An Application of Coprime Factor Based Anti-Windup and Bumpless Transfer Control to the Spark Ignition Engine Idle Speed Control Problem

Richard G. Ford and Keith Glover
(rgf21, kg@eng.cam.ac.uk)
University of Cambridge, Dept. of Engineering. England

Abstract

This paper presents an application of the coprime factor based anti-windup technique of Miyamato and Vinnicombe [MV96]. The application is an automotive engine idle speed control problem using a novel framework which inverts most of the system nonlinearities. A discussion of the design procedure is given, and extensions to the technique are developed to enable implementation. The final controller is implemented in the Matlab/dSPACE rapid prototyping suite and is used to control a 1.8l spark ignition engine. Results for both anti windup and bumpless transfer situations are presented and discussed.

1 Introduction

The idle speed control problem is a popular candidate for the application of modern control techniques with over 50 papers in the last 5 years alone. In its simplest form the problem is a disturbance rejection problem, it being desired to keep the engine speed constant whilst under the influence of various (known and unknown) load sources such as power steering and air-conditioning pumps, alternator load, and general cyclic variability of the combustion events. If these disturbances can be rejected effectively, then the driver will have a better impression of engine quality and, in addition, the engine idle speed can be lowered which will reduce fuel consumption. Recent research has applied a wide range of control techniques to solving the idle control problem. They range from linear control techniques based around an operating point (eg LQG [BGOS86], $\mathcal{H}_\infty$ [GB96]) to nonlinear techniques (eg sliding mode [BSDE99], feedback linearisation [KNVH94], fuzzy control [AD90], neural networks [PFD96], other [GF99], [NOM92]). The amount of torque produced in a combustion event is dependent on the amount of air in the cylinder (in this work we assume that the air/fuel ratio is always kept at its stoichiometric value), and the point in the cycle at which it is ignited. The Air Bypass Valve (ABV), controlled by the engine management system, controls the air flow into the engine under idle conditions. The spark timing

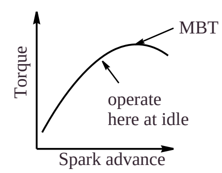

Figure 1: Combustion torque characteristics

controls the point of ignition with a spark timing which gives the Maximum (Brake) Torque (MBT) as indicated in figure 1. By retarding the spark away from MBT timing a region is obtained in which spark timing as an actuator of torque has authority in both directions. This actuator is fast, although still involving a single event delay between command and torque application but obviously has limited authority. The ABV is slow, involving valve and manifold air dynamics but has a lot of authority. Because of a) the nonlinear nature of the torque generation characteristics at a fixed operating point and b) the speed and severity of changes in characteristics due to operating point variations, the system as it stands is highly nonlinear. Traditionally the idle control problem has been thought of as a two input one output control problem, with the ABV and spark timing being used as actuators to keep the speed constant. In [For00], [FG00], [Ser00] the authors propose a new framework which both inverts the main nonlinearities involved and introduces a second output allowing the controller to vary the efficiency of the engine by operating closer to, or further from, MBT spark timing. The structure of the final system is illustrated in figure 2. In the new framework the outputs of the controller are combustion torque T_C and a manifold pressure related parameter $P_{M_{ss}}$. (This is the manifold pressure that would be reached if the manifold were allowed to reach equilibrium). The inputs to the controller are engine speed and the distance (in Nm) that it is operating from MBT, the *torque reserve* of the engine. The framework and controller evolve in (combustion) event based discrete time.

