Proceedings of 3rd
European Control Conference
Rome, Italy, September 1995

STRONGLY INPUT CONSTRAINED NONLINEAR CONTROL OF A HORIZONTAL AXIS WIND TURBINE

D.J.Leith, W.E.Leithead

Department of Electronic & Electrical Engineering,
University of Strathclyde, 50 George St,
GLASGOW G1 1QE, U.K.
Electronic Mail: system@icu.strath.ac.uk

Keywords: nonlinear control, switched control, constrained systems, wind turbines

Abstract

Two novel nonlinear control strategies are presented which substantially enhance the performance of a horizontal-axis constant-speed wind turbine, compared to linear control, by working the actuator at both its hard and soft constrained limits.

1. Introduction

Wind energy is one of the most promising sources of renewable energy for the U.K. and over the last two decades there has been rapid development of the technology. The standard commercial design of turbine is a horizontal-axis grid-connected up-wind medium-scale machine with a rating of approximately 300 kW to 500 kW. The rotor usually has two or three blades and in pitch-regulated machines the pitch angle of either the full-span of the blades, or just the outer tips, can be varied. The control design task for constant-speed pitch-regulated machines is to exploit this capability in order to regulate power output in the face of wind speed variations whilst minimising the load transients and thereby reducing fatigue damage. The objectives for the control system are discussed fully elsewhere [1][2][3].

Since the actuator characteristics are known to impose one of the main restrictions on the performance that can be achieved by a wind turbine controller it is important that the available actuator capability is exploited as fully as possible. The purpose of this paper is to investigate two novel nonlinear control strategies which improve performance by working the actuator at its constraints. A typical medium-scale 300 kW three-bladed constant-speed full-span pitch-regulated grid-connected wind turbine with an electro-mechanical actuator is considered. (It is straightforward to extend the methods presented to other configurations, including those with hydraulic actuators). A block diagram representation of the linearised wind turbine model is depicted in figure 1. There are hard limits on the position, velocity, and acceleration/torque which may be developed within the actuator. In addition, there is a restriction on the rate of work of the actuator in order to prevent overheating. This is a soft constraint in the sense that it does not directly limit the magnitude of quantities within the actuator, but rather places a restriction on the standard deviation of the actuator current (equivalently, acceleration/torque). Hence, while it is not acceptable to work the actuator continuously at its hard velocity/acceleration/torque constraints due to the overheating that would result, it is possible to intermittently work the actuator at these limits for short periods.

The paper is organised as follows. Section Two outlines the controller specification. In Sections Three and Four, the proposed nonlinear control strategies are discussed. In Section Five, results of extensive simulations are used to compare the performance of these controllers with one another and with conventional linear control. Conclusions are drawn in Section Six.

2. Controller Specifications

It is important that fair comparisons are made. To this end each controller studied is required to have similar stability margins and to operate within the same actuator restrictions, namely :

1. Gain margin of at least 10 dB
2. Phase margin of approximately 60 degrees.
3. Actuator pitch acceleration standard deviation no more than approximately 20 deg/s^2.
4. Maximum actuator pitch acceleration of 90 deg/s^2.
5. Maximum actuator pitch velocity of 15 deg/s.
6. Pitch angle is confined to the range [0,45] degrees.

The aerodynamic behaviour of wind turbine blades is highly nonlinear and strongly dependent on wind speed. In terms of a linearised plant description, as wind speed increases the gain of the plant increases since the rate of change of aerodynamic torque to pitch angle increases. It is standard practice for wind turbine controllers to include a nonlinear gain to compensate for this variation and make the control task essentially linear [1][2]. However, the representation of the aerodynamics is very basic and subject to considerable uncertainty. In addition, the gain of the controller is incapable of always being scheduled to match the varying wind speed. Consequently, good gain and phase margins are required to achieve adequate stability margins. If these are not achieved, the system must sometimes destabilise, although not necessarily become unstable, in which case the wind turbine would experience large load fluctuations.

In the context of the actuator model employed here, requirements 3, 4, and 5 impose limits which are typical of those present on comparable commercial machines. It should be emphasised that the pitch acceleration and velocity considered is not that of the actual turbine blades. Rather it is a normalised measure which permits valid comparisons to be made between differing designs of actuator. For example, blade pitching systems with different gearing ratios linking the actuator to the blades may be compared in an unbiased manner using this measure.

