and yaw and thus due to the bank-to-turn nature of the missile, to alter the heading of the missile a response is required in both of these channels.

$$\hat{G}_2(s) \sim \left[\begin{array}{ccc|cc} -0.8 & -1 & 0.3 & 0.1 & 0 \\ 80.3 & -0.6 & 0 & -194.4 & 37.6 \\ -2734 & 0.1 & -2.1 & -2716 & -1093 \\ \hline 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \end{array} \right] \quad (7.15)$$

$$\hat{K}(s) \sim \left[\begin{array}{ccccccc|cc} -0.3 & -107.8 & 6.67 & -2.6 & -0.4 & 2.3 & 0.5 & 0 & 0 \\ 107.7 & -97.8 & 64.0 & -4.5 & -5.4 & -40.8 & 2.1 & 0 & 0 \\ -6.7 & 64.8 & -54.2 & -40.8 & 5.1 & 18.5 & -0.2 & 0 & 0 \\ 3.2 & 2.1 & 29.6 & -631.2 & 429.9 & -2.1 & -44.7 & 0 & 0 \\ 0.4 & -3.4 & 3.1 & -460.0 & -0.74 & -0.98 & -1.2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ \hline 0.9 & 8.5 & -1.7 & 43.9 & 1.1 & 0 & 0 & 0 & 0 \\ 2.2 & 39.9 & -18.4 & -8.5 & 1.0 & 0 & 0 & 0 & 0 \end{array} \right] \quad (7.16)$$

As shown in Figure 7.5, the standard full order compensator performs very well with the result that linear performance is almost completely recovered in the second channel. The main problem with this design is the undesirable oscillation which appears during mode III when the energy stored in the compensator is dissipated through the disturbance filter. In order to suppress these oscillations a shelving type high pass filter was used to penalise these frequencies in the \mathcal{L}_2 gain optimisation. Figure 7.6 shows how the oscillations in the output are attenuated as a result of the frequency weighted design using the filter described by (7.17). This performance improvement does come at the cost of larger compensator poles but the frequency weight can be altered to give a compromise between performance and pole magnitude.

$$W_1(s) = \frac{s + 19}{s + 30} \quad (7.17)$$

It is found that using a low pass filter frequency weight can be useful in obtaining closer reference tracking during saturation. A first order low pass filter described by (7.18) results in the performance shown in Figure 7.7. Note that while the tracking is slightly closer during mode II, the oscillations during mode III are of greater amplitude as they are beyond the bandwidth of the filter. In this example it is arguable that the costs outweigh the benefits but it does demonstrate the concept.

$$W_1(s) = \frac{0.1s + 0.1}{s + 0.1} \quad (7.18)$$




Figure 7.5: Missile autopilot doublet response - standard AW




Figure 7.6: Missile autopilot doublet response - HPF frequency weighted AW




Figure 7.7: Missile autopilot doublet response - LPF frequency weighted AW

7.2.3 Remarks on Tuning the Compensator

The extra freedom in this method compared to other anti-windup methods is the flexibility in choosing the weight W_1 , allowing explicit frequency weighting of the design. However, in contrast to standard \mathcal{H}_∞ design, there may be some difficulty in tuning the weight W_1 . There are several reasons for these difficulties:

- There is some conservatism in the method. A strong influence on the success of frequency shaping is imposed by the \mathcal{H}^∞ norm of the open-loop plant [79, 24], which can be large for some systems. As a consequence, the values of γ which can be obtained, are therefore much greater than unity and thus our specifications are not guaranteed to be satisfied¹.
- In the work to date, there seems to be significant “trial and error” required in the choice of the weights. In some systems, low pass filters are useful; in others, high pass filters are more appropriate. Other than observation of the system’s response, there does not seem to be a universal method of choosing the weight $W_1(s)$. This requires further investigation.

¹In standard \mathcal{H}_∞ optimisation such as S/KS design, a γ of unity indicates that our design specifications are guaranteed to have been met; anything greater than this does not provide this guarantee

7.2.4 Conclusions and Future Work

A method of tuning full-order AW compensators for improved performance by shaping an \mathcal{L}_2 gain optimisation over frequency has been presented. Through a number of case studies, this method has been shown to give the potential of improved tracking performance and/or the attenuation of undesirable frequency content in the output of the nonlinear system following saturation.

The order of the compensator produced using this approach is increased compared to standard designs by the order of the filter used. Although this is undesirable, in practice, Hankel model reduction could be employed to reduce the compensator order whilst maintaining the performance benefit. For such a reduced order design, a subsequent stability check would be required as it would no longer fit the coprime factorisation framework. However, since stability could be considered in isolation in this case, this would allow the use of less conservative stability guarantees if necessary, for example the multivariable Popov criterion.

It would be ideal if an alternative method could be developed to weight the \mathcal{L}_2 gain optimisation without introducing extra states. This would facilitate the use of higher order band-pass, shelving or notch filters in the frequency weighted design without additional cost. This could be achieved trivially for any optimal compensation strategy in which the compensator order is independent of the plant model. One such possibility is the static anti-windup approach reported in [91] although the lack of dynamics in the compensator may limit the effect that any frequency weighting may have. Similarly, for the low-order dynamic approach given in the same paper the optimisation is only able to influence a matrix gain in series with the filter dynamics chosen by the designer. It is thought that for the low order approach, the freedom that the designer has in choosing the filter dynamics manually may overshadow any subtle variations in the optimisation result caused by frequency weighting. Therefore the most fruitful outcomes are likely to be achieved either by incorporating a frequency weight in such a way as not to increase the order of the compensator, or by application to a dynamic anti-windup approach that does not make use of coprime factorisation of the plant.

Although performance improvement has been shown, this is not observed for all systems and the performance observed can be quite sensitive to small changes in the frequency weighting filter. For the examples presented here, the improvement is fairly modest, although this approach is expected to be more fruitful for higher order, complex systems. A practically verified example of this is the discrete anti-windup compensator of [35] which is designed for a hard disk servo system using a high pass type frequency weight. It appears that a good baseline compensator design is given using the standard full-order optimisation, but tuning for improved performance is complex. Frequency weighting is now one of a number of tools available to the designer with which to fine tune a successful design.

7.3 Reduced Sector Bounds

For the modern anti-windup compensator designs presented in Section 4.4, the design approach ensures that the nonlinear closed loop is globally absolutely stable, i.e. stable for all possible inputs to the nonlinearity and

all nonlinearities bound by the sector, which in the Weston and Postlethwaite framework [100], is the deadzone function. Since the sector bound is a conservative definition of the deadzone function, the Circle Criterion may be a stronger condition than required to provide stability and some conservatism is introduced into the resulting design. This means that the ‘aggressiveness’ of the compensator may be restricted to ensure satisfaction of the stability condition. If a representation of the deadzone nonlinearity can be used that more tightly defines the nonlinearity, there will be more freedom in the optimisation and hence a more aggressive compensator could be produced, potentially improving performance. This follows since the set of all possible nonlinearities bounded by the sector, for which stability must be guaranteed, is reduced when the size of the sector is reduced.

Firstly, consider the scalar deadzone function defined by (7.19) where $-\bar{u}$ and \bar{u} represent the bounds of the deadzone region where the gain is zero.

$$Dz(u) = \begin{cases} u - \bar{u} & \forall u \geq \bar{u} \\ 0 & \forall u \in (-\bar{u}, \bar{u}) \\ u + \bar{u} & \forall u \leq -\bar{u} \end{cases} \quad (7.19)$$

The function is adequately bound by the Sector $[0, 1]$, meaning that it lies between the u axis and a line of gradient equal to one on the graph of $Dz(u)$ against u depicted in Figure 7.8. However, notice that for inputs u less than a value u_{max} , the deadzone function could be adequately bound by the Sector $[0, \beta]$ where $\beta \in (0, 1)$ is a positive constant greater than zero. The relationship between u_{max} and β is dependent upon the saturation limit, \bar{u} , and is given by (7.20). This is shown in Figure 7.8 as the sector between the u axis and a line of gradient β .




Figure 7.8: Reduced sector $[0, \beta]$

$$u_{max} = \frac{\bar{u}}{1 - \beta} \quad (7.20)$$

More generally, if we consider the multivariable deadzone function

$$Dz(u) = \begin{bmatrix} Dz_1(u_1) \\ Dz_2(u_2) \\ \vdots \\ Dz_m(u_m) \end{bmatrix}, \quad Dz_i(u_i) = \begin{cases} u_i - \bar{u}_i & \forall u_i \geq \bar{u}_i \\ 0 & \forall u_i \in (-\bar{u}_i, \bar{u}_i) \\ u_i + \bar{u}_i & \forall u_i \leq -\bar{u}_i \end{cases} \quad (7.21)$$

it follows that, if we define

$$u_{max,i} = \frac{\bar{u}_i}{1 - \beta_i} \quad (7.22)$$

each component of the deadzone nonlinearity locally inhabits Sector $[0, \beta_i]$ for all $u_i \leq \beta_i u_{max,i}$. If we now let

$$\beta = \max_i \beta_i \quad (7.23)$$

it then follows that $Dz \in \text{Sector}[0, \beta I]$ for all $u \in \mathcal{U}_{max}$, where

$$\mathcal{U}_{max} := \left[-\frac{\bar{u}_1}{1 - \beta}, \frac{\bar{u}_1}{1 - \beta}\right] \times \dots \times \left[-\frac{\bar{u}_m}{1 - \beta}, \frac{\bar{u}_m}{1 - \beta}\right] \subset \mathbb{R}^m \quad (7.24)$$

$$= [-u_{1,max}, u_{1,max}] \times \dots \times [-u_{m,max}, u_{m,max}] \subset \mathbb{R}^m. \quad (7.25)$$

It then follows that for all $u \in \mathcal{U}_{max}$ the *reduced sector condition* holds:

$$Dz(u)'W(\beta u - Dz(u)) \geq 0 \quad (7.26)$$

where $W > 0$ is a diagonal matrix to be defined.

