COMPENSATING FOR PLANT NONLINEARITIES: GAIN-SCHEDULING OR FEEDBACK LINEARISATION ?

D.J.Leith, W.E.Leithead

Department of Electronic & Electrical Engineering,
University of Strathclyde, 50 George St,
GLASGOW G1 1QE, U.K.

Abstract: Controllers designed by the gain-scheduling and feedback linearisation methodologies are analysed and compared for a class of SISO plants which encompasses a wide variety of practical applications. These controllers are, on the face of it, quite different both in structure and with regard to their range of applicability. Nevertheless, the feedback linearisation and gain-scheduling solutions are, with regard to their linearising action, shown to be equivalent. The equivalence does not, however, extend to the robustness and ease of implementation. In particular, full plant state information is required for feedback linearisation but not gain-scheduling.

Keywords: nonlinear control, feedback linearisation, gain-scheduling

1. INTRODUCTION

Every real system is nonlinear and, in many applications, a linear controller is unable to achieve the performance specification or, indeed, ensure stability, across the entire operating envelope. In order to accommodate the nonlinear characteristics of the plant some form of nonlinear control is necessary. However, whilst a large body of both practical and theoretical knowledge has been accumulated concerning the design and analysis of linear controllers, design methods for nonlinear controllers are less well developed.

A widely employed approach to nonlinear controller design is so-called 'gain-scheduling'. This provides a conceptually simple approach to nonlinear controller design whereby the nonlinear controller is constructed by continuously interpolating, in some manner, between the members of a family of linear controllers. Each linear controller is typically designed on the basis of a linearisation, at a specific equilibrium operating point, of the plant. Whilst gain-scheduled controllers are traditionally adjusted with reference to a slowly varying externally measured quantity, the term is also widely applied to encompass a broad range of systems including those for which the scheduling variable varies rapidly and/or where an internal state of the controller is employed to implicitly schedule the system. Although the dynamic characteristics of these latter systems may, in general, bear little relation to those of slowly-varying exogeneously gain-scheduled controllers, the controllers are linked by the design approach employed. An attractive feature of the gain-scheduling design approach is that it builds upon, and provides continuity with, conventional, well-validated, linear control methods. Gain-scheduled controllers have been successfully employed in a wide range of applications over many years.

More recently, feedback linearisation methods have been developed for accommodating plant nonlinearities. These methods are based on the determination of a linearising transformation which, essentially, inverts the nonlinear plant behaviour and lead to a two-step nonlinear control design: first, the nonlinear plant characteristics are (nonlinearly)

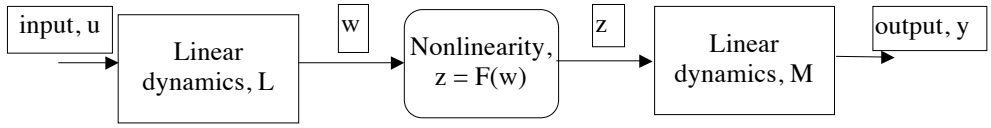

Fig. 1 Class of plants considered.

compensated and, second, a linear controller is designed for the resulting linearised plant dynamics. However, whilst great interest exists in this approach it is, as yet, less widely applied than gain-scheduling.

In this paper, controllers designed by the gain-scheduling and feedback linearisation methodologies are analysed and compared for SISO plants which may be represented in the form depicted in figure 1, where the static nonlinearity, $F(\bullet)$, is invertible and L is stable and minimum-phase. A wide variety of plants fall into this category, including wind turbines (Leith & Leithead 1995, 1996a), neutralisation processes (Klatt & Engell 1995), distillation columns (Shinskey 1987, Tsogas & McAvoy 1982), magnetic suspension systems (Fujita *et al*. 1995) and arc welders (Bjorgvinsson *et al*. 1993). The results are illustrated with reference to a typical wind turbine regulation problem.

The paper is organised as follows. The wind turbine example is briefly described in section 2. Gain-scheduling and feedback linearisation solutions for this example are discussed, respectively, in section 3 and 4. It is noted that the results of these sections, whilst illustrated with reference to the wind turbine example, are quite general. In section 5 the relationship between the gain-scheduling and feedback linearisation approaches is investigated and the conclusions are summarised in section 6.

