Simulation and Control of Windfarms

Christopher J. Spruce

University of Oxford

Department of Engineering Science

Lady Margaret Hall

Michaelmas Term 1993

Submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy.

Simulation and Control of Windfarms

Christopher J. Spruce

Lady Margaret Hall, Oxford

Submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy at the University of Oxford.
Michaelmas Term 1993

Abstract

This thesis examines the design of supervisory controllers for windfarms of pitch-controlled wind turbines. The control objectives are the maximisation of the financial income from the generated electricity and the minimisation of the turbines' fatigue damage. The design exploits the wide variations in the ratio of financial income to fatigue damage which are found both spatially across windfarms and as a function of time. The supervisory control strategy makes use of the ability of pitch-controlled turbines to operate with variable power set points; a capability which is rarely exploited in practice.

A windfarm simulation which has been developed for the purposes of testing supervisory controllers is described. It is shown that the simulation is a suitable test-bed for this application. Results are presented which demonstrate how the fatigue damage of a turbine's gearbox and structural components vary as functions of the mean wind-speed, turbulence intensity and power set point, both for isolated turbines and for turbines experiencing wake effects.

A lifetime performance function is proposed and 'ideal' power set point curves are evaluated using a genetic search algorithm. It is shown that significant improvements in performance can be achieved if the operation of the turbines is altered to take account of variable electricity tariffs.

A windfarm control strategy that splits the turbines into interacting and non-interacting categories is found to give good results. Using data generated by the simulation, it is shown that simple cost functions can be developed for non-interacting turbines which, when used in a controller, give performance that is close to the 'ideal'. A similar cost function is applied to a group of three interacting turbines, and it is found that substantial reductions in all measures of total annual fatigue damage are achieved for a small reduction in total annual financial income. The on-line implementation of windfarm supervisory controllers is discussed and the behaviour of a simple hill-climbing algorithm is examined using a simulated group of three interacting turbines.

Acknowledgements

I am greatly indebted to my supervisor, Dr. Arthur Dexter, for his interest in the initial idea, for his help in obtaining financial support, and for his constant encouragement, guidance and enthusiasm during the research leading to this thesis.

I should like to express my gratitude to all those at the Wind Energy Group Ltd. who have been involved with the work, for their farsighted and enthusiastic assistance. In particular, I should like to thank Dr. Charles Gamble, Dr. Ervin Bossanyi and Gordon Smith for providing much advice, as well as important data; Dr. Roger Haines, for useful discussions about the 'ideal' power curve technique; and Dr. David Lindley, for his advice and assistance in obtaining industrial sponsorship.

I should like to thank Dr. Bill Leithead, of Strathclyde University's Industrial Control Unit, for many useful discussions on the simulation and control of wind turbines and windfarms, and also for computer code of a model of the turbulent wind. I would like to thank Professor Werner Leonhard of the Institut für Regelungstechnik at the Technische Universität Braunschweig for allowing me to work at the institute as a visiting researcher in 1991. I would also like to thank my colleagues in Oxford for sharing their knowledge and experience, particularly Dr. Phil Haves for assistance with computing facilities and Simon Hepworth for proof-reading the thesis. The financial support of the Science and Engineering Research Council, the Wind Energy Group Ltd. and the Leverhulme Trust is gratefully acknowledged.

Last but not least, I should like to thank my parents for their constant encouragement and support over the years, and I should like to thank Meike for her support and understanding.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Renewable energy sources which are powered by the natural elements have zero input fuel costs and very low operations and maintenance costs. Therefore the capital cost of renewable energy power stations dominates the lifetime cost of the electricity generated, unlike in the case of conventional power stations, whose fuel costs form a significant proportion of the total lifetime cost of electricity. So for renewable energy power stations, the efficiency of operation of the machinery which is purchased with the initial capital has a considerable effect upon the lifetime cost of electricity. Consequently the optimisation of the designs of, and control strategies for, renewable energy power stations has significant potential for reducing the cost of the electricity generated.

In general terms, this thesis is concerned with the development of control strategies for windfarms, with the aim of achieving more efficient operation of the turbines within a windfarm.

A “wind farm” is defined by the Oxford English Dictionary as “a group of energy-producing windmills or wind turbines”. The smallest group of turbines that would generally be referred to as a windfarm consists of three turbines, and the largest would be of the order of hundreds of turbines. When thousands of turbines are involved, then the complete installation is referred to as a number of windfarms grouped together. A picture of windfarms in California is presented in Figure 1.1.

Figure 1.1: Windfarms in California (Source: Windpower Monthly/US Windpower)

The first windfarm to use a group of turbines whose combined electrical output was fed directly into a utility grid was probably the windfarm built in 1980 on Crotched Mountain in New Hampshire, USA, which consisted of twenty 30 kW turbines with a total capacity of 600 kW (Symanski, 1983). In the 13 years since that time approximately 1600 MW of wind power have been installed in California (Ancona et al., 1993) and 1100 MW in Europe (Lindley, 1993), almost all of it in windfarms. As the emergence of wind energy has taken place so recently, the technology of wind energy conversion is still relatively immature.