The upper limit on blade pitch angle specified in requirement 6 is rarely encountered in practice and may be neglected. The lower limit is encountered when the wind speed falls below the level at which rated power may be generated. When this occurs control action is suspended until the wind speed rises again. Smooth and timely start-up of the controller is important, and is discussed further in [1,2] together with various other implementation issues.

3. Maintaining Actuator Activity With Wind Speed

As wind speed rises, a linear controller places less demand on the actuator since the sensitivity of the aerodynamic torque to pitch changes increases faster than the sensitivity to wind speed changes. Hence, for a controller with fixed open-loop cross-over frequency, while the actuator may be worked to its full level at low wind speeds, it may not be used as fully at higher wind speeds. However, it is at these higher wind speeds that loads are greatest and therefore controller performance is most critical. The requirement is to design a controller which works the actuator as near as possible to its continuous operating limit in all wind speeds, counteracting the effect of variations in the sensitivity of aerodynamic torque to pitch changes. A complication is the lack of a measurement of wind speed. Indeed there is no such thing as 'the windspeed' experienced by a wind turbine, since the rotor experiences a spatially and temporally distributed wind field. Simple scheduling is therefore not appropriate, and the wind speed must be inferred from the plant dynamics via the pitch demand. If the controller is operating correctly, the demanded pitch angle is a good indicator of wind speed. (This approach is widely used to vary the previously noted nonlinear gain which compensates for variations in the aerodynamic torque sensitivity). Employing an internal state of the system, such as the pitch demand, to implicitly change the controller as wind speed varies must be treated with some caution, however, since it introduces additional feedback loops, thereby changing the plant dynamics. The design task is to develop a varying controller which induces the required closed-loop dynamics at any wind speed, despite the presence of the feedback loops. The resulting controller is nonlinear.

A family of linear controllers is designed for various wind speeds using classical loop-shaping design techniques. Some care was taken to minimise the differences between these controllers so that interpolation between them could be carried out as smoothly as possible. The continuous family of controllers thereby generated is:

$$g\frac{(s^2+4s+6.25)}{(s^2+as+b)}\;\frac{(s^2+7.12s+79.21)}{s(s+0.3)(s^2+140s+2500)(s+50)}$$

where

$$a = 1.839p + 4.406$$
$$b = 0.284p^2 + 14.443p + 111.320$$
$$g = 0.0296p^2 + 1.619p + 2.220$$

and p is the pitch angle demanded by the controller, in degrees. Upper and lower bounds are placed on a, b and g. When p is less than 4.51 degrees (corresponding to 12 m/s wind speed), a, b and g are held at their 4.51 degree values. Similarly, when p is greater than 25.89 degrees (24 m/s wind speed), a, b and g are held at their 25.89 degree values. It can be seen that these controller transfer functions differ only by a varying gain and a pair of varying poles.

A nonlinear controller is obtained by interpolating continuously between the members of the family of linear controllers as pitch demand varies. See [4,5,6] for a more detailed discussion and analysis of this nonlinear control strategy.

The controller may be simplified to a switched linear dual-mode type of controller, where at low wind speeds one linear controller is used, and at some point a switch is made to a second linear controller for use in higher wind speeds. Switched linear dual-mode controllers for the wind turbine application are considered further in [5][6].

4. Peak Rejection

The nonlinear controller design described in Section Three ensures that the average rate of work of the actuator, as measured by the standard deviation of the pitch acceleration, is kept near 20 deg/s^2 in all wind speeds. However, as noted previously it is possible to intermittently demand a high level of activity for short periods, taking the actuator up to its hard velocity and/or acceleration constraints. This may be exploited to improve the control system response to the worst peaks in the power output by working the actuator temporarily at its maximum level when the start of an unacceptably high peak in the power output is detected, therefore responding as rapidly as possible to the disturbance. The result is a switched nonlinear controller. It is well known that it is the occasional extreme loads experienced by a wind turbine during its life which contribute most to fatigue damage (see for example [7]) and this is also the case in many other applications. Any reduction in these loads is therefore strongly motivated.

The following design issues must be considered when implementing this switched nonlinear control strategy:

1. Working the actuator at its maximum level for the full duration of a power excursion will, in general, result in a degradation of the system stability margins. The design of suitable switching rules is therefore necessary.

2. The baseline controller is required to contain a pure integrator term in order to achieve acceptable disturbance rejection. If the actuator saturates at its hard limits, a loss of performance and a reduction in stability margins can result from wind-up of the pure integrator term. In order to compensate, anti-wind-up measures must be taken.