The closer to MBT timing that the engine is operated,

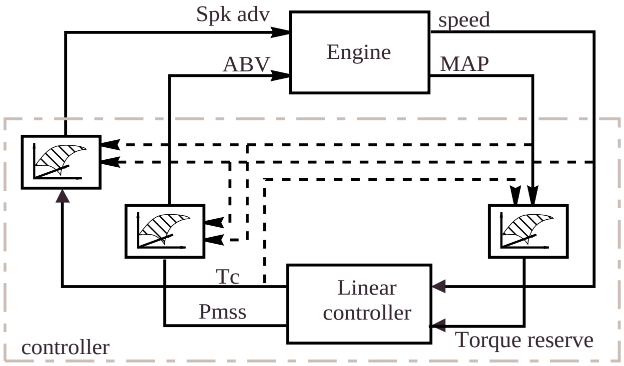

Figure 2: Linearising idle control framework

the more efficient the engine is however the less torque reserve it has when a load disturbance is met. From a control perspective, the less torque reserve the engine has, the more time the first input (T_C) spends in saturation. Because of this saturation a controller with anti-windup characteristics is necessary. As well as the anti-windup function of the controller, it is also necessary to have some kind of bumpless transfer mechanism. This serves two purposes: i) In the testing stage it is necessary to switch smoothly between the existing ECU based idle controller and the dSPACE based controller. ii) In production implementation there must be a reliable and smooth transition between non-idling and idling modes. To achieve this, again the controller states must be set to give sensible controller outputs before the idle controller is engaged.

2 Base Controller Design

The base controller is designed using the $\mathcal{H}_\infty$ optimal loopshaping design procedure of MacFarlane and Glover [MG90]. The designer shapes the plant using traditional loopshaping guidelines and then the optimisation creates a (stabilising) controller which is optimally robust to a broad class of uncertainties. Practical control designs using this synthesis technique are described in e.g. [Hyd95], [SPW$^+$98], [Bra96]. The loop transfer functions along with the singular values are shaped, to give the $T_C \rightarrow$ speed loop a slightly faster response than the $P_{M_{ss}} \rightarrow$ torque reserve loop.

3 Coprime Factor Based Anti Windup and Bumpless transfer (AWBT) Design

The unifying framework suggested in [MV96] is illustrated in figure 3. The original controller is split into coprime factors such that $C(s) = \tilde{V}_0^{-1}(s)\tilde{U}_0(s)$. Then, varying $Q(s)$ creates controller factorisations with different anti-windup characteristics. Note that if the system is not saturating then the original controller $C(s)$ is recovered. To ease some later manipulations, the $\tilde{V}_0$ and $\tilde{U}_0$ are chosen to satisfy the Bezout identity with the plant

$$\tilde{V}_0 M + \tilde{U}_0 N = I \tag{1}$$

where $P = NM^{-1}$ is a normalised coprime factorisation of P.

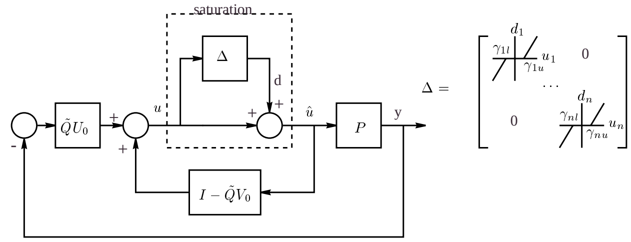

Figure 3: General layout of anti windup and bumpless transfer framework

To evaluate the performance of the AWBT system it is sensible to examine the effect of the disturbance on the input and output of the plant. Using (1)

$$\begin{aligned} T_{d\rightarrow\hat{u}} &= I + T_{d\rightarrow u} = MQ^{-1} && (2)\\ T_{d\rightarrow y} &= P(I + T_{d\rightarrow u}) = NQ^{-1} && (3)\end{aligned}$$

The Δ block has $\mathcal{L}_2$induced norm of 1 so a sufficient condition for stability, from the small gain theorem, is that $\|T_{d\rightarrow u}\|_\infty \equiv \left\|-I + MQ^{-1}\right\|_\infty < 1$. Using weights on the transfer functions to be shaped we obtain an $\mathcal{H}_\infty$ optimisation problem to find Q, it being that which achieves

$$\inf_{Q,Q^{-1}\in\mathcal{H}_\infty} \left\| \begin{bmatrix} 0 \\ 0 \\ I \end{bmatrix} - \begin{bmatrix} W_1 M \\ W_2 N \\ M \end{bmatrix} Q^{-1} \right\|_\infty \tag{4}$$

The optimisation is solved for Q^{-1} and then Q follows.