1.1 Basic operational characteristics of windfarms

In pure economic terms, and excluding wider benefits such as avoided pollution, the closest to a fundamental measure of windfarm performance is probably the discounted cost per unit of electricity (p/kWh) over the lifetime of the windfarm. Therefore one of the crucial factors in achieving high windfarm performance is the ability of the turbines within a windfarm, as a group, to achieve high energy capture for low capital and running costs per turbine. This ability is, in turn, heavily influenced by the operational characteristics of the windfarm, characteristics which differ widely from site to site.

The free stream wind is a windfarm's raw energy input, and therefore the dynamics of the wind-speed in the small section of the atmosphere up to approximately 100 m above ground level are of importance to the operation of windfarms. It has been shown that windfarm energy capture is more sensitive to the free stream wind characteristics than to any other factors affecting energy production (Harris, 1992). The free stream wind is a stochastic process and Figure 1.2 presents the broad range spectrum of wind-speed at a height of 100 m above ground level, taken from the classic study of van der Hoven (Van der Hoven, 1957). The major contributions to the variance of the wind-speed are due to the passing of identifiable weather systems (time period, T, $\approx$ 4 days), diurnal variations ($T \approx 12$ hours) and atmospheric turbulence (T between approximately 5 minutes and 5 seconds)*. The broad 'spectral gap' between the contributions due to diurnal variations and those due to atmospheric turbulence is widely used in wind energy analysis and design to justify

*Van der Hoven notes that "An anomalous wind situation was purposely chosen to represent the high-frequency end of the spectrum ... The data were taken during hurricane Connie ... The effect of the high wind speeds was to increase greatly the amplitude of the high-frequency end of the spectrum. More normal situations ... seldom show a spectral amplitude above 1.0 m^2/s^2 and none above 2.0 m^2/s^2. ... [Figure 1.2] shows a peak amplitude of 3.1 m^2/s^2 in this range".

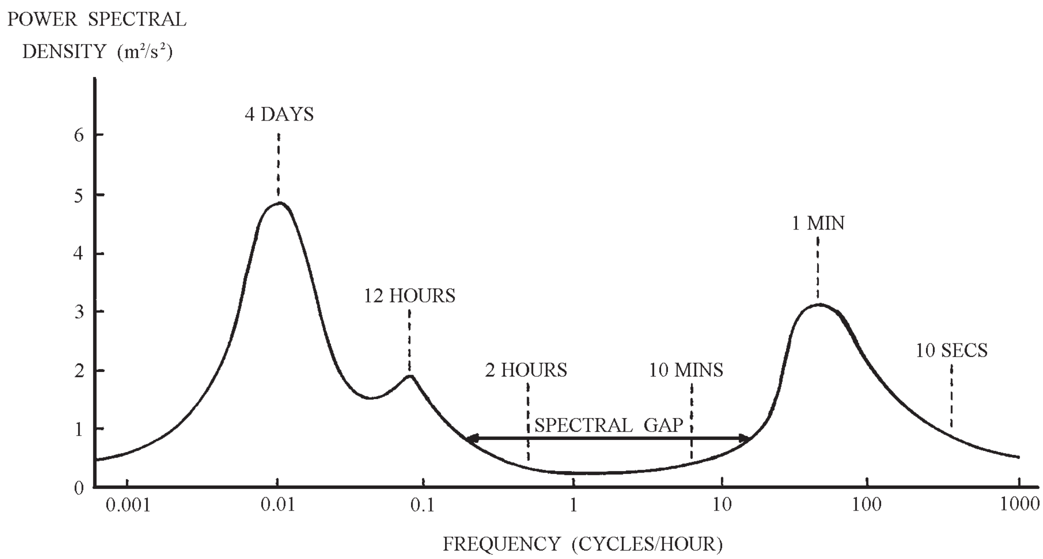

Figure 1.2: Spectrum of the free stream wind-speed at 100 m height, from (Van der Hoven, 1957)

regarding the latter as a stationary process (Fordham, 1985). The wind-speed due to atmospheric turbulence is generally represented by a mean ($\overline{v}$) and a normalised 'turbulence intensity' (I), where

$$I = \frac{\sigma_v}{\overline{v}}$$

Probability distributions such as the Weibull distribution are then used to extend estimates of energy capture or fatigue damage over a year (Eggleston and Stoddard, 1987)[pp. 67*ff*, 311*ff*]. In preparing such estimates, it is usually implicitly assumed that the probability distribution of the wind-speed at a given site is constant from year to year, even though in practice this may not be true (Eggleston and Stoddard, 1987)[p. 71].