The effect of reducing β from 1 can be seen by considering the graphical interpretation of the circle criterion in the SISO case as follows. Consider the closed loop interconnection of a linear system, $G(s)$, with the nonlinearity, ψ , bound by the Sector $[\alpha, \beta]$ and $\beta > \alpha \geq 0$. The circle criterion states that the nonlinear closed loop system is globally absolutely stable if the Nyquist locus $G(j\omega) \forall \omega \in (-\infty, \infty)$ does not penetrate the disc centred on the real axis with intercepts at $[\frac{-1}{\beta}, \frac{-1}{\alpha}]$ and encircles it as many times anticlockwise as $G(s)$ has unstable poles. Considering stable plants only and that for the sector bound sector $(0, \beta)$ defined in 7.8, $\alpha = 0$, the disc has infinite diameter and the requirement for stability is simply that the Nyquist locus must not pass to the left of the line $\Re(s) = -1/\beta$. A simple example of this is shown in Figure 7.9, where the Nyquist locus does not conform to the circle criterion with the standard sector. With a reduced sector $(0, \beta)$, the prohibited region for the Nyquist contour moves in the negative direction along the real axis, allowing the circle criterion to be satisfied. For anti-windup design, this leads to greater flexibility in the design of the linear system model whilst satisfying the stability guarantee. Of course, the stability guarantee only holds to the extent that the nonlinearity is bounded by the chosen sector. If the sector bound only holds locally, only local stability is guaranteed. For the deadzone nonlinearity depicted in Figure 7.8, the nonlinear closed loop is guaranteed to be absolutely stable for any $u \in (-u_{max}, u_{max})$. To assess the benefits of using reduced sectors in anti-windup synthesis, application is made to full-order anti-windup compensation, for which an LMI effecting anti-windup synthesis with a user-defined sector bound is derived.




Figure 7.9: Simple example of the graphical circle criterion bounds on the Nyquist contour for sectors $(0, I)$ and $(0, \beta)$

7.3.1 Reduced Sector Full-Order Anti-Windup Compensator LMI Derivation

For the synthesis of a full order anti-windup compensator, we require a sector bound and Lyapunov function candidate by which the Circle Criterion is satisfied, and an \mathcal{L}_2 gain condition by which performance of the nonlinear system is optimised. The following outlines equivalent conditions for the case where the sector bound $(0, I)$ is replaced by a smaller sector $(0, \beta I)$, albeit with the loss of global stability. For convenience of notation, a compact set $\mathcal{U} \in \mathbb{R}^m$ is defined in (7.27) such that when the controller output in the absence of saturation, u_{lin} , belongs to this set none of the plant inputs are saturated.

$$\mathcal{U} := [-\bar{u}_1, \bar{u}_1] \times \dots \times [-\bar{u}_m, \bar{u}_m] \quad (7.27)$$

The goal of anti-windup synthesis is much the same as in the standard full-order case, except that we enforce stability and \mathcal{L}_2 gain over the *reduced* Sector $[0, \beta I]$. The idea behind this is that, as such a sector bounds fewer nonlinearities, it will be a less conservative description and will therefore allow more aggressive anti-windup compensators to be synthesised. The derivation of the algorithm proceeds in a similar way to the standard full order case. We wish to ensure that $\|\mathcal{T}_p\|_{i,2}$ is minimised, but over a smaller sector. Similarly we wish again to ensure quadratic stability of the nonlinear system which means finding a Lyapunov function $V(x) > 0$ such that $\dot{V}(x) < 0$ when $u \in \mathcal{U}_{max}$. These requirements can be captured by the inequality:

$$\dot{V}(x) + y'_d y_d - \gamma^2 u'_{lin} u_{lin} + 2\tilde{u}' W [\beta u - \tilde{u}] < 0 \quad (7.28)$$

$$\frac{d}{dt}(x' P x) + y'_d y_d - \gamma^2 u'_{lin} u_{lin} + 2\tilde{u}' W [\beta u - \tilde{u}] < 0 \quad (7.29)$$

where $P > 0$ and $W > 0$ are diagonal matrices to be defined. Substituting for $u = u_{lin} - u_d$, y_d and x from

(4.20) it then follows that inequality (7.29) can be written as

$$\begin{bmatrix} x' \\ \tilde{u}' \\ u'_{lin} \end{bmatrix}' \begin{bmatrix} (C_p+D_pF)'(C_p+D_pF)+(A_p+B_pF)'P+P(A_p+B_pF) & (C_p+D_pF)'D_p+PB_p-\beta F'W & 0 \\ * & D'_pD_p-2W & \beta W \\ * & * & -\gamma^2I \end{bmatrix} \begin{bmatrix} x \\ \tilde{u} \\ u_{lin} \end{bmatrix} < 0$$

Noting the presence of nonlinear (bilinear) terms, the Schur Complement can be applied to obtain the inequality

$$\begin{bmatrix} A'_p + F'B'_pP + PA_p + PB_pF & PB_p - \beta F'W & 0 & C'_p + F'D'_p \\ * & -2W & \beta W & D'_p \\ * & * & -\gamma^2I & 0 \\ * & * & 0 & -I \end{bmatrix} < 0 \quad (7.30)$$

The matrix inequality of (7.30) is still nonlinear due to cross products between F and P , and also between F and W . These are removed by applying congruence transformations i.e. pre and post multiplying by a constant matrix which in this case is $diag(P^{-1}, W^{-1}, I, I)$. The resulting matrix maintains the same definiteness properties, hence the inequality of (7.31) also implies (7.30).

$$\begin{bmatrix} P^{-1}A'_p + P^{-1}F'B'_p + A_pP^{-1} + B_pFP^{-1} & B_pW^{-1} - \beta P^{-1}F' & 0 & P^{-1}C'_p + P^{-1}F'D'_p \\ * & -2W^{-1} & \beta I & W^{-1}D'_p \\ * & * & -\gamma^2I & 0 \\ * & * & 0 & -I \end{bmatrix} < 0 \quad (7.31)$$

By re-labelling variables $Q = P^{-1} > 0$, $V = W^{-1} > 0$, and applying a change of variables, $FQ = L$, we are left with the following LMI in $Q > 0$, $diag(V) > 0$, and $\gamma > 0$ (7.32):

$$\begin{bmatrix} QA'_p + L'B'_p + A_pQ + B_pL & B_pV - \beta L' & 0 & QC'_p + L'D'_p \\ * & -2V & \beta I & VD'_p \\ * & * & -\gamma I & 0 \\ * & * & 0 & -\gamma I \end{bmatrix} < 0 \quad (7.32)$$

F can then be recovered as $F = LQ^{-1} = LP$. Note that by setting $\beta = 1$ the standard sector bound is restored and the LMI yields standard full-order compensator designs for which global stability is guaranteed.

With the extra freedom allowed by using the reduced sector it is expected that performance will be improved locally i.e. improved when $|u| \leq |u_{max}|$. Any such improvement is gained at the expense of the global stability guarantee. However, due to the inherent conservatism of the Circle Criterion the domain of attraction for which the nonlinear system is locally stable may still be larger than the region in which the sector bound holds.

7.3.2 Reduced Sector Case Studies

Using the LMI of (7.32) and the solver present in the Mathworks Robust Control Toolbox, anti-windup compensators are designed for a selection of academic examples found in the anti-windup literature.

RC Electrical Network

The RC electrical network introduced in Section 7.2.2 is re-used, for which responses of the unconstrained system and the system constrained by magnitude limits $\pm\bar{u}$ are shown in Figure 7.10 where $\bar{u} = 1$. As the dynamics of the RC network are quite slow, high gains are used in the controller to improve the closed loop behaviour and these result in vary large control signals for the linear simulation; in the order of 200V. When compared to the saturation limits applied at $\pm 1V$, it is apparent that severe saturation will result.

The application of full order dynamic anti-windup is successful and greatly improves the behaviour of the nonlinear system as shown in Figure 7.11. By designing with a reduced sector bound defined by $\beta = 0.8$, performance is improved further still (Figure 7.11). The anti-windup compensator performance level, γ , the associated parameterising matrix, F , and compensator poles are given in (7.33-7.36) where the subscript of F denotes the sector size used. For the choice of $\beta = 0.8$ the range for which the sector condition bounds the deadzone nonlinearity is given by $u \in [-5, 5]$. In the simulation results shown, the control signal magnitude does exceed these limits quite considerably and so stability is not guaranteed, although it is observed in practice. The largest pole of the reduced sector compensator is significantly reduced in magnitude compared to that of the standard design, and so is more desirable for practical implementation. The performance improvement observed is accounted for by an increase in the magnitude of the dominant slow pole by 11%.




Figure 7.10: RC circuit unconstrained and constrained responses without anti-windup

$$F_{(0,I)} = [-3182.7 \quad -33378 \quad -15559], \quad \gamma = 33.333 \quad (7.33)$$

$$\lambda_i(A_p + B_p F_{(0,I)}) = \begin{bmatrix} -3182.8 \\ -9.9984 \\ -0.48894 \end{bmatrix} \quad (7.34)$$

$$F_{(0,0.8)} = [-1310.9 \quad -13816 \quad -7093.2], \quad \gamma = 0.44725 \quad (7.35)$$

$$\lambda_i(A_p + B_p F_{(0,0.8)}) = \begin{bmatrix} -1311 \\ -9.9978 \\ -0.54125 \end{bmatrix} \quad (7.36)$$

An interesting point to note is that the \mathcal{L}_2 gain bound is significantly influenced by the size of the sector. This is shown in Figure 7.12 where the performance level for the full order compensator design is plotted as a function of β , which defines the size of the sector. The shape was found to be typical of all systems tested although the extent to which the \mathcal{L}_2 gain changes does vary between systems.

The theoretical minimum \mathcal{L}_2 gain for the full sector is the \mathcal{H}_∞ norm of the open loop linear system so would not drop below this value when using a conventional design. However, when the reduced sector is used much lower gains than this are returned in the optimisation. The actual \mathcal{L}_2 gain value returned therefore is more difficult to interpret but its value relative to other optimisation results for the same system and sector bound is still useful.