3. Since the controller reacts more strongly to peaks in the power than to troughs, the skew of the probability distribution of the power is increased. This may depress the mean power output, which is clearly undesirable from an economic standpoint. The switching rules should ensure that this effect is restricted to an acceptably low level.

4. A predictor may be used to permit an earlier reaction to a peak in the power.

5. For the actuator being considered, the limit on the maximum acceleration forms the most restrictive constraint over the range of time scales of the power peaks experienced during normal operation. Hence, a maximal actuator response may be attained by operating at the maximum acceleration level.

4.1 Control Switching Rules

As noted above, it is essential to carefully consider the switching arrangements in order to maintain adequate stability margins. It is known from simulation studies that near sinusoidal limit cycles occur when the switched nonlinear closed-loop system is caused to destabilise. Describing functions are therefore appropriate for determining the stability margins. An analytical method for calculating the controller describing function is given in [8], but for flexibility a numerical approach is mainly used in the present study. A control switching rule of the following form is used:

if ($X<e_0$ AND $E < e_0$) OR ($Z<z_0$)
then use normal controller
else demand maximum pitch acceleration/velocity

where E is the power error and X, Z are obtained by filtering E appropriately. X is to be interpreted as a prediction of the future power error, E, while Z is an indicator of the rate of change of the E with time. If either E or X rise above their thresholds it is taken that a peak in the power is occurring, or about to occur. The threshold of Z is typically negative, so that if Z falls below this value it indicates that linear control may be resumed since the power is dropping sufficiently quickly. The value of this particular threshold clearly has a strong impact on the stability margins of the system.

For a given switching rule and a sinusoidal power error it is useful to plot against frequency the point on the waveform, if any, where maximum pitch acceleration is demanded, and the point where normal linear control is resumed. Such plots, for a sinusoidal power error with a range of amplitudes similar to the peaks encountered in normal operation, aid in designing the filters for X and Z since the interaction between switching and the gain and phase of the filters is clarified. The filter for X should be chosen to give the earliest possible switching at those frequencies where most peaks occur, subject to stability requirements and maintenance of mean power output close to the rated value. Describing function analysis determines, as a function of limit cycle amplitude, the frequencies at which the open-loop gain is unity and the frequencies where the phase is -180 degrees. At these frequencies, the filter for Z should give an early return to normal control in order to maintain adequate stability margins. An early return may also be desirable at other frequencies in order to restrict any reduction in the mean power output as a result of the strong reaction to peaks. Experience suggests that in order to achieve the best performance switching should occur earlier rather than for longer.

Employing the nonlinear controller described in Section Three as the baseline controller in conjunction with the predictor described in Section 4.2 as the filter for X, an appropriate filter for Z has the following transfer function :

$$1100 \frac{(s^2+0.8s+0.64)}{(s^2+25s+625)}$$

With a threshold of 45kW for e_0, and $-100x10^3$ for z_0, the gain margin is 10.2 dB and the phase margin is 52.1 deg (calculated by describing functions and confirmed by simulation). A plot of the switching points versus frequency for a power error amplitude of 70 kW is given in figure 2. For the wind turbine considered this amplitude lies within the range of higher magnitude peaks typically experienced.

4.2 Predictor

When a predictor is incorporated into the controller the predictor becomes part of a feedback loop where its output influences the power output which in turn acts upon the predictor. An analysis of the behaviour of a predictor in such a loop is not straightforward for a nonlinear system, and it is difficult to assess the performance without carrying out a full nonlinear simulation. In addition to anticipating the power output, the predictor is subject to requirements that high frequency noise should not be magnified too greatly and that, for power peaks which exceed the controller switching threshold, the predicted peaks should ideally also not exceed the threshold. Unnecessary controller activity is then kept to a minimum. The following first-order filter with phase lead concentrated from 3-6 rad/s satisfies these requirements.

$$0.64 \frac{(0.4s+1)}{(0.09s+1)}$$

4.3 Anti-Wind-Up Measures.

A conventional anti-wind-up scheme is used to compensate for the loss of stability and performance that would otherwise occur when the actuator saturates (see for example [9]).

4.4 Method of Demanding Maximum Actuator Activity

Whilst it is possible to send a signal directly to the actuator to immediately demand maximum acceleration, an alternative scheme is adopted where the signal is produced within the controller and passed to the actuator via the pitch velocity demand signal. To ensure a rapid response, the velocity demand signal is generated from a simple demand for maximum acceleration passed through a filter which is an approximate inverse of the actuator dynamics. This approach has several advantages. First, no actuator modifications are necessary. Second, analysis is tractable. Third, spurious brief demands for maximum activity, due for example to measurement noise, are filtered out by the pure integral term in the controller. Fourth, since demands for maximum actuator activity are not immediately achieved, the response to power excursions of longer duration is greater than to short duration peaks. This improves performance both by helping to maintain the mean power output at 300 kW, and by improving stability robustness and therefore allowing less conservative switching rules to be used.