2. WIND TURBINE EXAMPLE

Constant-speed pitch regulated wind turbines are of the form depicted in figure 1, where the nonlinear behaviour arises from the aerodynamic characteristics of the rotor. For example, a typical two-bladed medium-scale machine (Leith & Leithead 1995, 1996b) has the following parameters,

	TRANSFER FUNCTION
L	$20.7/(s+20.7)$
M	$\dfrac{46460.9}{(s^5+81.27s^4+3683.90s^3+120773.6s^2+1474504.0s+36857450.0)}$

$\nabla F(w) \approx 1254.92w + 250.89, \quad w \in [0, 25]$ degrees

In addition to compensation for the nonlinear aerodynamics, the controller incorporates linear loop-shaping dynamics with the transfer function,

$$871.229 \frac{(s+1.6)^2(s^2+7.243s+38.637)}{s(s+0.3)(s+3.7)(s+20)(s+50)\,(s^2+11s+104.04)} \times \frac{(s^2+1.5s+104.04)(s^2+6s+416.16)}{(s^2+10s+416.16)(s^2+65.8s+2209)}$$

3. GAIN-SCHEDULING

Locally to an equilibrium operating point at which w equals w_o,

$w = w_o + \delta w$, $u = u_o + \delta u$, $z = z_o + \delta z$, $y = y_o + \delta y$ (1)

and the wind turbine dynamics can be approximated by the linear system,

$$\delta y \approx M\, \nabla F(w_o)\, L \delta u \qquad (2)$$

The gain-scheduling approach to compensation of the plant nonlinearity is to incorporate the reciprocal of the 'gain' $\nabla F(x_o)$ within the controller and schedule this gain with respect to a variable which parameterises the equilibrium operating points. In the present case, the input to the plant nonlinearity, w, is an obvious choice for the scheduling variable.

It is known that, in general, the realisation chosen for a gain-scheduled controller can have a substantial influence on the controller performance (Leith & Leithead 1995, 1996a, b). For example, in the wind turbine case it can be seen from figure 2 that the performance is rather sensitive to the position of the scheduled gain relative to the pure integrator

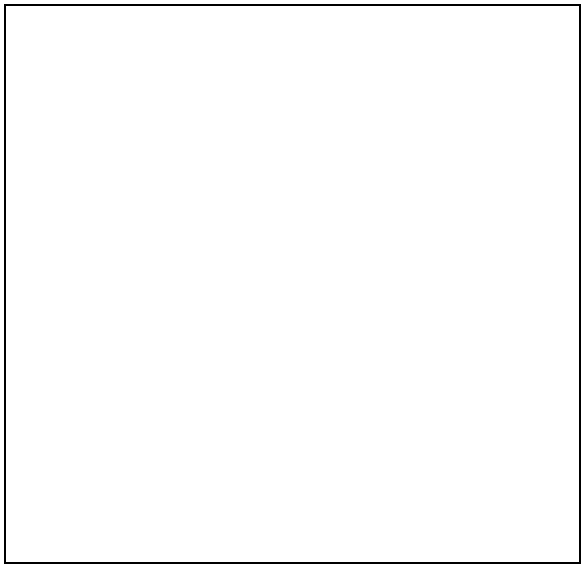

Fig. 2 Typical power time histories with nonlinear gain positioned before & after pure integrator in controller.