3.1 The algebraic loop issue

From figure 3 it will be seen that if the $Q\tilde{V}_0$ term has a through term then, when not in saturation, an algebraic loop will occur. In [MV96] it is pointed out that a strictly proper plant and W_1 will result in $Q(\infty)\tilde{V}_0(\infty) = I$. In the continuous time case this plant constraint is not severe since a sufficiently fast low-pass filter can be added. However for a discrete time implementation this is not the case. In any physically realisable system $C(\infty)P(\infty) = 0$, and from (1), $\tilde{V}_0(I + CP)M = I$. Therefore, for no algebraic loop we need $Q(\infty) = M(\infty)$. To guarantee this in the discrete time case we can instead solve the modified problem

$$\inf_{\hat{Q},\hat{Q}^{-1}\in\mathcal{H}_\infty} \left\| \begin{bmatrix} 0 \\ 0 \\ I \end{bmatrix} - \begin{bmatrix} W_1 M \\ W_2 N \\ M \end{bmatrix} (z^{-1}\hat{Q} + D_M)^{-1} \right\|_\infty \tag{5}$$

and then reform $Q = z^{-1}\hat{Q} + D_M$.

3.2 Designing the anti-windup controller

In normal operation the only saturating input is the first, T_C. The closer to MBT that the engine is operated, the more time it will spend in saturation. The performance of the AWBT system will to a large extent dictate how efficiently the engine can operate. Brief saturations are of no concern, since they will not cause windup. If the period of saturation increases, for example when a heavy load disturbance is encountered, it is important to increase the air flow into the engine so that the combustion torque ceiling is increased. The air into the engine is controlled by the second $P_{M_{ss}}$ input which itself is mainly affected by the second MBT-T_c, torque reserve output. Thus for improved recovery from saturation, it is important that the second synthetic output is driven by the *unsaturated* rather than saturated version of T_C. It should also be noted that if the anti-windup action is sufficiently fast that the desired T_C tracks accurately the saturated T_C, then recovery from saturation, effected by the $P_{M_{ss}}$ input, will be slower. Hence there is a trade-off to be had between rate of recovery from saturation and degree of anti-windup allowable.

To recover quickly from saturation, it would be desirable to include a cross-coupling term from d_1 to u_2 in the controller when in saturated mode, thus when T_C saturates the controller immediately boosts the $P_{M_{ss}}$ value. This would require increasing the size of the $T_{d_1 \to u_2}$ term. To make transfer functions *large* using the standard $\mathcal{H}_\infty$ optimisation is difficult and it may be that a more suitable framework is possible.

One of the main advantages of optimisation based control synthesis techniques, from a practical perspective, is that if the control problem fits the synthesis technique, then the engineer has simply to shape some cost function in a sensible and intuitive way. The optimisation then creates a 'sensible' controller for a control problem which may be impossible to solve by a traditional technique. However, if the cost function is not directly applicable to the control problem then optimisation based techniques may be harder to use since it may not be clear how the cost functions are to be shaped to create the desired controller behaviour. This was the case to some extent in the work presented here. Thus diagonal weights with varying first order characteristics were tried and the resulting AWBT controller tested. In the authors' experience, little benefit was obtained by using the W_1 weight and it was set to 0. To illustrate the effect of a simple variation, the system disturbance rejection characteristics were investigated with the AWBT functionality synthesised using a range of constant multipliers for a typical W_2 weight.

Figures 4 to 7 indicate the different transfer functions and characteristic responses to a 10Nm step load disturbance obtained in nonlinear simulation. The load is