Missile Autopilot with LQG Control

The well-known model of a bank to turn missile introduced in Section 7.2.2 is re-used and a simulation was set up in which both inputs to the plant are magnitude limited to ± 8 degrees. The unconstrained and constrained without anti-windup simulation responses to a doublet demand applied to both channels is shown in Figure 7.13. The controller performs well with the linear (unconstrained) system, with the system responding very smoothly to track the reference and excellent decoupling between the channels. Note also that in order to track the reference, the majority of the control effort is applied in a single actuator channel. When the saturation constraints are applied, decoupling between the roll and yaw channels is lost and performance is degraded significantly.

A full order dynamic anti-windup compensator is designed for this system using the standard sector bound and performance and robustness weights equal to the identity matrix. The compensator produced is defined in (7.37), and the corresponding simulation performance is shown in Figure 7.14. Note that the subscript of F denotes the sector bound used for the design to aid the distinction between different designs. Due to the restriction on control energy imposed by the constraints, neither channel is able to meet the set-point, however each channel tracks quite closely to its reference. Notice though that when saturation has ceased and the system begins recovery to linear behaviour (Mode III), both output channels exhibit a somewhat oscillatory return to zero. Recall that during Mode III, in the full-order case, the output of the disturbance filter is governed by $N(s)$




Figure 7.11: RC circuit responses with full order AW




Figure 7.12: L_2 gain variation with sector bound for RC circuit example




Figure 7.13: Missile doublet response for the unconstrained model and the constrained model without anti-windup

and is thus $y_d(t) = C_p \exp(A_p + B_p F) \tilde{x}(0)$. Hence the poles of the full order compensator are central to the type of return to linear behaviour observed. The compensator poles given in (7.38) are underdamped and thus explain the oscillations observed in Figure 7.14.

$$F_{(0,I)} = \begin{bmatrix} 2.2795e-2 & 1.7686e-1 & 4.5274e-3 \\ -1.7586e-3 & -3.1289e-3 & -2.3714e-4 \end{bmatrix}, \quad \gamma = 376.59 \quad (7.37)$$

$$\lambda_i(A_p + B_p F_{(0,I)}) = \begin{bmatrix} -1.1589 + 32.482j \\ -1.1589 - 32.482j \\ -47.713 \end{bmatrix} \quad (7.38)$$

The control signals in the simulation without saturation violate the saturation level by 25% which could be enclosed by a sector bound on the deadzone function of $[0, 0.2]$. This sector bound is chosen as an initial trial and if the control signals do not exceed this level in the anti-windup compensated system, we might expect to see the greatest available benefit of using a reduced sector. Using LMILAB as the solver and $\beta = 0.2$, the compensator produced is given by (7.39). Performance in simulation is improved over the standard full order design, with closer tracking of the reference and a smoother response (Figure 7.15). This is explained partially by the improved damping of the compensator, shown in (7.40). In the constrained simulation, the magnitude of the unconstrained control signals (not shown) do remain less than u_{max} and hence the stability condition does hold. However, for larger inputs this may not be the case.

$$F_{(0,0.2I)} = \begin{bmatrix} 1.6956 & 9.2176 & 0.33344 \\ -0.024106 & -0.3306 & 0.025955 \end{bmatrix}, \quad \gamma = 1.0328 \quad (7.39)$$

$$\lambda_i(A_p + B_p F_{(0,0.2I)}) = \begin{bmatrix} -2697.8 \\ -21.877 + 23.503j \\ -21.877 - 23.503j \end{bmatrix} \quad (7.40)$$

When investigating the effect of using other sector bounds, some unexpected results were found. It was thought that for a system giving stable results with the Sector $[0, I]$, performance would progressively increase as the sector was reduced up until the point at which the bound is violated and as a consequence, stability no longer guaranteed. However, poor performance was also observed for some sector bounds where use of a smaller or larger sector did provide good performance. An example of this is the Sector $[0, 0.8]$. A compensator designed using this sector and the same performance and robustness weighting matrices is defined in (7.41). This design causes limit cycles in simulation when the same reference demand sequence is applied as shown in Figure 7.16. In this simulation the magnitude of the unconstrained control signals *do* exceed u_{max} and so the stability guarantee does not hold. This is understood to be the cause of the limit cycle behaviour, although, the very large real pole shown in (7.42) may also be causing difficulties for the simulation solver, and this is likely to be impractical to implement.




Figure 7.14: Missile simulation with standard sector bound. - - Nominal linear, – Saturated + AW




Figure 7.15: Missile simulation with reduced sector AW. - - Nominal linear, – Saturated + AW




Figure 7.16: Missile simulation with reduced sector AW that causes limit cycles

$$F_{(0,0.8I)} = \begin{bmatrix} 15.859 & 84.607 & 3.0364 \\ -0.2424 & -2.7334 & 0.24027 \end{bmatrix}, \quad \gamma = 1.4172 \quad (7.41)$$

$$\lambda_i(A_p + B_p F_{(0,0.8I)}) = \begin{bmatrix} -24693 \\ -3.6955 \\ -364.04 \end{bmatrix} \quad (7.42)$$

Increasing the magnitude of the robustness weight in the optimisation can cause the compensator poles to be reduced in magnitude. When applied to the design with sector bound $[0,0.8]$, performance equivalent to that achieved with sector $[0, 0.2]$ is provided and due to the larger sector, u_{max} is also increased significantly. In this case, $u_{max} = 40$ and so local stability is not restrictive, particularly for this system as the LQG controller does not cause the control signals to saturate heavily. Clearly, the impact on the optimisation result of reducing the size of the sector bound is not trivial, but improvements to performance can be gained.

Lightly damped system with PI control

This next case study uses the lightly damped plant and PI controller of (7.43) taken from [23] and a magnitude saturation limit, $\bar{u} = 0.5$. A nominal linear (unconstrained) response and constrained response without anti-windup are shown in Figure 7.17. Whether or not the nominal linear response should be considered to be *good*




Figure 7.17: Lightly damped system responses without AW

is subjective but for the purposes of anti-windup design we seek a compensator that minimises the deviation from the intended linear response. The saturated response is oscillatory due to the lightly damped poles of the plant and although the system remains stable, the settling time is long and the control response of the saturated system is characterised by large amplitude oscillations.

$$K(s) \sim \left[\begin{array}{c|c} 0 & -1 \\ \hline 2 & -2 \end{array} \right], \quad G(s) \sim \left[\begin{array}{cc|c} -0.2 & -0.2 & 1 \\ 1 & 0 & 0 \\ \hline -0.4 & -0.9 & -0.5 \end{array} \right] \quad (7.43)$$

A full order compensator is designed for this system using the Sector $[0, I]$, performance weight $W_p = 1$ and robustness weight $W_r = 0.1$ and is defined by (7.44).

$$F_{(0,I)} = \left[\begin{array}{cc} -99.777 & -11.5585 \end{array} \right], \quad \gamma = 10.7602 \quad (7.44)$$

$$\lambda_i(A_p + B_p F_{(0,I)}) = \left[\begin{array}{c} -99.859 \\ -0.11775 \end{array} \right] \quad (7.45)$$

The behaviour of the system with this anti-windup compensator applied is shown in Figure 7.18. Although this

design is successful in bringing the system out of saturation and causing linear behaviour to be regained, the return to linear behaviour is not as smooth or swift as one might hope, perhaps due to the slow pole present in the compensator. The response of an additional full order design using the Sector $[0, 0.6]$ defined by (7.46) is also shown in Figure 7.18. This compensator was designed using the same performance and robustness weights as previously and although the compensated response still is not smooth, the deviation from linear behaviour is reduced: the recovery to linear behaviour is a little swifter, partly due to the increase in magnitude of the compensator's slow pole. An important point to note about this example is that although the control signal magnitude in the linear simulation was not greater than 2, in the constrained simulation without anti-windup the magnitude of the control signal exceeded 8. Finally, in the anti-windup compensated simulations the peak control signal magnitude approaches 70.

$$F_{(0,0.6)} = \begin{bmatrix} -55.6844 & -11.3833 \end{bmatrix}, \quad \gamma = 0.42384 \quad (7.46)$$

$$\lambda_i(A_p + B_p F_{(0,0.6)}) = \begin{bmatrix} -55.6763 \\ -0.2080 \end{bmatrix} \quad (7.47)$$

7.3.3 Discussion

These case studies show that use of a reduced sector bound can lead to improved anti-windup performance. There are also examples for which the reduced sector affords no improvement, although no cases have been found for which reduced sectors result in degraded local performance. By local performance, we mean performance of the system whilst the sector adequately bounds the nonlinearity. Of course, for control signals that cause the output of the nonlinearity to exceed the sector bound, the stability condition is not satisfied and therefore neither performance or stability can be assured. This region of operation has been observed with each of the case studies but only in the missile autopilot system did this cause detrimental performance.

A critical consideration in the use of reduced sectors is whether or not the sector is large enough to describe the nonlinearity for the control signals that will feature during the systems operation. In this respect the case studies show some interesting characteristics. With the RC circuit, the control signals of the linear unconstrained system greatly exceeded the magnitude saturation limits that were to be applied and so there was little likelihood of being able to reduce the size of the sector significantly and still bound the nonlinearity. As shown in Table 7.1 the peak control signal magnitude, u_{pk} , actually increased when the constraints were applied, and increased further still when anti-windup was applied. However, although the sector bound did not hold globally and stability was only guaranteed locally, the system was observed to be stable even for large deviations from the sector bound and a small performance benefit was gained.