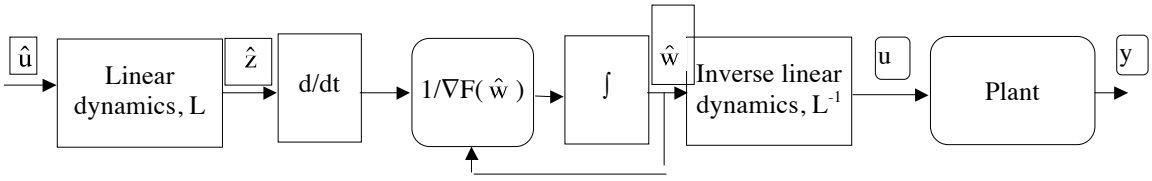

Fig. 3a Gain-scheduling solution

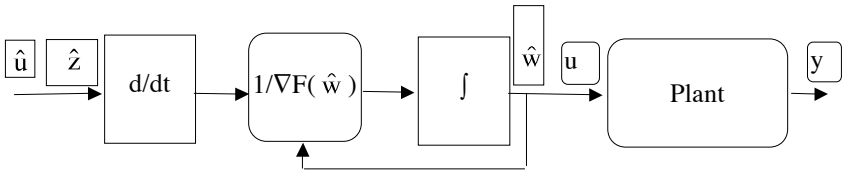

Fig. 3b Simplified gain-scheduling solution

within the controller. Leith & Leithead (1994, 1995, 1996b) propose that an appropriate gain-scheduled controller realisation should satisfy an extended local linear equivalence condition. (Briefly, the gain-scheduled controller is obtained by interpolating between the members of a family of linear controllers designed for specific operating points. The extended local linear equivalence condition requires that the linearisation of the gain-scheduled controller should match the appropriate member of the linear family at every operating point, both equilibrium operating points and non-equilibrium operating points). In the present context, with w the scheduling variable, this is achieved by placing a pure integrator after the scheduled gain and incorporating a model of the linear dynamics, L, and its inverse[1] in the controller as depicted in figure 3a. Of course, when the dynamics, L, have a sufficiently high bandwidth they may be neglected and the gain-scheduled controller reduces to figure 3b. Typically, in order to achieve adequate disturbance rejection and steady-state tracking, the controller contains integral action, in which case the derivative element introduced in this realisation is unnecessary. Otherwise, a proper controller realisation may be attained by reformulating the controller in accordance with the discussion in section 5 below (see, also, Leith & Leithead 1996b).

The gain-scheduling approach of figure 3, or approaches which are essentially equivalent, are widely employed in wind turbine regulation to compensate for the plant nonlinearity (Leith & Leithead 1995, 1996a). The validity of this choice of realisation is confirmed by empirical studies, and the resulting gain-scheduling solution is remarkably successful in practice (Leith & Leithead 1995, 1996a).

[1] When a first order model of the linear dynamics, L, is sufficient, the inverse model is combined with the pure integrator within the controller to attain a proper realisation. When a higher-order model is necessary an approximate, realisable, inverse model is employed (Leith & Leithead 1996b).

4. FEEDBACK LINEARISATION

Whilst the foregoing gain-scheduling solution is successful in practice, it can be seen that the plant 'gain', ∇F(w), varies by more than an order of magnitude over the operating range of w. Moreover, the wind speed fluctuations are highly stochastic and the operating point of the wind turbine varies rapidly and continuously over the whole operating envelope. Whilst the bandwidth of the linearised closed-loop system is typically about 3 r/s, the operating point might cover its full range, corresponding to an order of magnitude or greater change in the aerodynamic gain, in one or two seconds. In addition, large, rapid fluctuations in wind speed are common, in particular gusts; that is, steady increases or decreases in the wind speed which persist for relatively long periods and produce substantial and prolonged perturbations from equilibrium. Hence, *a priori*, it cannot be assumed that the system is slowly varying as required for conventional gain-scheduling analysis.

Feedback linearisation (see, for example, Isidori 1989, Slotine & Li 1991) is an alternative approach to compensating for plant nonlinearities which does not require that the system is slowly varying and which may, therefore, be more appropriate for wind turbine regulation.

Notation: Let $h:\Re^n\rightarrow\Re$ be a smooth (sufficiently differentiable) scalar function and $\mathbf{f}$, $\mathbf{g}:\Re^n\rightarrow\Re^n$ be smooth vector fields on $\Re^n$. The Lie derivative, or directional derivative, with respect to $\mathbf{f}$ is the scalar function defined by $L_f h = \nabla h\cdot\mathbf{f}$. Higher order Lie derivatives are defined recursively by $L_f^i h = L_f^o L_f^{i-1} h$, $L_f^o h = h$; similarly, $L_g L_f h = \nabla L_f h\cdot\mathbf{g}$.