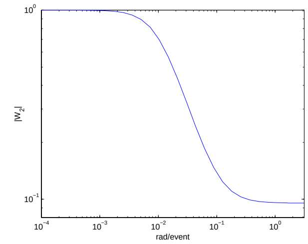

Figure 4: Typical W_2

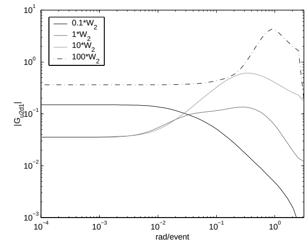

Figure 5: $|T_{d_1 \to u_2}|$ for various W_2

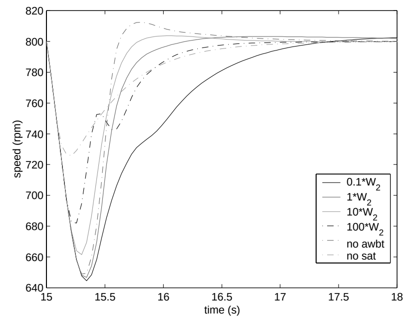

Figure 6: Speed response for various W_2 weights

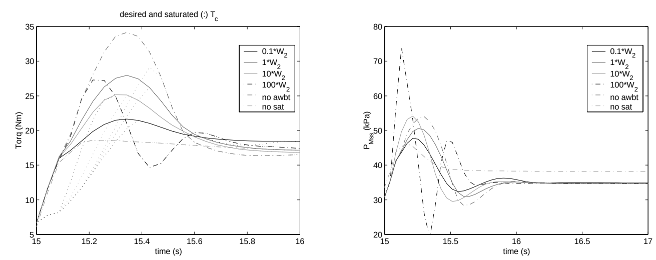

Figure 7: Controller output for various W_2 weights

applied at 15s. For comparison the behaviours of the system without anti windup control and when operated with enough torque reserve so that no saturation occurs are also shown. Some points to note are:

The winding up of T_C when no AWBT is used causes the second output to increase rapidly, increasing the manifold pressure and raising the limiting value of T_C. This in turn increases the rate at which the system crosses back through the reference speed.

It is evident that increasing the scaling of the W_2 weight from the minimum increases the speed of recovery. This is due to the cross-term $T_{d_1 \to u_2}$ which increases the $P_{M_{ss}}$ value as a result of the saturation. Beyond a certain level the increased interaction causes a lightly damped oscillation to appear.

3.3 Anti-Windup Performance in Hardware

To test the AWBT behaviour in the test cell, a heavy load application and removal is applied via the alternator in a square waveform and the average behaviour found. The load is near the upper end of those which the engine would experience. It is not surprising that the largest W_2 multiplier, 100, which seemed close to instability in noise free simulation could not control the system in practice and no results could be obtained when using it. As a reference, the behaviour without anti windup at both 2Nm and 6Nm torque reserve is also shown. For clarity, only the $0.1W_2$ and $10W_2$ tests are shown.

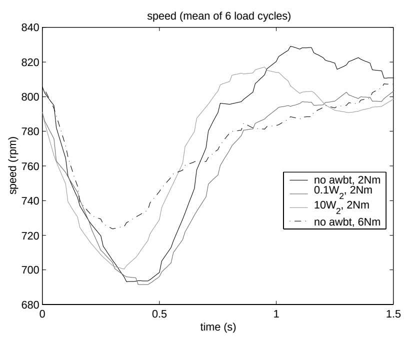

Figure 8: Speed response of system under load application and removal with anti-windup

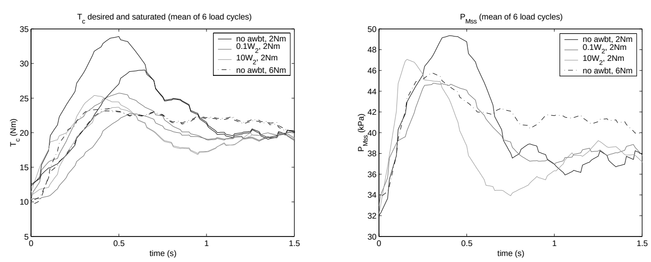

Figure 9: Controller outputs in hardware

The trends observed in the speed responses taken from the engine are similar to those observed in simulation, when the smallest value of the weight is used, the anti-windup performs a traditional role, keeping the controller output close to the actual input. This does cause the more sluggish recovery observed in the speed because the $P_{M_{ss}}$ variable is not boosted significantly (either by the cross-coupling term $T_{d_1 \to u_2}$ or by the larger negative torque reserve). The speed response with the higher gain on the weight exhibits some overshoot upon recovery. This can be seen also in the torque demand and $P_{M_{ss}}$ traces. In many control systems this would be a disadvantage however in the idle control system, speed errors above the setpoint are much less important than those below it.