For the lightly damped system with PI control, the magnitude of the control signal in the linear simulation and the constrained simulation without anti-windup were relatively small. This suggested that an anti-windup




Figure 7.18: Lightly damped system responses with full order AW

Model	\bar{u}	u_{pk}		
		Linear	No AW	Full-order AW
RC circuit	1	165	240	2200
Lightly damped	0.5	2	8	70
Missile autopilot	8	10	20	10

Table 7.1: Comparison of control signal magnitudes

compensator designed using a reduced sector bound may be successful. However, the application of full-order anti-windup, regardless of the size of the sector, served to increase the severity of saturation by an order of magnitude, meaning that only a small reduction in the size of the sector could be accommodated if the sector condition was required to hold during operation. This shows that the size of the control signal in linear and constrained simulations without anti-windup is not an indicator of the size of sector required for the stability guarantee to apply for normal operation. However, stability was observed in practice.

The missile autopilot system was a much more successful application of the reduced sector bound as a performance improvement was gained and the peak magnitude of the control signal was maintained close to that of the linear simulation. As a result, it is likely that local stability may be sufficient to ensure stability during normal operation. For compensator designs that did cause the sector bound to be violated during operation, limit cycle behaviour was observed but by suitable tuning, a performance benefit could be gained whilst still maintaining a relatively large sector. Of course, the peak magnitude of the control signals that may occur in a given system are also dependent upon the exogenous inputs applied to it and so some form of global stability guarantee may still be desired. However, this may be provided by a less conservative method such as the Popov Criterion.

The application of reduced sectors has only been shown for the application to full-order anti-windup but it is equally applicable to both the low-order and static compensation strategies. Further benefits of reduced sectors alluded to in [84] are that they may enable solution to the anti-windup stabilisation and synthesis problem for non-Hurwitz plants although this area has not been explored here.

Another method available is to optimise local performance using a sector bound that matches the expected size of the control signals, but guarantee global stability by incorporating the full sector into the stability condition. This can be achieved by posing both problems as separate LMIs and then combining them together to be solved simultaneously. Forays in this direction revealed that when the global stability condition was present, the outcome of the optimisation varied very little from the standard case when performance and stability were considered globally. This reveals that focusing the performance optimisation on a smaller set of signals has little effect on achieved performance. Furthermore, it suggests that the size of the sector has a more significant impact on the existence of a stability guarantee for a given system, and it is generally this that restricts the attainable performance level for an anti-windup design.

7.3.4 Conclusions

The following statements are the conclusions reached from this research into the use of reduced sectors.

- In most cases, use of the reduced sector bound was observed to have a positive effect on performance in simulation, although its effect was generally quite minor. The improved full-order performance may still lag behind that of a well tuned low-order compensator.
- Stability of the anti-windup compensated system is only guaranteed while the sector condition holds and therefore for a practical application, an *a priori* check of global stability is desirable. However, due to the conservatism of the circle criterion, global stability is often afforded in practice and so a single check using a more complex but less conservative method may be all that is required.
- The \mathcal{L}_2 performance level returned by compensator synthesis using a reduced sector reduces sharply with the size of the sector and loses its relation to the \mathcal{H}^∞ norm of the connected linear system. Therefore, its value is only of use for comparing compensator designs that employ the same sector condition.
- Whether a given sector bound is sufficient to envelope the given nonlinearity is dependent upon a number of characteristics of the system. These include the dynamics of the nominal linear system, the type of controller, the design of the anti-windup compensator, the saturation limits and also the exogeneous inputs applied.
- Any performance gained by the use of a reduced sector in synthesis appears to result from the associated relaxation of the stability condition rather than focusing the optimisation toward performance for conditions in which the control signal is bounded in magnitude by u_{max} .

7.4 Overall Conclusions

This chapter has considered two ways in which the performance of anti-windup compensators could, in principle, be improved. In both cases, modifications to the existing full-order anti-windup synthesis algorithm of [88] have been proposed, one involving frequency weighting the \mathcal{L}_2 performance objective, the other by considering reduced sector bounds in order to obtain improved local performance.

Both methods have showed some merit in the simulation examples considered, although the performance benefit seen has been modest. The frequency weighted approach seems able to bestow some minor performance improvements when compared to a baseline full-order anti-windup design and may best be used to fine-tune designs, once a standard full order compensator has been computed. The reduced-sector approach demonstrated inconsistent results in simulation, although in some examples perhaps showed the best potential of the two approaches.

Combined with the success of similar approaches, for instance in [77, 6, 102], the reduced sector technique

seems able to deliver real performance benefits, although these may be better demonstrated on more realistic, practical systems.

All in all, the theoretical results here should be considered as useful forays into the difficult territory of designing optimal, technically rigorous compensators and they lay the foundations for further research of this type.

Chapter 8

Conclusion

The main contributions of the thesis and some areas for future research can be summarised as:-

- The main contribution of this thesis is the application of modern optimal anti-windup compensation to a high performance Permanent Magnet Synchronous Motor (PMSM) speed control system that is employed within an Electrically Powered Hydraulic Steering system. Anti-windup is applied to compensate for violation of a constraint on the norm of a current demand vector signal within the control system. A novel approach allows this complex multivariable constraint to be considered as a time-varying SISO constraint for the purposes of compensator synthesis, allowing the application of modern optimal anti-windup approaches.
- Following a review of anti-windup compensation strategies, a number of classical and modern approaches were applied to simulation models of the PMSM speed control system. A selection of these designs were tested extensively on a complex multi-rate discrete time nonlinear model of the EPHS system and the most appealing classical and modern designs were subjected to practical testing on an EPHS test rig. For practical testing the compensator designs were incorporated into a prototype software build and run on a production ECU using fixed point arithmetic. The optimal low order anti-windup design yielded arguably the best time-domain performance whilst also providing rigorous stability guarantees for the nonlinear system. In spite of variation in the system dynamics with motor speed, theoretical analysis showed that a single linear compensator design was able to provide such stability guarantees for in excess of 95% of the operating envelope. Furthermore, the \mathcal{L}_2 performance bound was shown to be virtually invariant with the operating condition of the system. In addition to providing new results for PMSM speed control, the success and robustness of the low-order dynamic design for this application gives greater confidence for the application of this approach to other engineering systems. Furthermore, the application results presented here are believed to be the first application of modern anti-windup compensation techniques to a real PMSM system.

One opportunity for future work in the context of PMSM control is the application of nonlinear anti-windup methods to the current limitation problem. Techniques based on Nonlinear Dynamic Inversion

such as presented in [57] promise effectively to cancel the nonlinearity in the plant. This should reduce variation of the nonlinear anti-windup closed loop dynamics with the operating condition and as a result, may allow stability to be guaranteed over the whole operating envelope. In addition, performance improvements may also be achieved.

A similar but less computationally intensive option would be to design a gain scheduled linear anti-windup compensator. Both the static and low-order designs could be adapted quite easily into gain-scheduled designs. This may allow the stability to be guaranteed across the whole operating envelope without a significant sacrifice of performance and with little additional complexity for implementation.

- It was identified that the problem of limiting current with the EPHS system case study is more naturally cast as an override problem as the currents are states of the plant. A control structure was proposed applying a similar simplification of the vector norm saturation problem as in the anti-windup application work. This structure allowed the application of optimal override compensation approaches to the EPHS system model and static and dynamic designs were produced. The dynamic override design gave encouraging results, showing more consistent limitation of the current than afforded by the anti-windup approach in simulation testing. Again, it is believed that this chapter reports the first application of modern, theoretically rigorous override techniques, to PMSM systems.

Although override control is considered to be an appealing conceptual approach, some minor problems remain that warrant further research. The first is that transient violation of the limit can still occur and although this is more short-lived than in the anti-windup approach, it may exceed the limit by a greater extent. A key issue with this difficulty is that in the override approach, limit violation is necessary to trigger the activation of the compensator. Thus, confining attention to linear override controllers (in which the reaction will be linearly proportional to the constraint violation) may be too conservative. Some recent work on nonlinear override controllers may help to address this [36, 95], although, of course, such controllers have an associated complexity increase. Another alternative is that if the conditions that give rise to such overshoots can be reliably identified, an ad-hoc modification could be developed that temporarily reduces the saturation threshold in this condition. Such an approach may be difficult to tune but the stability guarantee afforded by the Circle Criterion does accommodate such time-varying limits and so the stability properties of the system should not be affected. A second problem is that, depending upon how the voltage and current limits are set, saturation of the voltage limit may prevent violation of the current limit and allow the outer loop to wind up. Further study between the interplay of the constraints is required to improve understanding of this phenomenon and develop a rigorous solution.

Another opportunity for future work is to remove the anti-windup compensator in the inner loop and tackle this problem via override compensation. This could be achieved by passing the d and q axis voltages through the plant as dummy outputs. This approach is appealing as it would allow the design of a single override compensator to compensate for both current and voltage saturation.

- A theoretical contribution toward improving anti-windup synthesis by the application of reduced sector

descriptions of the isolated nonlinearity was given. This technique was applied to full-order anti-windup compensation and shown to provide performance improvements during Mode III when saturation has ceased and the plant output converges toward that of the nominal linear system. Although the improvements seen were modest, the increased flexibility in the synthesis routine that results may provide more significant performance benefits in other applications.

Further work in this area could explore applications to both low-order and static compensation. For these more simple compensator designs, the extra flexibility allowed in tuning should be more transparent, allowing the characteristics of systems for which reduced sectors provide significant performance improvement to be identified.

- A second theoretical contribution considered the idea that the \mathcal{L}_2 gain objective can be too general a measure of anti-windup performance to ensure that a design with optimal \mathcal{L}_2 gain provides what the designer would consider to be best performance. Specifically, the concept of weighting performance objectives over frequency was borrowed from linear robust control and applied to the \mathcal{L}_2 design problem for full order anti-windup compensation. The technique revealed that by giving certain frequency ranges greater significance in the performance optimisation, characteristics of the resulting design could be influenced that allowed a design to be fine-tuned for improved performance. Again, modest performance improvements were observed in simulation but it is thought that certain, more complex and more realistic systems might demonstrate the improvement more dramatically.