To achieve high windfarm performance, the wind turbines should be capable of the highest possible financial income from generated electricity per annum for the lowest possible installation and running costs. Both of these factors depend on the design of the turbines. In this thesis, the turbines that are considered are upwind, horizontal axis machines of 'medium' or 'large' size (greater than approximately 100 kW), as illustrated in Figure 1.3. The configuration consists of a rotor, a gearbox which steps

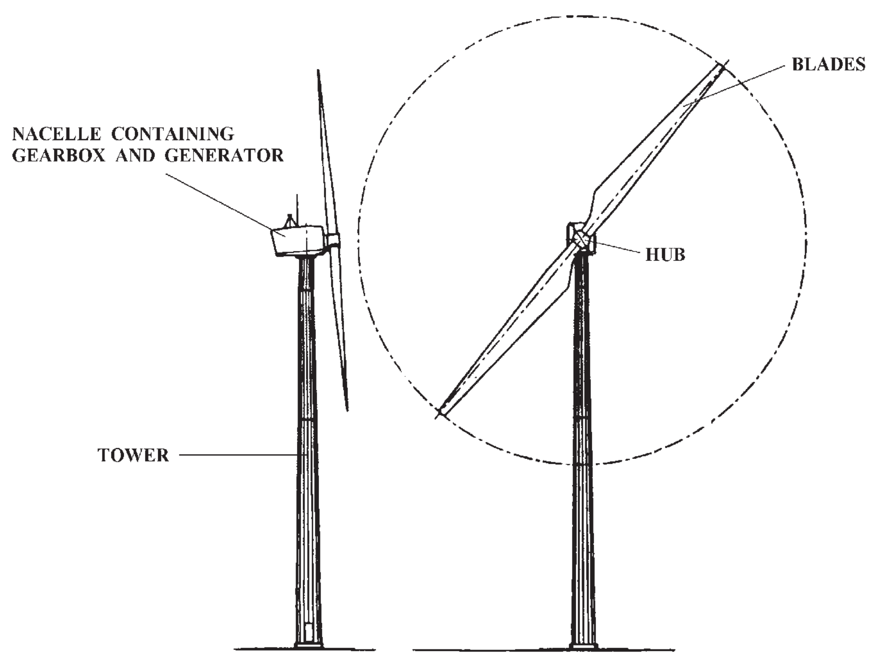

Figure 1.3: Typical medium-sized horizontal axis wind turbine, from (Spruce, 1988b)

up the rotational speed of the drive-train, and an induction generator. The grid to which the turbines supply power is assumed to be infinite, so the connection of the turbines to the grid locks the speed of rotation of the turbine to the grid frequency. As a result, this type of turbine is called a constant speed wind turbine. The great majority of commercial turbines are of this type.

As the wind-speed incident upon a turbine increases, the power in the wind increases as approximately the cube of the wind-speed. Higher wind-speeds are experienced much less frequently than lower wind-speeds, and as a result it is not worth building turbines that are capable of extracting the total energy available in the wind at all wind-speeds. The high ratings that would be required for the components in the turbine would not be economically worthwhile for the small amounts of extra energy that would be generated per annum and, consequently, two measures are taken in turbine design. Firstly, when the wind-speed (v) exceeds a certain level (the rated wind-speed, v_{rated}), aerodynamic power regulation takes place. Secondly, when the wind-speed exceeds a certain high wind-speed (the upper cut-out, v_{hi}) the turbine is shut down altogether. Curves of power set point and mean generated power against average wind-speed are presented in Figure 1.4, illustrating the region where power regulation takes place as well as the upper cut-out. Not all turbines have a power

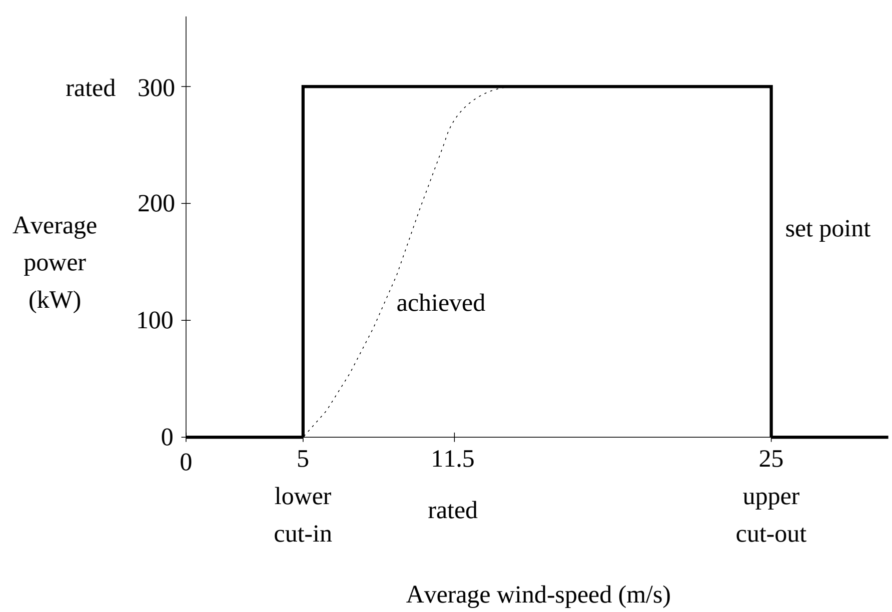

Figure 1.4: Typical power set point and actual mean power curves

set point which can be varied, but those that do almost invariably use a constant set point throughout the turbine's operating range, as illustrated by Figure 1.4.