5. Controller Performance

A well-validated simulation methodology is used to assess the performance of the nonlinear controllers. The performance achieved is compared with that of a conventional linear controller designed using classical loop-shaping methods [1] to meet the same specifications and possessing the following transfer function:

$$9.6916 \frac{(s^2+4.753s+5.8806)(s^2+6.403s+51.29)}{s(s+0.3)(s+15)(s+18)(s+20)(s+50)^2}$$

(gain margin 9.99 dB, phase margin 56.14 degrees, cross-over frequency 2.78 r/s).

Simulations are run with the controllers over a range of wind speeds and turbulence levels to reproduce the real machine conditions noted in [10], and to predict performance at higher wind speeds. Four mean wind speeds of 12, 16, 20 and 24 m/s are used at three nominal turbulence levels of 10, 15 and 20 %. The simulations are run for 260 seconds, to provide four one minute periods of data. The nominal turbulence level only applies over a long time period, and the range of turbulence levels for the one minute samples was 6 - 26 %. For each one minute sample within some specified turbulence range, the maximum power is found and a linear fit to these one minute sample power maxima provides an indication of the trend with wind speed. In addition, if the standard deviation of the residues of the maxima about the linear fit is determined, then the power maxima experienced under normal operating conditions are unlikely to exceed the linear fit by more than three times the standard deviation [10]. These runs produce only 48 data points to cover the whole operational range of the machine, but this approach has nevertheless been found empirically to be a good indicator of the comparative performance between controllers [10].

The plot in figure 3 shows the 3 standard deviation line associated with each controller. The linear controller's maxima are greatest and increase at the fastest rate, followed by the nonlinear controller and then the switched nonlinear controller. Note that a linear fit is not used for the switched nonlinear controller. Due to the large reduction in the magnitude of the power maxima experienced at high wind speeds a quadratic fit is more appropriate. It can be seen from figure 4 that a small reduction of around 3kW in the mean output power is associated with this improved performance, but is confined to very high wind speeds. The pitch acceleration standard deviations are shown in figure 5. All of the controllers can be seen to respect the actuator restrictions. With the linear controller the pitch acceleration standard deviation falls as wind speed rises, due to the increase in the sensitivity of the aerodynamic torque to pitch changes. In contrast, the standard deviation for the nonlinear and switched nonlinear controllers remains roughly constant as wind speed rises, exploiting the extra actuator capacity available at higher wind speeds, as intended. The switched nonlinear controller produces only a small increase in actuator activity at high wind speeds compared to the nonlinear controller.

6. Conclusions

For a typical three-bladed configuration of wind turbine, extensive simulations using a well validated assessment methodology indicate that significant performance improvements may be gained by exploiting the actuator more fully when compared with a conventional linear controller. In particular, both the peak power and rate of increase of peak power with wind speed are substantially reduced with a consequent reduction in drive-train loads. All physical actuators impose constraints on the control input signal that may be applied to the plant, and the application considered in this paper involves constraints which are typical of those encountered in many real situations. The control strategies presented permit these constraints to be incorporated into the controller design, leading to an improvement in the performance achieved.

Acknowledgements

The SERC, DTI (formerly D.En) and ETSU are gratefully acknowledged for supporting the work presented here, and by whose permission this paper is published.

References

1.Leithead, W.E., de la Salle, S.A., Reardon, D.L., Grimble, M.J., 'Wind Turbine Control Systems Modelling and Design Phase I and II', DTI Report No. ETSU WN 5108, 1991.

2.Leithead,W.E., de la Salle, S.A., Reardon, D., 'Classical Control of Active Pitch Regulation of Constant Speed HAWTs', Int. J. Control, 1992, Vol. 55, pp845-876.

3.Leithead, W.E., de la Salle, S.A., Reardon, D., 'Role and Objectives of Control for Wind Turbines', IEE Proc. C, 1991, Vol. 138, No. 2, pp135-148.

4.Leith, D.J., Leithead, W.E., 'Application of Nonlinear Control to a HAWT', Proc. 3rd IEEE Conference on Control Applications, Glasgow, August 1994.