Other thoughts for future work on performance improvement for full-order compensators are as follows. For the low-order dynamic anti-windup compensator, one of the main benefits appears to be that the designer can choose the dynamics and therefore have a significant effect on the outcome of the synthesis. It is believed that this flexibility allows the true performance potential to be exploited, which may not correspond to the lowest possible \mathcal{L}_2 gain. Adopting a similar philosophy in the design of full-order compensators, improved performance may be afforded by allowing the designer to impose certain constraints on the dynamics during synthesis. One possible method may be to restrict the poles of the disturbance filter to certain regions. Similar work for linear systems and imposing quadratic stability can be found in [10].

- A significant area of future work is the topic of robustness. The designs presented in this thesis are robust in the sense that a single linear design has been shown to be robust to the variation in dynamics present in the nonlinear system and also to the differences between the model and a real system. However, the topic of robustness in terms of frequency domain uncertainty models as found in \mathcal{H}^∞ control has not been addressed. The lack of rigorous robustness analysis is typical of the majority of work in the field of anti-windup due to the difficulties in applying such approaches to nonlinear systems. However, some preliminary work can be found in [88] and [94]. Robustness to structured uncertainty could also be investigated.

Appendix A

PMSM Control System

A.1 Linearisation of the $d - q$ Axis PMSM Model

By making the substitutions

$$x = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} i_d \\ i_q \\ \omega_m \end{bmatrix} \quad u = \begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix} = \begin{bmatrix} V_d \\ V_q \\ T_l \end{bmatrix}$$

the nonlinear $d - q$ axis state equations of (3.10) can be expressed as

$$\frac{d}{dt} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} f_1(x, u) \\ f_2(x, u) \\ f_3(x, u) \end{bmatrix} = \begin{bmatrix} 1/L_s (u_1 - R_s x_1 + PL_s x_3 x_2) \\ 1/L_s \left(u_2 - R_s x_2 - PL_s x_3 x_1 - \frac{ke}{\sqrt{2}} x_3 \right) \\ 1/L_s \left(\frac{\sqrt{3}}{2} ke x_2 - u_3 - B x_3 \right) \end{bmatrix}$$

Linearising about an equilibrium point (x_0, u_0) gives

$$\dot{\tilde{x}} = A\tilde{x} + B\tilde{u}$$

where

$$A = \begin{bmatrix} \partial f_1/\partial x_1 & \partial f_1/\partial x_2 & \partial f_1/\partial x_3 \\ \partial f_2/\partial x_1 & \partial f_2/\partial x_2 & \partial f_2/\partial x_3 \\ \partial f_3/\partial x_1 & \partial f_3/\partial x_2 & \partial f_3/\partial x_3 \end{bmatrix} \quad B = \begin{bmatrix} \partial f_1/\partial u_1 & \partial f_1/\partial u_2 & \partial f_1/\partial u_3 \\ \partial f_2/\partial u_1 & \partial f_2/\partial u_2 & \partial f_2/\partial u_3 \\ \partial f_3/\partial u_1 & \partial f_3/\partial u_2 & \partial f_3/\partial u_3 \end{bmatrix}$$

Evaluating these derivatives gives

$$A = \begin{bmatrix} -R_s/L_s & Px_3(0) & Px_2(0) \\ -Px_3(0) & -R_s/L_s & -Px_1(0) - ke/L_s\sqrt{3} \\ 0 & ke\sqrt{3}/2J & -B/J \end{bmatrix} \quad B = \begin{bmatrix} 1/L_s & 0 & 0 \\ 0 & 1/L_s & 0 \\ 0 & 0 & -1/J \end{bmatrix}$$

where $x_i(0)$ represent the equilibrium value of state x_i . The output equation is given by

$$\begin{bmatrix} \omega_m \\ i_d \\ i_q \end{bmatrix} = \begin{bmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix}$$

and if required, the electromagnetic torque can be determined by

$$T_e = \begin{bmatrix} 0 & 0 & \frac{K_e\sqrt{3}}{2} \end{bmatrix} \begin{bmatrix} x1 \\ x2 \\ x3 \end{bmatrix}$$

A.2 Model Parameters

A full set of normalised parameters used for modelling and simulation of the PMSM speed control system is included in the following tables.

Model Parameter	Normalised Value
R_s	0.2541
L_s	0.1277
B	0.0323
J	1.4837
K_e	1
P	5
τ_{crt}	$200e^{-6}$ units
τ_{spd}	$1e^{-3}$ units

Table A.1: Model parameters

Controller parameter	Normalised value
$k_{p,crt}$	2
$k_{i,crt}$	2
$k_{p,spd}$	9.87
$k_{i,spd}$	413
τ_{crt}	$200e^{-6}$
τ_{spd}	$1e_{-3}$

Table A.2: Normalised controller parameters used for design and simulation

Parameter	Normalised value
Drive stage voltage limit	1
Supply Voltage	2
PWM modulation depth	1

Table A.3: Additional parameters relating to the Simulink model

Normalised speed	Phase advance angle [Rad]	Normalised speed	Phase advance angle [Rad]
0	0	2.05794015	1.1244
0.5879829	0	2.11673844	1.14
0.64678119	0	2.17553673	1.1548
0.70557948	0	2.23433502	1.1689
0.76437777	0.0154	2.29313331	1.1823
0.82317606	0.1827	2.3519316	1.1951
0.88197435	0.3115	2.41072989	1.2074
0.94077264	0.4155	2.46952818	1.2191
0.99957093	0.5022	2.52832647	1.2303
1.05836922	0.5761	2.58712476	1.2411
1.11716751	0.64	2.64592305	1.2515
1.1759658	0.6961	2.70472134	1.2614
1.23476409	0.7459	2.76351963	1.2711
1.29356238	0.7904	2.82231792	1.2804
1.35236067	0.8305	2.88111621	1.2894
1.41115896	0.8669	2.9399145	1.2981
1.46995725	0.9001	2.99871279	1.3065
1.52875554	0.9306	3.05751108	1.3147
1.58755383	0.9587	3.11630937	1.3226
1.64635212	0.9846	3.17510766	1.3303
1.70515041	1.0088	3.23390595	1.3379
1.7639487	1.0312	3.29270424	1.3452
1.82274699	1.0522	3.35150253	1.3524
1.88154528	1.0719	3.41030082	1.3593
1.94034357	1.0905	3.46909911	1.3662
1.99914186	1.1079	3.5278974	1.3729

Table A.4: Phase advance map

Appendix B

Anti-Windup Compensation

B.1 State Space Matrices for Minimal Realisation of Static and Low Order Anti-Windup Compensators

Minimal state space realisations for the static and low order anti-windup compensators presented in 4.4 are parameterised by the following matrices where the nominal controller and plant are given by (B.1) and (B.2) and the matrices Δ and $\tilde{\Delta}$ are defined as shown in (B.3) and (B.4).

$$K(s) \sim \begin{cases} \dot{x}_c = A_c x_c + B_c y + B_{cr} r \\ u = C_c x_c + D_c y + D_{cr} r \end{cases} \quad (\text{B.1})$$

$$G(s) \sim \begin{cases} \dot{x}_p = A_p x_p + B_{pu} u_m + B_{pd} d \\ y = C_p x_p + D_p u_m + D_{pd} d \end{cases} \quad (\text{B.2})$$

$$\Delta = (I - D_p D_c)^{-1} \quad (\text{B.3})$$

$$\tilde{\Delta} = (I - D_c D_p)^{-1} \quad (\text{B.4})$$

B.2 Static Compensator Matrices

$$\bar{A} = \begin{bmatrix} A_p + B_{pu} \tilde{\Delta} D_c C_p & B_{pu} \tilde{\Delta} C_c \\ B_c \Delta C_p & A_c + B_c \Delta D_p C_c \end{bmatrix}, \quad B_0 = \begin{bmatrix} B_{pu} \tilde{\Delta} \\ B_c \Delta D_p \end{bmatrix}$$

$$\bar{B} = \begin{bmatrix} B_{pu} \tilde{\Delta} & -B_{pu} \tilde{\Delta} D_c \\ B_c \Delta D_p & -B_c \Delta \end{bmatrix}, \quad \bar{C}_1 = \begin{bmatrix} \tilde{\Delta} D_c C_p & \tilde{\Delta} C_c \end{bmatrix}$$

$$D_{01} = \tilde{\Delta} D_c D_p, \quad \bar{D}_1 = \begin{bmatrix} \tilde{\Delta} & -\tilde{\Delta} D_c \end{bmatrix}$$

$$\bar{C}_2 = \begin{bmatrix} \Delta C_p & \Delta D_p C_c \end{bmatrix}, \quad D_{02} = \Delta D_p, \quad \bar{D}_2 = \begin{bmatrix} \Delta D_p & -\Delta D_p D_c \end{bmatrix}$$

B.3 Low-Order Compensator Matrices

$$\bar{A} = \begin{bmatrix} A_p + B_{pu}\tilde{\Delta}D_cC_p & B_{pu}\tilde{\Delta}C_c & B_{pu}\tilde{\Delta}C_1 & -B_{pu}\tilde{\Delta}D_cC_2 \\ B_c\Delta C_p & A_c + B_c\Delta D_pC_c & B_c\Delta D_pC_1 & -B_c\Delta C_2 \\ 0 & 0 & A_1 & 0 \\ 0 & 0 & 0 & A_2 \end{bmatrix}$$

$$B_0 = \begin{bmatrix} B_{pu}\tilde{\Delta} \\ B_c\Delta D_p \\ 0 \\ 0 \end{bmatrix}, \quad \bar{B} = \begin{bmatrix} B_{pu}\tilde{\Delta}D_1 & -B_{pu}\tilde{\Delta}D_cD_2 \\ B_c\Delta D_pD_1 & -B_c\Delta D_2 \\ B_1 & 0 \\ 0 & B_2 \end{bmatrix}$$

$$\bar{C}_1 = \begin{bmatrix} \tilde{\Delta}D_cC_p & \tilde{\Delta}C_c & \tilde{\Delta}C_1 & -\tilde{\Delta}D_cC_2 \end{bmatrix}, \quad \bar{C}_2 = \begin{bmatrix} \Delta C_p & \Delta D_pC_c & \Delta D_pC_1 & -\Delta D_pD_cC_2 \end{bmatrix}$$

$$D_{01} = \tilde{\Delta}D_cD_p, \quad \bar{D}_1 = \begin{bmatrix} \tilde{\Delta}D_1 & -\tilde{\Delta}D_cD_2 \end{bmatrix}$$

$$D_{02} = \Delta D_p, \quad \bar{D}_2 = \begin{bmatrix} \Delta D_pD_1 & -\Delta D_pD_cD_2 \end{bmatrix}$$

where the filter dynamics chosen by the designer are given by

$$F_1(s) \sim \left[\begin{array}{c|c} A_1 & B_1 \\ \hline C_1 & D_1 \end{array} \right], \quad F_2(s) \sim \left[\begin{array}{c|c} A_2 & B_2 \\ \hline C_2 & D_2 \end{array} \right]$$