4 Bumpless Transfer

Bumpless transfer functionality in a controller is used so that when the controller is offline the states can be conditioned, or controlled, in such a way that when the controller takes over the position of the active controller the transition at the input to the plant will be smooth. In the absence of a reference input (for both offline and online controllers), the anti-windup scheme will, without any modifications perform a bumpless transfer role. It will adjust the states of the controller to try to get the controller output to match the plant input. The rate at which the controller output converges to the true plant input depends on the Q term.

To aid in the understanding of the operation of the AWBT scheme in bumpless transfer mode, i.e. when another controller is actually controlling the plant, it has been observed ([Pax99]) that the framework with static Q is equivalent to a Kalman filter observer. It may be easier to understand the way the conditioning scheme operates if one assumes that the true online controller is similar to the offline controller, then the Kalman filter is just doing its usual task of estimating the states of the other system, given the inputs and outputs of that system.

For the most freedom in tracking response characteristics, the reference should be injected as proposed in [Vin93]. This corresponds to the layout of figure 10, where the reference filter ϕ is used to get steady state accuracy and shape the response. In the case when the online reference is different to the offline, although the states of $\tilde{U}_0$,$\tilde{V}_0$ will be compatible with the true plant inputs and outputs, in general the offline controller output will not be close to the true plant input. To ensure bumpless transfer we could feed the online reference into the offline controller, however in many cases the online reference is unknown or is incompatible with the offline controller references.

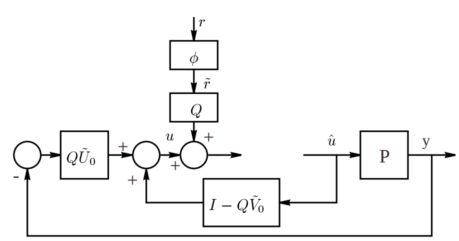

Figure 10: AWBT framework with reference

Fortunately, there can be a simple solution. Referring to figure 10, for bumpless transfer we require that at the point of switching, the states of $Q\tilde{U}_0$ and $I - Q\tilde{V}_0$ are compatible with the plant output and input respectively and that the states of the reference filter $Q\phi$ are such that $u_k = \hat{u}_k$. From [MV96] we have that $Q^{-1} \in \mathcal{H}_\infty$ so if we can achieve $\tilde{r} = Q^{-1}(\hat{u} - u)$ then the goal will be attained. In general, ϕ will not itself be invertible. Let ϕ be described in state space form by

$$\left[\frac{x_\phi(k+1)}{\tilde{r}(k)}\right] = \left[\begin{array}{c|c} A_\phi & B_\phi \\ \hline C_\phi & D_\phi \end{array}\right] \left[\frac{x_\phi(k)}{r(k)}\right] \qquad (6)$$

Then, in implementation, if C_ϕ is square (i.e the same number of states as outputs) and invertible the states of ϕ can be manipulated directly to achieve

$$x_\phi(k) = C_\phi^{-1}(Q^{-1}(\hat{u}(k) - u(k)) - D_\phi r(k)) \qquad (7)$$

which gives true bumpless transfer characteristics in an implementable form.

4.1 Bumpless transfer requirements in engines

The transition from drive to idle is difficult from a control perspective for a number of reasons:

1. The transition region during which the idle controller has to match with the online controller and plant outputs is an area of fast dynamic change, the engine typically decelerating quickly towards the idle regime.
2. Although a 'bumpless transfer' may be achieved, because the previous controller had different objectives (to decrease the speed of the engine) the controller state may be a long way from that needed to halt the deceleration at the reference speed.
3. The load on the engine will be unknown and therefore sensible steady state values for the idle controller cannot be estimated.