The aerodynamic power regulation that takes place for $v_{rated} < v < v_{hi}$ is usually achieved in one of two ways. The first is passive power regulation, whereby the blades are designed to stall at close to rated wind-speed, and the average power output then remains equal to the rated power for $v > v_{rated}$. The second is active power control, whereby the pitch angle of the blades is varied continuously to maintain the power as close to the rated power as possible. If the blades are feathered, for example, their lift forces decrease and hence the rotor torque decreases. There are two different ways, mechanically, of achieving pitch control. Either the whole length of the blades are pitched (full-span pitch control) or only part of the blades are pitched (tip pitch control). The adjustment of the pitch angle is usually carried out on the basis of a power measurement.

Both the power regulation that takes place when $v > v_{rated}$ and the shut-down of the turbine when $v > v_{hi}$ are critical design choices which involve trading off the energy capture against fatigue loads. The energy capture affects the income produced by the turbine and the fatigue loads affect the capital cost of the turbine. The trade-off is not made directly, but via iterations through the design process and by evolution from one generation of turbine to the next. A further trade-off between energy capture and fatigue loadings is made when choosing the lower cut-in wind-speed. The turbine should not start generating unless the power in the wind is greater than the losses experienced in the drive-train and, as the wind is a stochastic process, the power in the wind at low wind-speeds rises and falls above the necessary threshold. To maximise energy capture, the turbine would come on-line every time the available power rose above the threshold and would shut down every time the power fell beneath it again. This would give rise to a large number of start/stops, with an associated high level of fatigue damage, and therefore such an approach is never carried out. Instead a qualitative trade-off between energy capture and fatigue damage is made by specifying a cut-in wind-speed which is greater than 0 m/s.

The efficiency with which the turbines in a windfarm convert the kinetic energy in the wind into electrical energy has a significant effect upon windfarm performance. This efficiency is very sensitive to the wind conditions that are experienced by the turbines. It is measured by the power coefficient C_p, which is the ratio of the power produced by the turbine's rotor to the power in the wind that is incident upon the turbine's rotor (Eggleston and Stoddard, 1987)[p. 30]:

$$C_p = \frac{P_{aero}}{\frac{1}{2}\rho A_r v^3}$$

P_{aero} is the power produced by the rotor; ρ is the density of air; A_r is the cross-sectional area of the rotor; v is the wind-speed.

Power coefficient data is normally presented as a function of the fundamental speed parameter for wind turbine rotors, the tip-speed ratio λ, which is the ratio of the velocity of the blade tip (ωR) to the velocity of the incident wind:

$$\lambda = \frac{\omega R}{v}$$

A typical C_p–λ curve is presented in Figure 1.5, for tip-speed ratios across the operating range of wind-speeds. The figure illustrates the substantial variation in C_p across the operating range.

Maximising the performance of a windfarm depends on maximising its financial income. For the purposes of wind turbine design, the maximising of financial income is always (implicitly) assumed to equate with maximising the efficiency with which the kinetic energy in the wind is converted into electrical energy. This assumption is not, however, usually valid. There is a second issue which should be taken into account, that of the tariff paid by the electricity utilities for the power that is generated by windfarms. If a tariff varies with time, then the variable to be maximised, the finan-

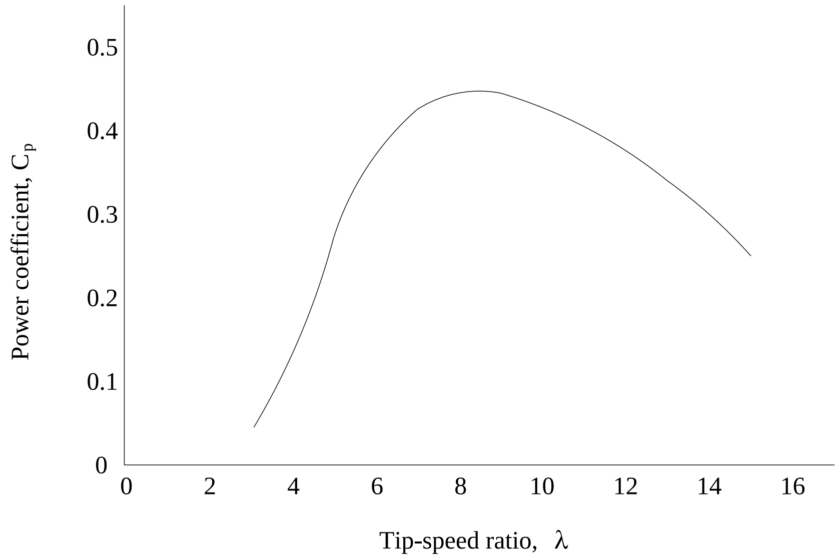

Figure 1.5: Typical curve of power coefficient vs. tip-speed ratio

cial income, is not proportional to the electrical energy generated by the windfarm. As a consequence, the trade-off between energy capture and fatigue damage must also change with time.