The linear plant dynamics, L, relating w and u can be described by,

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})u, w = \mathbf{C}\mathbf{x} \qquad (3a)$$

with $\mathbf{f}(\mathbf{x}) = \mathbf{A}\mathbf{x}$ and $\mathbf{g}(\mathbf{x}) = \mathbf{B}$. Hence,

$$z = F(w) = F(\mathbf{C}\mathbf{x}) = H(\mathbf{x}) \qquad (3b)$$

When u satisfies

$$u = \frac{-L_f^r H + v}{L_g L_f^{r-1} H} \qquad (4)$$

where r is the relative degree of the system (defined by $L_g L_f^i H=0$ for $0 \le i < r-1$, $L_g L_f^{r-1} H \neq 0$), the dynamics relating z to u are linearised in the sense that

$$z^{(r)} = v \qquad (5)$$

where $z^{(r)}$ denotes $d^{(r)}z/dt^{(r)}$. In order to confine the linearisation to the nonlinearity, F(•), thereby preserving the linear dynamics, L, v may be chosen according to

$$v = \hat{z}^{(r)} - k_{r-1}\varepsilon^{(r-1)} - \ldots k_1 \dot{\varepsilon} - k_o \varepsilon \qquad (6)$$

where, $\dot{\hat{\mathbf{x}}} = \mathbf{A}\hat{\mathbf{x}} + \mathbf{B}\hat{u}$, $\hat{z} = \mathbf{C}\hat{\mathbf{x}}$ and ε equals $z - \hat{z}$. Combining (5) and (6) it can be seen that ε satisfies

$$\varepsilon^{(r)} + k_{r-1}\varepsilon^{(r-1)} + \ldots k_1 \dot{\varepsilon} + k_o \varepsilon = 0 \qquad (7)$$

Hence, for suitable k_i, i = 0, 1, .., r-1, ε decays to zero after some initial transients. It follows that,

$$y = Mz = M(\hat{z} + \varepsilon) = ML\,\hat{u} + M\varepsilon \rightarrow ML\,\hat{u} \qquad (8)$$

as required. The error, ε, is associated with the observable, but uncontrollable, dynamics relating z and $\hat{z}$.

Whilst the choice of v, (6), which gives rise to the error dynamics, (7), is widely employed, it is not, of course, unique and many others also achieve the required linearisation action. For example, when the transfer function of the linear dynamics, L, has no zeroes, the relative degree, r, equals the number of poles and the dynamics can, be described by the differential equation,

$$a_r w^{(r)} + a_{r-1} w^{(r-1)} + \ldots + a_1 \dot{w} + a_0 w = u \qquad (9)$$

Hence, choosing v as

$$v = (a_{r-1} z^{(r-1)} + \ldots + a_1 \dot{z} + a_0 z - \hat{u})/a_r \qquad (10)$$

then the linearised dynamics, (5), are

$$a_r z^{(r)} + a_{r-1} z^{(r-1)} + \ldots + a_1 \dot{z} + a_0 z = \hat{u} \qquad (11)$$

and $y = Mz = ML\,\hat{u}$ as required. Notice that for this choice of v there is no initial transient error, ε, and no auxiliary dynamics, (8). Hence, it is somewhat simpler and more direct than the previous choice of v, (6). However, whilst (10) may be extended to encompass the case when the transfer function, L, involves zeroes, the expression for v then involves derivatives of u and $\hat{u}$. Consequently, (4) is a differential equation in u; that is, the linearisation involves dynamic feedback. Since, in addition, it is usually desirable to avoid measuring or estimating high-order derivatives of $\hat{u}$, the conventional choice of v, (6), is, perhaps, more attractive in these circumstances.

It should be noted that, in general, calculation of the input, (4), requires measurements, or estimates, of the states, **x**, associated with the linear plant dynamics, (3a).