To achieve good transitioning, four engine operating modes are defined: *Transition in:* In this mode the idle controller is enabled but offline. *Idle:* In this mode the idle controller is enabled and on-line. *Transition out:* In this mode the idle controller is disabled and the ABV moved to a value which ensures a smooth pullaway. *Drive:* In this mode the idle controller is disabled.

4.2 Bumpless transfer tests in hardware

As a simple test of the transition characteristics of the controller the throttle was 'blipped' causing the engine speed to race then drop back to idle. This was performed once with no load and then again with a heavy (400W) load introduced just when the controller was capturing the speed. This is a worst case scenario.

Several remarks can be made with reference to the transitional behaviour of the system.

Most importantly, the behaviour of the engine speed upon transition into idle is excellent. In particular there is little of no undershoot which may lead to a stall situation, even though the speed reference drops significantly when the idle controller takes command of the inputs.

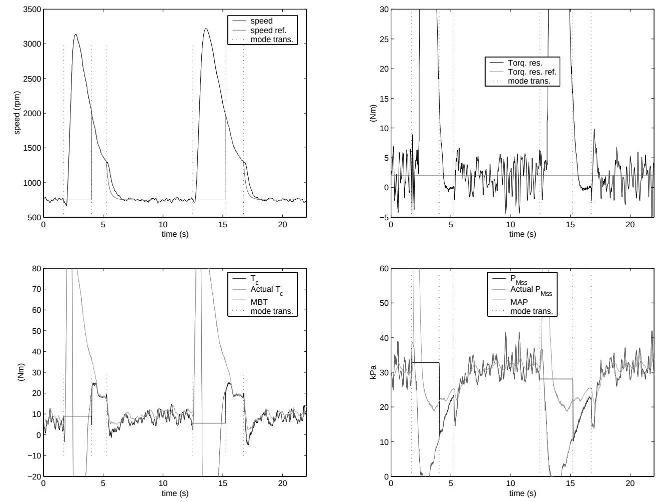

Figure 11: Transition characteristics with no extra electrical load

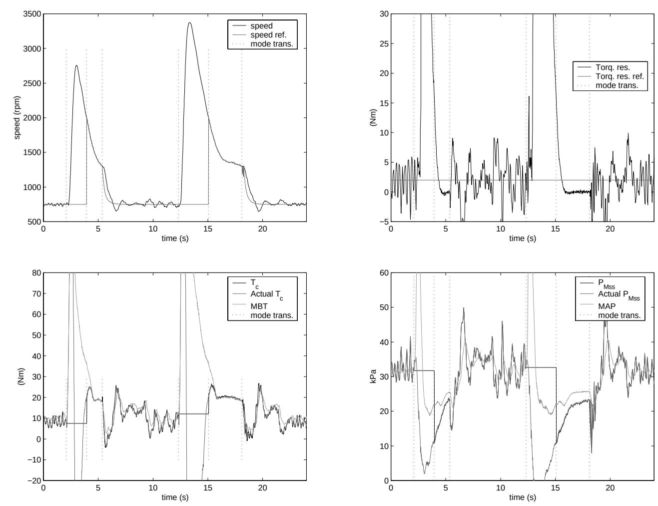

Figure 12: Transition characteristics with extra 400W electrical load switched in at critical time

The incompatibility of the engine model when outside the idle mode can be seen in the T_C and $P_{M_{ss}}$ traces when in drive mode. This does not however affect the transitional behaviour.

The addition of the extra load when the speed is settling to the reference speed has the expected effect of causing a slight dip in speed, however the controller rejects the disturbance well with minimal upset. By reading the T_C axis, it can be seen that the load application corresponds to approximately a 6-8Nm torque application.

5 Conclusions

In this paper, an application of the AWBT technique described in [MV96] to a well known automotive problem (albeit in a new framework) is given. A technique for guaranteeing the strictly proper nature of $T_{\hat{u}\to u}$ in the discrete time case has been introduced. The application has a number of interesting features not typical of anti-windup and bumpless transfer problems which have been discussed. The improvement in performance of the system when actuators saturate has been clearly demonstrated as have the bumpless transfer capabilities.