B.4 State-Space Matrices for BCAT Stability Check LMI

$$\Delta = \begin{bmatrix} I - D_pD_c & 0 \\ -D_pD_c & I \end{bmatrix}^{-1}, \quad \tilde{\Delta} = (I - D_cD_p)^{-1}$$

$$\bar{A} = \begin{bmatrix} A_p + B_{pu}\tilde{\Delta}[D_c \ 0][C'_p \ C'_p]' & B_{pu}\tilde{\Delta}C_c \\ [0 \ B_c]\Delta[C'_p \ C'_p]' & A_c + [0 \ B_c]\Delta[D'_p \ D'_p]'C_c \end{bmatrix}$$

$$B_0 = \begin{bmatrix} B_{pu}\tilde{\Delta} \\ [0 \ B_c]\Delta[D'_p \ D'_p]' \end{bmatrix}, \quad \bar{B} = \begin{bmatrix} B_{pu}\tilde{\Delta} & -B_{pu}\tilde{\Delta}[D_c \ 0] \\ [0 \ B_c]\Delta[D'_p \ D'_p]' & -[0 \ B_c]\Delta \end{bmatrix}$$

$$\bar{C}_1 = \begin{bmatrix} \tilde{\Delta}[D_c \ 0][C'_p \ C'_p]' & \tilde{\Delta}C_c \end{bmatrix}, \quad D_{01} = \tilde{\Delta}[D_c \ 0][D'_p \ D'_p]', \quad \bar{D}_1 = \begin{bmatrix} \tilde{\Delta} & -\tilde{\Delta}[D_c \ 0] \end{bmatrix}$$

$$\bar{C}_2 = \begin{bmatrix} \Delta[C'_p \ C'_p]' & \Delta[D'_p \ D'_p]'C_c \end{bmatrix}, \quad D_{02} = \Delta[D'_p \ D'_p]', \quad \bar{D}_2 = \begin{bmatrix} \Delta[D'_p \ D'_p]' & -\Delta[D'_p \ D'_p]'[D_c \ 0] \end{bmatrix}$$

Appendix C

Override Compensation

C.1 State-Space Matrices of $G_{cl}(s)$ and $\tilde{G}_{cl}(s)$ for Static and Dynamic Override Compensation

$$\begin{aligned}
 C_y &= \begin{bmatrix} \Delta D_p C_c & \Delta C_p \end{bmatrix} & D_{y0} &= \begin{bmatrix} \Delta D_p D_{cr} & \Delta D_{pd} \end{bmatrix} & \bar{D}_y &= \begin{bmatrix} 0 & \Delta D_p \end{bmatrix} \\
 C &= \begin{bmatrix} D_{pc} C_c + D_{pc} D_c \Delta D_p C_c & C_{pc} + D_{pc} D_c \Delta C_p \end{bmatrix} \\
 D_0 &= \begin{bmatrix} D_{pc} D_{cr} + D_{pc} D_c \Delta D_p C_c & D_{pdc} + D_{pc} D_c \Delta D_{pd} \end{bmatrix} \\
 \bar{D} &= \begin{bmatrix} 0 & D_{pc} + D_{pc} D_c \Delta D_p \end{bmatrix} & A &= \begin{bmatrix} A_c + B_c \Delta D_p C_c & B_c \Delta C_p \\ B_p C_c + B_p D_c \Delta D_p C_c & A_p + B_p D_c \Delta C_p \end{bmatrix} \\
 B_0 &= \begin{bmatrix} B_c r + B_c \Delta D_p D_{cr} & B_c \Delta D_{pd} \\ B_p D_{cr} + B_p D_c \Delta D_p D_{cr} & B_{pd} + B_p D_c \Delta D_{pd} \end{bmatrix} & \bar{B} &= \begin{bmatrix} I & B_c \Delta D_p \\ 0 & B_p + B_p D_c \Delta D_p \end{bmatrix}
 \end{aligned}$$

where

$$\begin{aligned}
 \Delta &= (I - D_p D_c)^{-1} & \tilde{A} &= \begin{bmatrix} A_f & 0 \\ \bar{B} C_f & A \end{bmatrix} & \tilde{B}_0 &= \begin{bmatrix} 0 \\ B_0 \end{bmatrix} \\
 \tilde{B}_1 &= \begin{bmatrix} B_f \\ \bar{B} D_f \end{bmatrix} & \tilde{C}_1 &= \begin{bmatrix} \bar{D} C_f & C \end{bmatrix} & \tilde{D}_{01} &= D_0 \\
 \tilde{D}_1 &= \begin{bmatrix} \bar{D} & D_f \end{bmatrix} & \tilde{C}_2 &= \begin{bmatrix} C_f & 0 \end{bmatrix} & \tilde{D}_2 &= D_f
 \end{aligned}$$

C.2 Linearisation of the Phase Advance Function

The nonlinear closed loop system formed by the interconnection of the speed loop PI regulator (K), phase advance function (ψ) and linearised model of the current control system (G) is shown in Figure C.1. We desire to linearise the phase advance function and absorb this into the current control system model. To achieve this, the nonlinear modulus functions must be omitted. Ignoring the modulus function applied to the motor speed measurement has no influence on the result as for application in EPHS, the motor only operates in one direction

i.e. $\omega_m \geq 0$. On the other hand, the omission of the modulus function on the q -axis current demand means that for operation with $i_q < 0$, the d -axis current demand will be positive in the linear model as opposed to zero in the nonlinear system. Thus, the linearisation will only be accurate when $i_q \geq 0$. This should not be problematic as for the operating conditions in which the model is used for compensator design, this will always be the case.




Figure C.1: Nonlinear closed loop system with phase advance

From the block diagram we have

$$i_{d,dmd} = f(i_{q,dmd}, \omega_m) = -|i_{q,dmd}| \times \tan(\psi(\omega_m))$$

and we desire to arrive at a linear relationship

$$i_{d,dmd} = [K i_q \quad K \omega_m] [i_{q,dmd} \quad \omega_m]'$$

for which

$$K i_q = \left. \frac{\partial f(\cdot)}{\partial i_{q,dmd}} \right|_{i_{q,dmd}(0), \omega_m(0)} \quad (C.1)$$

$$K \omega_m = \left. \frac{\partial f(\cdot)}{\partial \omega_m} \right|_{i_{q,dmd}(0), \omega_m(0)}. \quad (C.2)$$

The solution to (C.1) is given by

$$K i_q = \left. \frac{\partial f(\cdot)}{\partial i_{q,dmd}} \right|_{i_{q,dmd}(0), \omega_m(0)} = -\tan(\psi(\omega_m)) \Big|_{\omega_m(0)}. \quad (C.3)$$

The solution to (C.2) is given by differentiating by parts as follows. Defining $\alpha = \psi(\omega_m)$, we have that

$$\begin{aligned} \frac{\partial f(\cdot)}{\partial \omega_m} &= -i_{q,dmd} \frac{\partial \tan(\alpha)}{\partial \alpha} \cdot \frac{\partial \alpha}{\partial \omega_m} \\ &= -i_{q,dmd} \sec^2(\alpha) \frac{\partial \psi(\omega_m)}{\partial \omega_m}. \end{aligned}$$

Thus, by evaluating at $(\omega_m(0), i_{q,dmd}(0))$ and substituting in for $\alpha = \psi(\omega_m)$ we arrive at

$$K \omega_m = \left. \frac{\partial f(\cdot)}{\partial \omega_m(0)} \right|_{\omega_m(0), i_{q,dmd}(0)} = -i_{q,dmd}(0) \sec^2(\psi(\omega_m(0))) \frac{\partial \psi(\omega_m(0))}{\partial \omega_m(0)} \quad (C.4)$$

C.3 Linear Plant Model Matrices for PMSM Override Case Study

The linear plant model is parameterised by the following matrices.

$$\begin{aligned}
 A_p &= \begin{bmatrix} -2.844 & 1.263 & -0.009809 & -0.158 & 0.05071 \\ 8.07 & -136.9 & 118.2 & 261.2 & -33.44 \\ 3.53 & -145 & -117 & -35.66 & 175.9 \\ 7.564 & -278.5 & -487 & -1312 & -329.8 \\ -3.429 & 116.1 & 213.1 & 1255 & -1362 \end{bmatrix} \\
 B_{pd} &= \begin{bmatrix} 1.543 \\ -2.204 \\ -0.9575 \\ -2.052 \\ 0.9302 \end{bmatrix} & B_p &= \begin{bmatrix} -0.001138 \\ -0.008051 \\ 0.004471 \\ 0.03089 \\ -0.03901 \end{bmatrix} \\
 C_p &= \begin{bmatrix} -1.542 & 0.3585 & 0.03594 & 0.0286 & -0.002264 \end{bmatrix} \\
 C_{pc} &= \begin{bmatrix} -0.0008787 & -0.3851 & -0.5068 & 1.028 & 0.8355 \\ 0.05575 & -2.14 & 0.8116 & 1.776 & -0.4106 \end{bmatrix} \\
 D_{pd} = D_p &= 0 & D_{pdc} = D_{pd} &= \begin{bmatrix} 0 & 0 \end{bmatrix}'
 \end{aligned}$$