In the present wind turbine example, the dynamics, L, are described by the first order system

$$\dot{w} = -a^{-1}w + a^{-1}u \qquad (12)$$

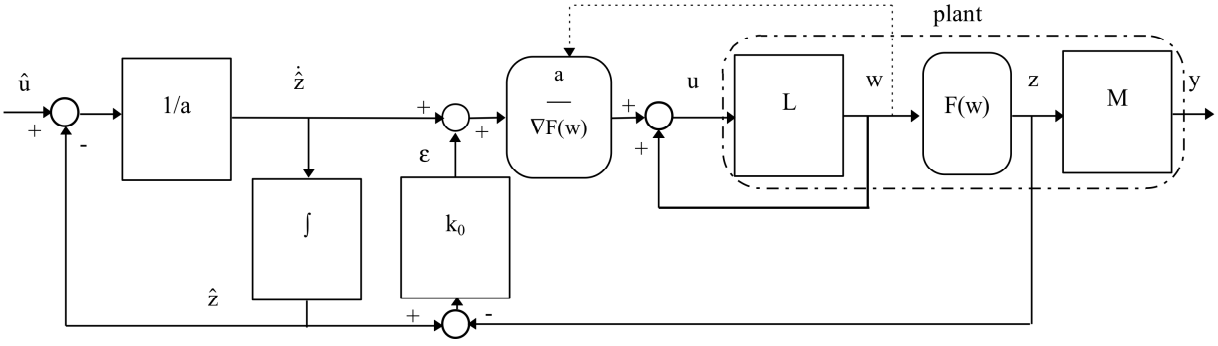

Fig. 4a Feedback linearisation solution for wind turbine example

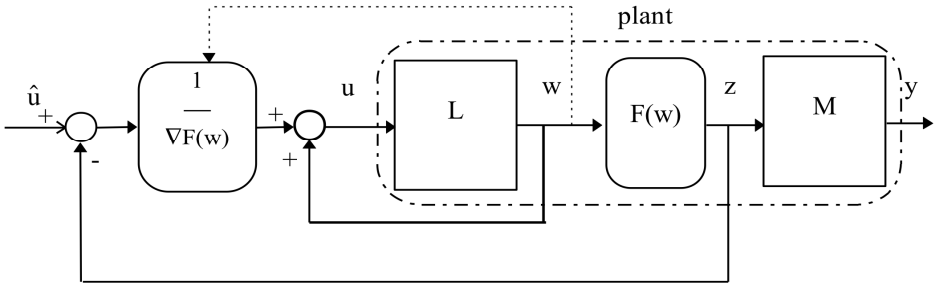

Fig. 4b Simplified feedback linearisation solution for wind turbine example

with $a^{-1}\neq 0$; that is, in (3a), $\mathbf{x} = w$, $f(\mathbf{x}) = -a^{-1}w$, $g(\mathbf{x}) = a^{-1}$, $\mathbf{C}=1$. The first order Lie derivatives of (12) are

$$L_gH = \nabla F(w)\, a^{-1},\ L_fH = -\nabla F(w)\, a^{-1} w \quad (13)$$

From (10), the relative degree, r, of the system is unity when $\nabla F(w)\neq 0$, $a^{-1}\neq 0$. The former condition is necessary for the wind turbine to be controllable, whilst, from (12), the latter is satisfied. Hence, employing the choice of v in (6), the feedback linearising input is

$$u = (a^{-1}\nabla F(w)\, w + \dot{\hat{z}} + k_0\, \hat{z} - k_0 z)/\, a^{-1}\nabla F(w) \quad (14)$$

$$\dot{\hat{z}} = -a^{-1}\, \hat{z} + a^{-1}\, \hat{u} \quad (15)$$

for some suitable, positive, choice of k_0. This corresponds to the block diagram representation depicted in figure 4a.

In the present example, the linear dynamics, L, have no zeroes and the alternative choice of v, (10), leads to the simpler linearisation

$$u = (\nabla F(w)\, w - z + \hat{u})/\nabla F(w) \quad (16)$$

which is depicted in figure 4b.

The linear dynamics, L, model the characteristics of the wind turbines actuator. It is noted that for both choices of input a positive feedback loop encloses the actuator and essentially cancels with the actuator internal feedback to leave it open-loop.

The controllers based on gain-scheduling and feedback linearisation are, on the face of it, quite different both in structure and with regard to their range of applicability. In particular, the gain-scheduled controller design is explicitly based on a family of local linear designs whereas it is difficult to discern any relationship between the feedback linearisation controller and linear control designs. Furthermore, whilst the gain-scheduled controller is based on local plant linearisations and appears to be restricted to situations where the scheduling variable is sufficiently slowly-varying, the feedback linearisation controller is essentially a global one for this example.