6 Acknowledgements

Richard Ford is funded by the Engineering and Physical Sciences Research Council. The engine test facilities at Cambridge are jointly funded by EPSRC and Ford Motor Company. To both the authors are grateful. Technical help of a theoretical nature from Glenn Vinnicombe and of a practical nature from Dave Gautrey is gratefully acknowledged.

References

[AD90] M. Abate and N. Dosio. Use of fuzzy logic for engine idle speed control. In *Proceedings of SAE Congress*, number 900594. Society of Automotive Engineers, 1990.

[BGOS86] C.E. Baumgartner, H.P. Geering, C.H. Onder, and E. Shafai. Robust multivariable idle speed control. In *Proceedings of Amercan Control Conference*, pages 258–265. American Control Council, 1986.

[Bra96] M. Brandstetter. *Robust Air-Fuel Ratio Control for Combustion Engines.* PhD thesis, University of Cambridge, 1996.

[BSDE99] A.I. Bhatti, S.K. Spurgeon, R. Dorey, and C. Edwards. Sliding mode configurations for automotive engine control. *International Journal of Adaptive Control and Signal Processing*, (13):49–69, 1999.

[FG00] R.G. Ford and K. Glover. Spark ignition engine idle speed control using a novel framework and facilitating control of the trade-off between fuel consumption and disturbance rejection capability. In *Proceedings of AVEC International Symposium on Advanced Vehicle Control*, 2000.

[For00] R. G. Ford. *Robust Automotive Idle Speed Control in a Novel Framework.* PhD thesis, University of Cambridge, 2000.

[GB96] K. Glover and M. Brandstetter. Idle speed control: Case studies with controllers designed via $\mathcal{H}_\infty$ loop shaping and μ-synthesis. Technical report, University of Cambridge, 1996. Internal Report for Ford Motor Company.

[GF99] J.W. Glass and M.A. Franchek. Narmax modelling and robust control of internal combustion engines. *International Journal of Robust Control*, 72(4):289–304, 1999.

[Hyd95] R. Hyde. *$\mathcal{H}_\infty$ Aerospace Control Design - A VSTOL Flight Application.* Advances in Industrial Control. Springer Verlag, 1995.

[KNVH94] L. Kjergaard, S. Nielsen, T. Vesterholm, and E. Hendricks. Advanced nonlinear engine idle speed control systems. In *Proceedings of SAE Congress*, number 940974. Society of Automotive Engineers, 1994.

[MG90] D.C. McFarlane and K. Glover. *Robust controller design using normalised coprime factor plant descriptions.* Lecture notes in Control and Information Sciences. Springer Verlag, 1990.

[MV96] S. Miyamoto and G. Vinnicombe. Robust control of plants with saturation nonlinearity based on coprime factor representations. In *Proceedings of IEEE Conference on Decision and Control.* IEEE, 1996.

[NOM92] M. Nasu, A. Ohata, and T. Meguro. Idle speed control by nonlinear feedback. *JSAE Review*, 13(2), 1992.

[Pax99] J. Paxman. Conditioning schemes for robust multivariable controllers. Mphil thesis, University of Cambridge, Dept. of Engineering, 1999.

[PFD96] G.V. Puskorius, L.A. Feldkamp, and L.I. Davis. Dynamic neural network methods applied to on-vehicle idle speed control. In *Proceedings of IEEE*, volume 84. IEEE, October 1996.

[Ser00] Cambridge University Technical Services. Procedure for controlling internal combustion engines. British Patent Application, 2000. number 0003871.1.

[SPW+98] A. Smerlas, I. Postlethwaite, D.J. Walker, M.E. Strange, J. Howitt, and R.I. Norton. Design and flight testing of an $\mathcal{H}_\infty$ controller for the nrc bell 205 experimental fly-by-wire helicopter. In *Proceedings of AIAA Guidance, Navigation and Control Conference.* AIAA, 1998.

[Vin93] G. Vinnicombe. *Measuring the Robustness of Feedback Systems.* PhD thesis, University of Cambridge, 1993.