In (Lissaman et al., 1989), it is reported that some US utilities have variable tariffs which are related to electricity demand, and Summer on-peak electricity sales in California are over three times higher than during the off-peak period. In Britain, there is a wide variation in demand for electricity between maxima on cold mornings in Winter and minima in the early hours of the morning in Summer. Some specific tariffs for the South Western Electricity Board from April 1988† are presented in Table 1.1 (South Western Electricity Board, 1988). These tariffs indicate the wide differences in the value of generated power at various times of the day and of the year, with a maximum tariff that is 4.6 times larger than the minimum. This degree of variation implies that the trade-off between energy capture and fatigue damage

†Before the British electricity supply industry was privatised.

	Tariff	
	(pence)	normalised
For each unit supplied between 00.30 and 07.30 hours	1.70	1
For each unit supplied between 07.30 and 20.00 hours on Monday to Friday during:		
November and February	4.94	2.9
December and January	7.85	4.6
For each unit supplied at other times	2.27	1.3

Table 1.1: South Western Electricity Board tariffs from 1988

changes substantially as the tariff varies.

Maximising the performance of a windfarm also depends on minimising the installation and operation costs of the turbines in the windfarm. The power regulation carried out by a turbine when $v > v_{rated}$ cannot be perfect and hence some overrating of the components is necessary. To minimise installation and operation costs, the overrating of the turbines' components should be reduced as far as possible without compromising the components' reliability.

Large service factors** are necessary for gearboxes in wind turbines, as illustrated by Table 1.2. In general, large service factors necessitate more expensive gearboxes

Turbine manufacturer	Rated power of turbine (kW)	Service factor of gearbox*
Ecotéchnià	150	1.8†
HMZ–Windmaster	200	1.7‡
Nedwind	500	1.6§
Vestas–DWT	225	1.8§
Wind Energy Group	300	1.7¶

Table 1.2: Gearbox service factors of five commercial constant speed wind turbines

and are caused by long cumulative periods of operation at power levels above the rated power. The ex-works component costs for a typical medium-sized turbine, as fractions of the total ex-works cost, are presented in Table 1.3 (Bossanyi, 1992). It can be seen that the gearbox is the most expensive component in the turbine, so minimising the overrating of the gearbox is a prime target in the trade-off between financial income and turbine costs. The tower, blades and generator are the next

**The service factor of a gearbox is defined as q_d/q_∞. q_d is the design torque for $\geq n_\infty$ cycles equivalent to the variable duty, where n_∞ is the number of cycles at the lower 'knee' of the S–N curve. q_∞ is the endurance limit torque, that is, the torque that gives the stress at the lower 'knee' of the S–N curve. The design torque is a function of the cumulative time spent operating at torques above the endurance limit, and a calculation method is given in (BS 436:Part 3, 1986).

*All service factor calculations assume that the drive-train losses are given by $P_e = 0.95 \times P_{lss} - 0.04 \times P_{e(rated)}$, where P_e is electrical power and P_{lss} is low-speed shaft power.

† (Spruce, 1990)

‡ (De Wilde et al., 1988)

§ From observation at exhibition of the European Wind Energy Conference 1991

¶ (Haines, 1993)

Components, in order of decreasing cost		Fractional cost (%)
1	Gearbox	16
2	Tower	13
3	Blades (2)	12
4	Generator (2-speed)	10
	Other	49

Table 1.3: Ex-works component costs for a typical medium-sized wind turbine

most important components to consider. However, it is significant that half of the the capital cost of the turbine is spent on components other than the gearbox, tower, blades and generator, and therefore the other components should not be neglected when making the trade-off between financial income and fatigue damage.

The average wind-speed that a turbine experiences has a large effect on its energy capture (Justus, 1985) and fatigue damage (Madsen and Frandsen, 1984b). Also, higher turbulence intensities cause higher fatigue loads (Garrad and Hassan, 1986) for a negligible change in energy capture (Justus, 1985), as well as lower 'power quality'[††]. Windfarms, however, are usually characterised by average wind-speeds and turbulence intensities which vary significantly from turbine to turbine. There are two principal mechanisms that cause these variations and that therefore have a significant effect upon windfarm performance: firstly, any wake effects that are experienced by turbines that stand downstream of, and sufficiently close to, one or more other turbines; secondly, the topography (surface configuration) of the site

[††]Power quality is the variability of a windfarm's power output. A definition of what exactly is good power quality is not completely clear, but a simple measure is the standard deviation of power normalised by the average power level (Garrad and Hassan, 1988).

upon which the turbines stand.

A review of wake effects and the effects of operating turbines in clusters is presented in (Milborrow, 1988). Wind conditions within the wakes of wind turbines are characterised by lower average wind-speeds and higher turbulence intensities relative to those of the free stream, and these effects are compounded when the wakes of multiple turbines overlap (Milborrow, 1988). Wake effects lessen with the distance downstream of the turbine creating the wake as mixing takes place with the atmospheric boundary layer. Therefore attempts are made when designing windfarms to avoid powerful wake effects by the optimal design of the layout of the turbines with regard to the probability distribution of the wind-direction (Lissaman et al., 1989).