Nevertheless, with regard to the wind turbine regulation, where the operating point varies rapidly and continuously over the entire operating envelope, it is observed that the performance of both controllers is indistinguishable (Leith & Leithead 1995, 1996a).

5. RELATIONSHIP BETWEEN GAIN-SCHEDULING & FEEDBACK LINEARISATION

5.1 Wind turbine example

Owing to the positive feedback around the first-order dynamics, L, it follows that for the feedback linearisation approach of figure 4b,

$$\dot{w} = (\hat{u} - z)/a\nabla F(w) \quad (17)$$

Since, from figure 1,

$$\dot{z} = \nabla F(w)\, \dot{w} \quad (18)$$

then it follows that,

$$a\dot{z} = \hat{u} - z, \quad (19)$$

Hence, from (12),

$$z = L\, \hat{u} \quad (20)$$

and, in accordance with the analysis of section 4, linearisation is achieved . The linearising plant input is (figure 4b),

$$u = w + (\hat{u} - z)/\nabla F(w) \quad (21a)$$

$$= w + a\,(L\, \dot{\hat{u}})/\nabla F(w) \quad (21b)$$

With the gain-scheduling approach of figure 3a,

$$w = \hat{w} = \int_0^t \frac{L\dot{\hat{u}}}{\nabla F(\hat{w})}\, d\tau \quad (22)$$

(where any initial transients associated with the dynamics, LL^{-1}, are neglected since they are associated only with controller start-up). Noting that,

$$F^{-1}\big(z(t)\big) = \int_0^t \frac{1}{\nabla F\big(F^{-1}(z(\tau))\big)}\, \dot{z}(\tau) d\tau \quad (23)$$

it follows that,

$$w = F^{-1}(L\, \hat{u}) \quad (24)$$

and so,

$$z = F(w) = L\, \hat{u} \quad (25)$$

Hence, the gain-scheduling approach also achieves, albeit inadvertently, global linearisation of the plant dynamics. Moreover, with the gain-scheduling approach the plant input is,

$$u = L^{-1}\, \hat{w} = \hat{w} + a\, \dot{\hat{w}} \quad (26)$$

Since,

$$\dot{\hat{w}} = \dot{w} = (L\, \dot{\hat{u}})/\nabla F(w) \quad (27)$$

the plant input, u, may be reformulated as,

$$u = w + a(L\, \dot{\hat{u}})/\nabla F(w) \quad (28)$$

which is identical to the feedback linearisation input, (21); that is, the gain-scheduling and feedback linearisation approaches of figure 3a and figure 4b are equivalent (neglecting robustness to plant uncertainty and ease of implementation).

5.2 General case

The analysis of section 5.1 is specific to wind turbine example. In this section, general plants of the form in figure 1 are considered.

Feedback linearisation achieves the open-loop dynamics,

$$y = ML\, \hat{u} + M\varepsilon \quad (29)$$

where the error, ε, is either identically zero (for example, with the choice of v, (10)) or decays to zero after some initial transients (for example, with the choice of v, (6)).

Now, consider the direct linearisation approach, depicted in figure 5, for which the plant input, u, is

$$u = L^{-1}F^{-1}(L\, \hat{u}) \quad (30)$$

and the error associated with the observable, but uncontrollable, states of LL^{-1} is,

$$\varepsilon = z - \hat{z} = F(LL^{-1}\, F^{-1}(\hat{z})) - \hat{z} \quad (31)$$

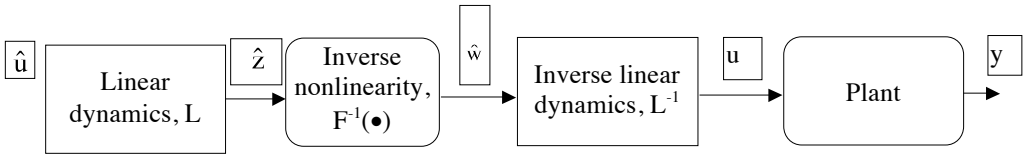

Fig. 5 Direct linearisation

With this direct formulation,

$$\begin{aligned} y &= MF(Lu) = MF(LL^{-1}F^{-1}(L\,\hat{u})) \\ &= MF(LL^{-1}F^{-1}(\hat{z})) = M(\hat{z}+\varepsilon) \\ &= ML\,\hat{u} + M\varepsilon \end{aligned} \qquad (32)$$

Hence, the feedback linearisation solutions may, equivalently, be reformulated as the direct linearisation of figure 5 with appropriate inverse dynamics, L^{-1}. Whilst the feedback linearisation approach is usually expressed in state-space terms, the foregoing relationship provides a, rather straightforward, operator based interpretation.