References

- [1] K.J. Aström and B. Wittenmark, *Computer controlled systems, theory and design*, Prentice Hall, Englewood Cliffs, NJ, 2nd Ed., 1984.
- [2] K.J. Aström and L. Rundqwist, Integrator windup and how to avoid it, *Proceedings of the Automatic Control Conference*, Vol. 2:1693-1698, 1989.
- [3] K.J. Aström and T. Hägglund, *Advanced PID Control*, Instrumentation, Systems and Automation Society, 2006.
- [4] L. Bei, W. Fen and S. Kim, Linear parameter-varying-antiwindup compensation for enhanced flight control performance, *Journal of Guidance, Control and Dynamics*, Vol.28(3):494-505, 2005.
- [5] ROBERT BOSCH GmbH *CAN Specification 2.0*,
www.semiconductors.bosch.de/pdf/can2spec.pdf, Postfach 30 02 40, D-70442 Stuttgart, 1991.
- [6] O. Brieger, M. Kerr, D. Leissling, I. Postlethwaite, J. Sofrony and M.C. Turner, Anti-windup compensation of rate saturation in an experimental aircraft, *Proceedings of the American Control Conference*, pp. 924-929, 2007.
- [7] F. Bruzelius, *Linear Parameter-Varying Systems, an approach to gain scheduling*, Thesis for the degree of Doctor of Philosophy, Chalmers University of Technology, Sweden, 2004.
- [8] J. Chandler, *PMSM technology in high performance variable speed applications*,
http://www.servo-motors-controls.com/PMSM_Technology_In_High_Performance_Variables.htm,
presented at the SPS Electric Automation America Conference, Chicago, Illinois, 25-26 May, 2005.
- [9] C.Y. Chen and M.H. Perng, Optimal anti-windup control of saturating discrete-time MIMO systems, *International Journal of Control*, Vol. 67(6):933-959, 1997.
- [10] M. Chilali, P. Gahinet and P. Apkarian, Robust pole placement in LMI regions, *IEEE Transactions on Automatic Control*, Vol.44(12):2257-2270, 1999.
- [11] Copley Controls Inc. *What is 'Field Oriented Control' and what good is it?*,
www.copleycontrols.com/motion/downloads/pdf/Field-Oriented-Control.pdf

- [12] S. Crawshaw, Global and local analyses of coprime-factor based anti-windup for stable and unstable plants, *Proceedings of the European Control Conference*, 2003.
- [13] C. Edwards and I. Postlethwaite Anti-windup schemes with closed-loop stability considerations, *Proceedings of the European Control Conference*, 1997.
- [14] C. Edwards, and I. Postlethwaite, Anti-Windup and Bumpless Transfer Schemes, *Automatica*, Vol. 34(2):199-210, 1998.
- [15] H.A. Fertik and C.W Ross *Direct digital control algorithms with anti-windup feature*, Instrument Society of America, Preprint Number 10-1-1-CQS-67.
- [16] A.E. Fitzgerald, C. Kingsley Jr and A. Kusko, *Electric Machinery* McGraw-Hill, 1971.
- [17] F. Forni and S. Galeani, Model based, gain-scheduled anti-windup control for LPV systems, *46th IEEE Conference on Decision and Control*, pp. 1174-1179, 2007.
- [18] G. Franklin, J.D. Powell and A. Emami-Naeini, *Feedback Control of Dynamic Systems* Prentice Hall, 2005.
- [19] A.H. Glattfelder and W. Schaufelberger, *A Path from Antiwindup to Override Control*.
- [20] A.H. Glattfelder and W. Schaufelberger, *Control systems with input and output constraints* Springer, 2003.
- [21] R. Goebel, Stabilizing linear systems with saturation through optimal control, *IEEE Transactions on Automatic Control*, Vol. 50(5):650-655, May 2005.
- [22] G. Grimm, I. Postlethwaite, A.R. Teel, M.C. Turner and L. Zaccarian, Linear Matrix Inequalities for full and reduced order anti-windup synthesis, *Proceedings of the American Control Conference*, (2001).
- [23] G. Grimm, I. Postlethwaite, A.R. Teel, M.C. Turner and L. Zaccarian, Case Studies Using Linear Matrix Inequalities in Anti-Windup Synthesis for Stable Linear Systems with Input Saturation, *European Journal of Control*, Vol 9:459-469, 2003.
- [24] G. Grimm, et al. Antiwindup for stable linear systems with input saturation: an LMI-based synthesis, *IEEE Transactions on Automatic Control*, Vol. 48(9):1509-25, 2003.
- [25] G. Grimm, A. Teel and L. Zaccarian, Robust linear anti-windup synthesis for recovery of unconstrained performance, *International Journal of Robust and Nonlinear Control*, Vol. 14:1133-1168, 2004.
- [26] A. Hansson, P. Gruber and J. Tödli Fuzzy anti-reset windup for PID controllers, *IFAC 12th Triennial World Congress*, pp. 917-920, 1993.
- [27] R. Hanus A new technique for preventing control windup, *Journal A*, Vol. 21(1), 1980.
- [28] R. Hanus, M. Kinnaert and J.L. Henrotte Conditioning technique, a general anti-windup and bumpless transfer method, *Automatica*, Vol. 23(6):729-739, 1987.

- [29] R. Hanus, Anti-windup and bumpless transfer: a survey, *Proc. 12th IMACS World Congress on scientific computing, Paris*, pp. 59-65, 1988.
- [30] J.R. Hendershot Jr and T.J.E. Miller, *Design of brushless permanent-magnet motors*, Magna Physics Publications, 1994.
- [31] G. Herrmann, M.C. Turner and I. Postlethwaite, Some new results on anti-windup-conditioning using the Weston-Postlethwaite approach, *43rd Conference on Decision and Control*, 2004.
- [32] G. Herrmann, M.C. Turner, I. Postlethwaite and G. Guo, Practical implementation of a novel anti-windup scheme in a HDD-dual-stage servo-system, *IEEE/ASME Transactions on Mechatronics*, Vol. 9(3), 2004.
- [33] H. Herrmann, M.C. Turner and I. Postlethwaite, Performance oriented anti-windup for a class of neural network controlled systems, *Proceedings of the IEEE RAM/CIS Conference*, 2004.
- [34] G. Herrmann, M.C. Turner, P.P. Menon, D.G. Bates and I. Postlethwaite, Anti-Windup Synthesis for Nonlinear Dynamic Inversion Controllers, *Proceedings of the IFAC Symposium on Robust Control Design, Toulouse*, 2006.
- [35] G. Herrmann, M.C. Turner and I. Postlethwaite, Improvement of a novel dual-stage large-span track-seeking and track following method using anti-windup compensation, *Proceedings of the American Control Conference*, 2006.
- [36] G. Herrmann, M.C. Turner and I. Postlethwaite, A robust override scheme enforcing strict output constraints for a class of strictly proper systems, *Automatica*, Vol.44(3):753-760, 2008.
- [37] P. Hippe *Windup in control: Its effects and their prevention*, Springer, 2006.
- [38] T. Hu, A. Teel and L. Zaccarian, Non-quadratic Lyapunov functions for performance analysis of saturated systems, *44th IEEE Conference on Decision and Control*, 2005.
- [39] D. Ito, D.T. Ward and J. Valasek, Robust Dynamic Inversion Controller Design and Analysis for the X-38, AIAA-2001-4380, *AIAA Guidance, Navigation, and Control Conference, Montreal, Canada*, 6-9 August 2001.
- [40] N. Kapoor, A. Teel and P. Daoutidis, *An anti-windup design for linear systems with input saturation*
- [41] N. Kapoor, P. Daoutidis, An observer-based anti-windup scheme for non-linear systems with input constraints, *International Journal of Control*, Vol. 72(1):18-29, 1999.
- [42] M.V. Kothare, P.J. Campo, M. Morari and C.N. Nett, *A Unified Framework for the Study of Anti-Windup Designs*, CDS Technical Report No. CIT-CDS 93-011, 1993.
- [43] M.V. Kothare, P.J. Campo, M. Morari and C. N. Nett, A unified framework for the study of anti-windup designs, *Automatica*, Vol. 30(12):1869-1883, 1994.

- [44] N.E. Kahveci, P.A. Ioanou and M.D. Mirmirani, A robust adaptive control design for gliders subject to actuator saturation nonlinearities, *Proceedings of the American Control Conference, New York*, 2007.
- [45] H.K. Khalil, *Nonlinear Systems*, Third Edition, Prentice Hall, 2002.
- [46] M.L. Kerr, M.C. Turner and I. Postlethwaite, Practical Approaches to Low-Order Anti-Windup Compensator Design: A Flight Control Comparison, *Proceedings of the IFAC World Congress, Seoul, Korea*, 2008.
- [47] N.J. Krikelis, State feedback integral control with ‘intelligent’ integrators *International Journal of Control*, Vol. 32(3):465-473, 1980.
- [48] R. Krishnan, *Electric motor drives, modelling, analysis & control*, Prentice Hall, 2001.
- [49] W.H. Kwon and S. Han Seoul *Receding Horizon Control: Model Predictive Control for State Models*, Springer-Verlag London Ltd, 2005.
- [50] J.E. Larsson, A simple anti-windup method, *International Journal of Control*, Vol. 60(5):1025-1033, 1994.
- [51] Z. Lin and A. Saberi, Semi-global low-and-high gain design technique for linear systems with input saturation - stabilization and disturbance rejection, *International Journal of Robust and Nonlinear Control*, Vol.5(5):381-398, 1995.
- [52] Z. Lin and A. Saberi, Low-and-high gain design technique for linear systems subject to input saturation - a direct method, *International Journal of Robust and Nonlinear Control*, Vol.7(12):1071-1101. 1997.
- [53] A. Aceves Lopez and J. Aguilar Martin, Using multivariable nonlinear stability theory for override control systems, *Proceedings of the European Control Conference*, 1999.
- [54] A. Aceves Lopez and J. Aguilar Martin, *On the stability of override control systems*.
- [55] A. Marcos, M.C. Turner, D.G. Bates and I. Postlethwaite, Robustification of static and low order anti-windup designs, *5th IFAC Symposium on Robust Control Design*, Vol.5(1), 2006.
- [56] V.R. Marcopoli and S.M. Phillips, Analysis and synthesis tools for a class of actuator-limited multivariable control systems: a Linear Matrix Inequality approach, *International Journal of Robust and Nonlinear Control*, Vol. 6:1045-1063, 1996.
- [57] P.P. Menon, G. Herrmann, M.C. Turner, D.G. Bates and I. Postlethwaite, General Anti-windup synthesis for input constrained nonlinear systems controlled using nonlinear dynamic inversion, *Proceedings of the 45th IEEE Conference on Decision and Control*, San Diego, 2006.
- [58] P. March, M.C. Turner, G. Herrmann and I. Postlethwaite, Frequency weighted L2 gain optimisation for improved anti-windup performance, *Proceedings of the International Control Conference*, 2006.