5.Leith, D.J., Leithead, W.E., 'Comparison of Various Control Strategies for a Two-Bladed Wind Turbine', Proc. British Wind Energy Association Conference, Stirling, June 1994.

6.Leith, D.J., Leithead, W.E.,'Benefits of Optimising the Controller with Windspeed for a Constant Speed HAWT', Proc. European Wind Energy Conference, Greece, Oct. 1994.

7.Warren, J.G. (Ed), 'Prediction of Fatigue Damage in Wind Turbines', report to D.En. under contract no. E/5A/CON/5041/1263,. 1987.

8.Zhiltsov, K.K., 'Harmonic Linearisation of Systems with Logical Control Rule', Automation & Remote Control, 1966, No. 2, pp286-289.

9.Campo, P.J., Morari, M., Nett, C., 'Multivariable Anti-Windup & Bumpless Transfer: A General Theory', Proc. American Control Conference, Pittsburgh, 1989.

10.Bossanyi, E.A., Smith, G.J., Leithead, W.E., Agius, P., 'Design and Testing of a Classical Controller for the MS-3 Wind Turbine', Final Report E/5A/CON/6033/2184, Energy Technology Support Unit, Harwell, Oxfordshire OX11 ORA, U.K., 1992.

Figure 1 Linearised control model

dQ/dp: sensitivity of aerodynamic torque to pitch changes.
dQ/dV: sensitivity of aerodynamic torque to wind speed.

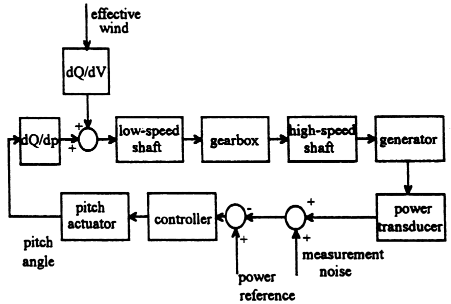

Figure 2 Switching point vs. frequency for a sinusoidal power error of amplitude 70 kW.

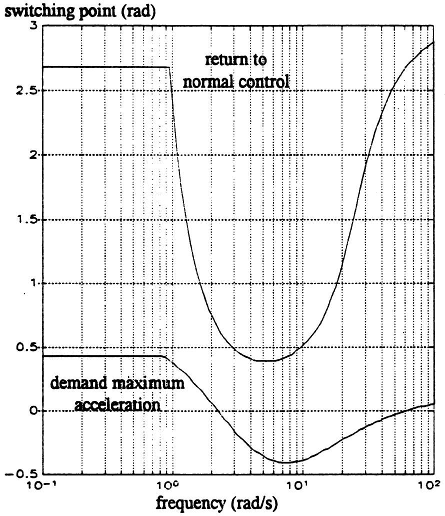

Figure 4 Mean output power
+ nonlinear o switched nonlinear

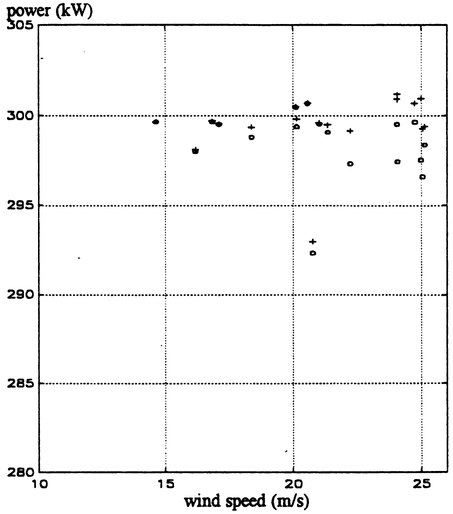

Figure 3 Three standard deviation lines for fits to power maxima

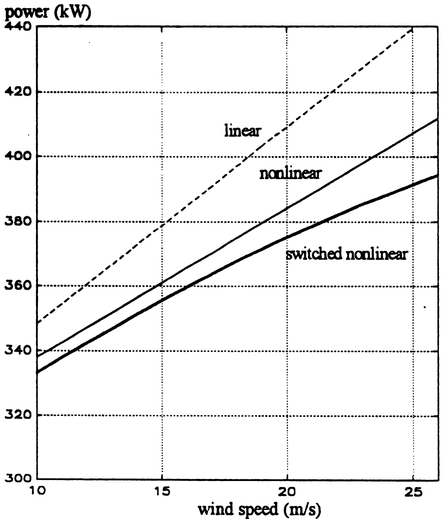

Figure 5 Pitch acceleration standard deviation
+ linear o nonlinear * switched nonlinear