With regard to the gain-scheduling approach of figure 3, it is noted that

$$F^{-1}\left(z(t)\right) = \int_0^t \frac{1}{\nabla F\left(w(t)\right)}\,\dot{z}(t)dt \qquad (33)$$

Hence, it follows immediately that figure 3 may also be reformulated as the direct linearisation of figure 5 (although, unlike the feedback linearisation approach, the inverse dynamics, L^{-1}, in figure 5 are implemented as a linear system for which the transfer function is the inverse of that of L; when the inverse transfer function is not proper, it is suitably augmented so that L^{-1} is realisable). Consequently, the gain-scheduled controller in fact achieves, albeit inadvertently, global compensation of the plant nonlinearity regardless of the rate of change of the scheduling variable. Of course, this relationship with the direct linearisation of figure 5 is dependent on the choice of scheduling variable and the gain-scheduling realisation adopted. It therefore provides a rigorous basis for the analysis of the influence of the choice of realisation on the performance of the gain-scheduled controller. In particular, for the class of plants considered, it provides rigorous support for realisations which satisfy the extended local linear equivalence condition of Leith & Leithead (1994, 1995, 1996b).

The foregoing analysis indicates that the feedback linearisation and gain-scheduling controllers correspond to particular realisations of the direct linearisation of figure 5 and for both approaches the plant output, y, is related to the input, $\hat{u}$, by the same linear dynamics, ML. Of course, different realisations of the inverse dynamics, L^{-1}, may be employed and, consequently, the initial error, ε, may decay with different transients. Nevertheless, these initial transients, which are associated with controller start-up, may be neglected during normal controller operation. Hence, with regard to their linearising action, during normal controller operation, the feedback linearisation and gain-scheduling realisations are equivalent.

The equivalence does not, however, extend to the ease of implementation of these two realisations. For example, whilst full state information of the plant dynamics, L, is required for the feedback linearisation realisation, the gain-scheduling realisation does not require any plant measurements. The gain-scheduling realisation is, however, of higher order than the feedback linearisation realisation. Nevertheless, since full state information is often extremely difficult to obtain, it is a distinct advantage of the gain-scheduling realisation that it does not require such information.

Moreover, the robustness of the feedback linearisation and gain-scheduling realisations differs. The feedback linearisation realisation requires full state information and its performance is, therefore, influenced by measurement noise and/or errors in the estimates of any states. Other differences between the feedback linearisation and gain-scheduling realisations can also influence their relative robustness. For example, the feedback linearisation solutions of figure 4 enclose the linear dynamics, L, in a positive feedback. These dynamics model the behaviour of the wind turbine actuator and the positive feedback, essentially, cancels with the actuator internal feedback leaving it open-loop; under nominal conditions, the negative feedback of z ensures that z is dynamically linearly related to $\hat{u}$. However, the linearity of the negative feedback is dependent on mutual cancellation, *via* the open-loop actuator dynamics, of the nonlinear functions, F and $1/\nabla F$. Significant nonlinearities, which can be expected to influence the foregoing mutual cancellation in a complex manner, are, typically, present internally within real actuators. It follows that this approach can be expected to be rather less robust, in general, with respect to the actuator nonlinearities, than the gain-scheduling realisation (which benefit from the linearising action on the actuator of its internal high gain feedback).