- [59] P. March and M.C. Turner, Anti-windup compensator designs for permanent magnet synchronous motor speed regulation *International Electric Machines and Drives Conference*, 2007.
- [60] P. March and M.C. Turner, Anti-windup compensator designs for permanent magnet synchronous motor speed regulation *IEEE Transactions on Industrial Applications*, to appear, 2009.
- [61] P. March and M.C. Turner, Practical considerations for override compensator synthesis and implementation *International Federation of Automatic Control World Congress, Seoul, Korea*, 2008.
- [62] S. Miyamoto and G. Vinnicombe, Robust control of plants with saturation nonlinearity based on coprime factor representations, *Proceedings of the 35th Conference on Decision and Control*, 1996.
- [63] M.Morari and E.Zafiriou, *Robust Process Control*, Englewood Cliffs, NJ, Prentice Hall.
- [64] E.F. Mulder and M.V. Kothare, Synthesis of stabilizing anti-windup controllers using piecewise quadratic Lyapunov functions, *Proceedings of the American Control Conference*, 2000.
- [65] E.F. Mulder, M.V Kothare and M. Morari, Multivariable anti-windup controller synthesis using linear matrix inequalities, *Automatica*, Vol. 37: 1407-1416, 2001.
- [66] D.W. Novotny and T.A. Lipo, *Vector control and dynamics of AC drives*, Oxford Science Publications, 2000.
- [67] D.Y. Ohm, *Dynamic model of PM Synchronous Motors*,
http://www.drivetechinc.com/articles/IM97PM_Rev1forPDF.pdf, Drivetech Inc., Blacksburg, Virginia.
- [68] S.K. Panda, Performance comparison of feedback linearised controller with PI controller for PMSM speed regulation, *Proceedings of the 1996 International Conference on Power Electronics, Drives and Energy Systems for Industrial Growth*, Vol.1(1):353-9, 1996.
- [69] S.K. Panda, J.M.S. Lim, P.K. Dash and K.S. Lock Gain-scheduled PI Speed Controller for PMSM Drive, *23rd Conference on Industrial Electronics, Control and Instrumentation*, 1997, Vol. 2(2):925-30, 1997.
- [70] R.M. Phelan, *Automatic control systems*, Ithaca and London: Cornell University Press, 1977.
- [71] S.J. Qin and T.A. Badgwell, *An Overview of Industrial Model Predictive Control Technology*,
<http://www.che.utexas.edu/qin/cpcv/cpcv14.html>, 1996.
- [72] S. Rajaram, S.K. Panda and L.K. Sang, Performance comparison of feedback linearised controller with PI controller for PMSM speed regulation, *Proceedings of the International Conference on Power Electronics, Drives and Energy Systems for Industrial Growth*, Vol.1(1):353-359, 1996.
- [73] J. Reiner, G.J. Balas and W.L Garrard, Flight control design using robust dynamic inversion and time-scale separation, *Automatica*, Vol. 32(11): 1493-1504, 1996.

- [74] A.A. Rodriguez and J.R. Cloutier, Control of a bank-to-turn missile with saturating actuators, *Proceedings of the American Control Conference*, 1994.
- [75] J.A. Rossiter, *Model-Based Predictive Control: A Practical Approach*, CRC Press Control Series, 2003.
- [76] L. Rundqwist, Anti-reset windup for PID controllers *Preprints of the 11th IFAC World Congress*, Vol. 8:146-151, 1990.
- [77] J.Sofrony, M.C.Turner and I.Postlethwaite, Anti-windup synthesis for systems with rate-limits: A Riccati equation approach, *Proceedings of the SICE Annual Conference*, pp. 689-694, 2005.
- [78] J. Sofrony, M.C. Turner, I. Postlethwaite, O. Brieger and D. Leissling, Anti-windup synthesis for PIO avoidance in an experimental aircraft, *45th IEEE Conference on Decision and Control*, pp. 5412-5417, 2006 .
- [79] J. Sofrony, M.C. Turner and I. Postlethwaite, Anti-windup synthesis using riccati equations, *International Journal of Control*, Vol. 80(1):112-128, 2007.
- [80] A. Soltoggio, *A comparison of genetic programming and genetic algorithms in the design of a robust, saturated control system*, GECCO, LNCS 3103, pp. 174-185, 2004.
- [81] S. Skogestad and I. Postlethwaite, *Multivariable feedback control, analysis and design*, John Wiley and Sons, 1996.
- [82] S. Tarboriech and G. Garcia, *Control of uncertain systems with bounded inputs*, Lecture Notes in Control and Information Sciences, Springer-Verlag, 1997.
- [83] A.R. Teel, Global stabilization and restricted tracking for multiple integrators with bounded controls, *Systems and Control Letters*, Vol. 18(3): 165-171, 1992.
- [84] A.R. Teel and N. Kapoor, The L_2 anti-windup problem: Its definition and solution, *Proceedings of the European Control Conference*. (1997).
- [85] M.C. Turner and I. Postlethwaite, *Output violation compensation for systems with output constraints*, University of Leicester Technical Report 02-02, March 2002.
- [86] M.C. Turner and I. Postlethwaite, Output Violation Compensation for Systems With Output Constraints, *IEEE Transactions on Automatic Control*, Vol. 47(9):1540-1546, 2002.
- [87] M.C. Turner, G. Herrmann and I. Postlethwaite, Discrete-time anti-windup, Part 1: Stability and Performance, *Proceedings of the European Control Conference*, 2003.
- [88] M.C. Turner, G. Herrmann and I. Postlethwaite, Accounting for uncertainty in anti-windup synthesis, *Proceedings of the American Control Conference*, Vol. 6:5292-5297, (2004).

- [89] M.C. Turner, G. Herrmann and I. Postlethwaite, *An introduction to linear matrix inequalities in control*, University of Leicester Internal Report no. 02-04, July 16, 2004.
- [90] M.C. Turner and I. Postlethwaite, Multivariable override control for systems with output and state constraints, *International Journal of Robust and Nonlinear Control*, Vol. 14:1105-1131, 2004.
- [91] M.C. Turner and I. Postlethwaite, A new perspective on static and low order anti-windup synthesis, *International Journal of Control*, Vol 77(1):27-44, 2004.
- [92] M.C. Turner, I. Postlethwaite and G. Herrmann, Further results on full-order anti-windup synthesis: exploiting the stability multiplier, *Proceedings of the IFAC Nonlinear Control System Design Symposium (NOLCOS)*, 2004.
- [93] M.C. Turner and I. Postlethwaite, Local anti-windup compensation using output limiting, *Proceedings of the International Federation of Automatic Control World Congress*, 2002.
- [94] M.C. Turner, G. Herrmann and I. Postlethwaite, Incorporating robustness requirements into anti-windup design, *IEEE Transactions on Automatic Control*, Vol. 52(10):1842-1855, 2007.
- [95] M.C. Turner, G. Herrmann and I. Postlethwaite, Improving sector-based results for systems with dead-zone nonlinearities and constrained control applications, *Automatica*, Vol. 45:155-160, 2009.
- [96] M. Saeki and N. Wada, Synthesis of a static anti-windup compensator via linear matrix inequalities, *International Journal of Robust and Nonlinear Control*, Vol. 12:927-953, 2002.
- [97] O. Wallmark, *On control of Permanent-Magnet Synchronous Motors in Hybrid-Electric Vehicle applications*, Thesis for the degree of Licentiate of Engineering, Chalmers University of Technology, Sweden, 2004.
- [98] M.J. Walsh and M.J. Hayes, A Robust Throughput Rate Control Mechanism for an 802.15.4 Wireless Sensor Network - An Anti-Windup Approach, *Proceedings of the American Control Conference, New York*, 2007.
- [99] P.F. Weston and I. Postlethwaite, *Analysis and design of linear conditioning schemes for systems containing saturating actuators*, Technical report 97-12, Department of Engineering, Leicester University, UK. (1997).
- [100] P.F. Weston and I. Postlethwaite, Linear conditioning for systems containing saturating actuators, *Automatica*, Vol 36:1347-1354, 2000.
- [101] F. Wu and B. Lu, Anti-windup control design for exponentially unstable LTI systems with actuator saturation, *Systems and Control Letters* Vol. 52: 305-322, 2004.
- [102] L. Zaccarian and A. Teel, Nonlinear scheduled anti-windup design for linear systems, *IEEE Transactions on Automatic Control*, Vol. 49(11), 2004.

- [103] K. Zhou and J. Doyle, *Essentials of robust control*, Prentice Hall, 1998.
- [104] A. Zheng, M.V. Kothare and M. Morari, Anti-windup design for internal model control, *International Journal of Control*, Vol. 60(5):1015-1024, (1994).