There are, however, several reasons why turbines in windfarms should be sited as close together as possible. One is the need to position all the turbines sufficiently far from all habitation so that the visual impact and noise experienced by local residents is negligible. Another is the desire to minimise the general visual impact of windfarms by siting the turbines within as small an area as possible. Additional reasons include the costs of building access roads and installing cables on windfarm sites. These economic factors alone give optimal inter-turbine spacings of between 4 and 8 times the turbines' rotor diameter (Föllings, 1989), which is small enough to cause significant wake effects. The combination of limitations on land use with economic limitations means that in any commercial windfarm there will be certain wind-directions for which the turbines must operate within the wakes of one of more other turbines. Therefore for these wind-directions there will be substantial variations from turbine to turbine in the rates of fatigue damage and energy capture, which is undesirable for maximal windfarm performance.

Many windfarms are situated in complex topography, particularly those in upland regions. For example, in California large numbers of windfarms are installed in the

Altamont Pass, where the topography is generally hilly, and in the Techachapi Pass, where the topography varies from hilly to mountainous. In Europe, large numbers of windfarms have been built both on lowlands in Denmark, Germany and Holland, and in hilly land in Denmark, Germany, Spain and the United Kingdom.

As a result it is important to consider the influence of topographic effects on the operational characteristics of windfarms. The environmental and cost factors described above that lead to turbines being sited relatively close together in windfarms have a significant effect on operation via topographic, or 'micro-siting', effects. From turbine to turbine across a windfarm, micro-siting effects can cause significant variations in the free stream average wind-speed (Nierenberg, 1989) and turbulence intensity (Antoniou, 1991). Therefore there are also substantial variations in the rates of fatigue damage and energy capture from turbine to turbine, which is again undesirable for maximal windfarm performance.

The electricity grid to which the turbines are connected can also have a significant effect on the operational characteristics of a windfarm. The combined dynamics of the windfarm plus grid must be considered if the windpower forms a significant fraction of the total generating capacity, for example on an island system. The combined dynamics must also be considered if a windfarm is connected to any grid via a weak link, as sometimes occurs, for example, in remote rural areas. However, in this thesis the most common situation of connection to an 'infinite' grid is considered. The grid is assumed to operate at constant frequency and to be unaffected by the operation of the windpower connected to it.

1.2 Options for the control of windfarms

There are two possible control strategies for windfarms: hierarchical control and multivariable control.

Hierarchical windfarm control consists of plant level control and supervisory level control. The plant level controllers carry out dynamic control to the set points specified by the supervisory level controller. The objectives of, and techniques used for, plant controllers in windfarms are briefly reviewed in Section 1.3. The supervisory controller assumes that the system is in steady-state, and makes strategic adjustments of the set points of the plant level controllers. The objectives of, and techniques used for, windfarm supervisory controllers are reviewed in Section 1.4.

Multivariable control involves controlling all the turbines dynamically from a central computer with inputs from all turbines and outputs to all turbines. Information about the wind-speed at upstream turbines or at upstream anemometers surrounding the windfarm is fed forward to the downstream turbine or turbines. However, the turbulent nature of the wind heavily reduces the longitudinal coherence of high-frequency fluctuations in wind-speed from point to point (Frost et al., 1978)[pp. 4.87–4.89]. In addition, the low frequency fluctuations in wind-speed, which are coherent from turbine to turbine, can be compensated for by the integral action that standard plant level controllers provide. Consequently there is some doubt surrounding the ability of multivariable control of windfarms to improve performance (Leithead, 1989), and very little research work has been reported in this field. One possible scheme is described in (Wefelmeier, 1993), with details in (Rossen, 1990)[ch. 6], in which a wind-speed signal from an upwind anemometer or turbine is fed forward to a variable-speed turbine sited downstream. The objective of the control strategy is the maximisation of energy capture. The simulation results presented in (Wefelmeier, 1993) illustrate the passage of a deterministic gust of

20 second duration from the upstream position to the downstream turbine, and the consequent anticipation of the gust by the downstream turbine. The gust of wind is assumed to be 100% coherent from the upstream site to the downstream site. In the example given, the energy capture increases by 3%, although the variance of the power increases significantly.

1.3 Plant control

At the plant level, there are two basic control tasks for constant-speed wind turbines. One is to accelerate the turbine up to the rated rotational speed so that it can be connected to the grid ('speed control'), and the other is to regulate the output power level once synchronised to the grid ('power control'). Speed control is of relatively minor importance, principally because the torques transmitted through the drive-train are negligible. It is the regulation of power that is of prime importance. For an overview of the power control options for wind turbines, see (Madsen, 1988), and for a thorough review of power control strategies, see (de la Salle et al., 1990). A survey of the control of large turbines is given in (Mattsson, 1984).

1.3.1 Objectives

Clear statements of the control objectives are not always made in the literature on the control of wind turbines. However a wide discussion of the role and objectives of wind turbine control is given in (Leithead et al., 1991). The most common objectives of the turbine control techniques that are investigated in the literature are:

1. maximisation of the energy capture
2. smoothing of the power generated
3. minimisation of the rotor, drive train and structural loads
4. shaping of the dynamics of the drive-train to achieve satisfactory damping and compliance

Maximising the energy capture maximises the financial income from the turbine. Smoothing the power reduces possible fluctuations in voltage and frequency on the grid to which the wind turbines are connected. Minimising the loads experienced by components within the turbine should maximise the fatigue lives of the components. Shaping of the drive-train dynamics extends the fatigue lives of the turbine's components by minimising excitation of the drive-train modes.