6. CONCLUSIONS

In this paper, controllers designed by the gain-scheduling and feedback linearisation methodologies are analysed and compared for a class of SISO plants which encompasses a wide variety of practical applications. These methodologies have, essentially,

evolved independently and the controllers are, on the face of it, quite different both in structure and with regard to their range of applicability. In particular, the gain-scheduled controller design is explicitly based on a family of local linear designs whereas it is difficult to discern any relationship between the feedback linearisation controller and linear control designs. Furthermore, whilst the gain-scheduled controller is based on local plant linearisations and appears to be restricted to situations where the scheduling variable is sufficiently slowly-varying, the feedback linearisation controller is a global one for the class of plants considered.

Nevertheless, the feedback linearisation and gain-scheduling (satisfying the extended local linear equivalence condition) solutions are, with regard to their linearising action, shown to be equivalent. Indeed, they correspond to different realisations of a general direct linearisation (figure 5). Whilst the feedback linearisation approach is usually expressed in state-space terms, this relationship provides a, rather straightforward, operator based interpretation. Furthermore, for the class of plants considered, this result establishes a rigorous basis for understanding the widespread success of gain-scheduling controllers observed even when the scheduling variable is rapidly varying. In particular, the influence of the choice of realisation on the performance of gain-scheduled controllers, noted to be of considerable importance in empirical studies, may be rigorously analysed. It is found that, for the class of plants considered, rigorous support is provided for gain-scheduling realisations which satisfy the extended local linear equivalence condition of Leith & Leithead (1994, 1995, 1996b).

The equivalence does not, however, extend to the robustness and ease of implementation of the two realisations. In particular, whilst full plant state information is required for the feedback linearisation realisation, the gain-scheduling realisation does not require any plant measurements. Since full state information is often extremely difficult to obtain, this represents a distinct advantage of the gain-scheduling realisation for many applications. The gain-scheduling realisation is, moreover, rather straightforward and direct.

Of course, the impact of specific implementation and robustness issues is, in general, strongly dependant on the intended application. Hence, the choice of an appropriate realisation, either gain-scheduling or feedback linearisation, is also application dependent.

ACKNOWLEDGEMENTS

D.J.Leith gratefully acknowledges the support of the Royal Society for the work presented here. Both authors also gratefully acknowledge the support of the SERC, DTI (formerly D.En) and ETSU.

REFERENCES

Bjorgvinsson, J.B., Cook, G.E., Andersen, K., 1993, Microprocessor-Based Arc Voltage Control for Gas Tungsten Arc Welding Using Gain-Scheduling. *IEEE Transactions on Industry Applications*, **29**, pp250-255.

Fujita, M., Suzuji, R., Shimemura, E., Schmidt, G., 1995, H_∞ Control of a Nonlinear Magnetic Suspension System with Feedback Linearisation: An Experimental Study. *Proceedings of the European Control Conference*, Rome, pp252-257.

Isidori, A., 1989, *Nonlinear Control Systems: An Introduction* (Springer Verlag).

Klatt, K.U., Engell, S., 1995, Gain Scheduling Trajectory Control of Neutralisation Processes in Continuous Stirred Tank Reactors. *Proceedings of the European Control Conference*, Rome, pp2665-2669.

Leith, D.J., Leithead, W.E., 1994, Application of Nonlinear Control to a HAWT. *Proceedings of the 3rd IEEE Conference on Control Applications*, Glasgow.

Leith, D.J., Leithead, W.E., 1995, An Investigation of the Benefits of Nonlinear Control for Pitch-Regulated Wind Turbines. Industrial Control Centre, University of Strathclyde report prepared under DTI contract no. ETSU W/42/00348/00/00.

Leith, D.J., Leithead, W.E., 1996a, Implementation of Wind Turbine Controllers. To appear, *International Journal of Control*.

Leith, D.J., Leithead, W.E., 1996b, Appropriate Realisation of Gain-Scheduled Controllers with Application to Wind Turbine Regulation. *International Journal of Control*, **65**, 223-248.

Shinskey, F.G., *Distillation Control*, McGraw-Hill, 1977.

Slotine, ‚J.E., Li, W., 1991, *Applied Nonlinear Control* (Englewood Cliffs, N.J.: Prentice-Hall).

Tsogas, A., Mcavoy, T., 1982, Gain-Scheduling for Composition Control of Distillation Columns. *Proceedings of the 32nd Canadian Chemical Engineering Conference*, pp1292-1303.