1.3.2 Techniques

A wide variety of methods have been proposed for the power control of wind turbines. The solutions that are most capable of meeting the required objectives differ somewhat according to configuration: constant or variable speed; number of blades; teetered or rigid hub; size of turbine. See (de la Salle et al., 1990) for a review.

A large number of pitch control algorithms have been investigated (de la Salle et al., 1990). Almost all are for the case where the blades act in unison, although one recent study has investigated the potential advantages of independent blade control, which gives additional degrees of freedom to the controller (Leithead et al., 1993). For variable-speed turbines, generator torque is used as an additional control variable.

A common objective is to operate the turbine at optimum tip-speed ratio until a turbine speed, torque or power constraint is reached. At higher wind-speeds the rotor power is limited by either pitch control or stall regulation (de la Salle et al., 1990).

1.4 Supervisory control

The techniques that have been applied for the supervisory control of windfarms can be divided into three basic types: start/stop control, set point control and on/off control. The objectives of supervisory windfarm control are reviewed in the next section, followed by the techniques that are used in Sections 1.4.2 to 1.4.5.

1.4.1 Objectives

The most common objectives of supervisory control techniques that are investigated in the literature are:

1. maximisation of the energy capture of the windfarm
2. smoothing of the power generated by the windfarm
3. minimisation of the rotor, drive train and structural loads of the turbines in the windfarm
4. maintenance of the local voltage levels near a windfarm within prescribed limits
5. carrying out basic operations such as reporting malfunctions and protecting against over-speeding

The basic operations are discussed in (Eggleston and Stoddard, 1987)[ch. 15] and are not considered further here.

1.4.2 Start/stop control

As the power generated by the turbines in a windfarm is combined before being exported to the grid, the operational strategy of any isolated turbine within the windfarm is of direct consequence to the total windfarm performance. Therefore all supervisory control actions, including the start/stop algorithms which run at the individual turbines, can be considered as windfarm supervisory control actions.

The objectives of all start/stop algorithms are the maximisation of energy capture and the minimisation of fatigue damage caused by synchronisation to, and de-synchronisation from, the grid. Early work is reported in (McNerney, 1980) and (Bossanyi, 1981). The former analyses the effect of cut-in/out algorithms at the low wind-speed end of the operating range, while the latter considers the effect of algorithms at both low and high wind-speeds. Both papers consider the performance of possible algorithms in terms of the energy capture and the number of start/stops. Two start/stop algorithms for low wind-speeds that are used on wind turbines in Californian windfarms are reviewed and analysed in (Rizzi and Auslander, 1990), and results are presented that show the effects of the various parameters in the algorithms on the energy output and the number of start/stops.

The above start/stop techniques use an average wind-speed or average power output to effectively predict future wind conditions, and the decision as to whether to remain in the current state or to change states uses this prediction. Further work by Bossanyi examines the potential for wind-speed prediction via the use of Kalman filters (Bossanyi, 1985), using the application to high wind-speed start/stops as an example. And work following on from that of (Rizzi and Auslander, 1990) investigates the application of predictive control to start/stop control at the low wind-speed end of the operating range (Hu et al., 1991).

1.4.3 Set point control of individual turbines

Several strategies have been investigated for set point control where the turbines are considered as isolated units. Early work described in (Bossanyi, 1979) examines the effect on energy capture of the de-rating of a turbine in high wind-speeds. For a linear reduction in power set point (P_{sp}) from $P_{sp} = P_{rated}$ at $v = 1.5\ v_{rated}$ to $P_{sp} = 0$ at $v = 3\ v_{rated}$, the effect on the annual energy output is found to be from small to negligible, depending on the shape of the annual wind-speed probability distribution. This indicates that there may be some potential for de-rating turbines in high wind-speeds.

Two set point control strategies have been investigated by the Wind Energy Group in field tests on 250 kW and 300 kW turbines (Smith, 1988). The objective of both schemes was the prevention of large power spikes in high wind-speeds for the purposes of reducing gearbox fatigue loads. The first method reduced the power set point as a linear function of wind-speed. For $v \geq 15$ m/s, P_{sp} was lowered by 10 kW/m/s, so that at that upper cut-out wind-speed of 25 m/s the turbines were de-rated to 150 kW ($= 60\% \times P_{rated}$). This scheme was found to be inappropriate, however, as it reduced the energy capture even in those high-speed winds which had low turbulence levels and therefore did not cause particularly damaging torque fluctuations. The second method reduced the power set point as a function of the average amount by which the power exceeded a threshold. This threshold could be specified to be either a fixed power level or the modified set point. While the former method is steady-state supervisory control, the latter is dynamic cascade control, because the dependence of the magnitude of the power excursions upon the modified set point introduces dynamics into the supervisory level. The gearbox fatigue damage in this pair of algorithms is effectively represented by the mean difference between the generated power and a threshold, in the first case a fixed threshold and in the second case a time-varying one. Preliminary tests carried out with the two algorithms encountered some stability problems.

In industry in general, it appears that set point control schemes are not common (Vestas-DWT, 1992). However, one industrial example that has been reported operates on 600 kW Westinghouse turbines in the Kahuku Point windfarm on Hawaii (Hausfeld et al., 1987). The control system limits high instantaneous power spikes (caused by large gusts) by reducing the power set point by 10% whenever a power spike exceeds 800 kW ($= 1.33 \times P_{rated}$). This leads to the turbines being operated at 10%, 20% or 30% reduced power for extended periods of time. The objective is the reduction of fatigue damage caused by high-power load spikes that are transmitted through the turbine. A further example of a de-rating scheme used in industry is on Windmaster turbines in the IJsselmeer-dike windfarm in the Netherlands (Landman, 1991). The turbines are pitch-controlled and have a rated power of 300 kW. The pitch control algorithm adjusts the angle of the blades if the generated power strays outside a deadband. For $v < v_{rated}$ the deadband is 310–330 kW, but for $v > v_{rated}$ it is lowered to 290–310 kW, thereby de-rating the turbines.

1.4.4 Set point control of windfarms

A variety of studies have investigated set point control where the turbines in a windfarm are considered not as isolated units but rather as a group. Several studies were carried out in the USA at the time of the rapid expansion in the Californian windfarms which investigated the potential dynamic impacts of significant penetrations of wind power on utility grids (Bouwmeester et al., 1985) (Schlueter et al., 1983b) (Schlueter et al., 1984) (Simburger and Cretcher, 1983). It is concluded that, if the wind-power penetration exceeds 5% of a utility's capacity, modified power system unit and control strategies are required. Further studies investigated power system control strategies, including windfarm control, for significant penetrations of wind-power (Javid et al., 1985) (Schlueter et al., 1983a).

In (Javid et al., 1985), a hierarchical control scheme is proposed which describes how to operate groups of windfarms as a function of their capacity, the load on the power system, and the speed with which the conventional generating units can increase or decrease their output. The hierarchy contains three levels: local wind turbine control, windfarm control, and control of groups of windfarms. The purpose of the windfarm control is to smooth fast power fluctuations, and it achieves this by specifying the power set points of the individual turbines. The control of groups of windfarms uses wind predictions to determine the individual windfarms' set points.

The propagation of frontal systems across large groups of windfarms causes the greatest problems for utility operators, and a wind prediction methodology for use in unit commitment and generation control strategies is further investigated in (Schlueter et al., 1986). The methodology is for the prediction of variations of wind-speed due to the passage of fronts one or more hours ahead, using a set of meteorological towers that encircle the windfarms at a radius of 100 miles.

On the smaller scale, a method is proposed for the control of windfarms of variable-speed turbines in (Steinbuch et al., 1988). The control objective is purely the maximisation of energy capture. Whenever any turbines operate in the wakes of other machines, the proposed windfarm controller would maximise the total energy capture by varying the tip-speed ratios of the interacting turbines. Physically, this technique makes upstream turbines extract less power from the wind so that downstream turbines receive a higher incident wind-speed as well as operating closer to the peak of their C_p–λ curves. The net effect can be an increase in total power output. Minor effects are reported from simulation of the turbines in the Dutch experimental windfarm at Sexbierum, but when other turbines were simulated slight increases in energy capture of up to 4% resulted. The control of power quality is also discussed in (Steinbuch et al., 1988), the conclusions being that windfarm control for power smoothing is hardly necessary for the turbines in the Sexbierum windfarm when $v > v_{rated}$ if optimal controllers are used for the individual turbines, but some form of windfarm control could be desirable when $v < v_{rated}$, as then the power fluctuations are greater.

Other investigations have also been carried out into the possibility of windfarm control for power smoothing at the Sexbierum windfarm (Toussaint et al., 1992). A potential problem is described that could arise from the connection of windfarms to weak grids where fluctuations on the local low voltage network can be caused, and which can lead to the need for expensive strengthening of the grid connection. To avoid additional costs, windfarm control is identified as a possible alternative. It is pointed out that reductions in wind-generated income will result but that these reductions may be worth the savings in grid connection costs. In the proposed algorithm, the current generated by the windfarm is fed back to the central control computer which calculates the difference between the windfarm current set point and the actual current. This difference is fed through a PI (Proportional Integral) element and the result is divided by the number of turbines to generate the individual turbines' current set points. The feedback is required because some turbines may not be able to achieve their required current level at times. In such circumstances the windfarm controller increases the set points of those other turbines which can supply the necessary power. The controller knows which turbines are capable of supplying extra power by monitoring the pitch angles of their blades. Initial experimental results illustrate how a reduction in windfarm power set point significantly reduces the power fluctuations.

A study has been carried out which examines the possibilities of reducing windfarm power fluctuations and maximising windfarm energy capture by controlling the electrical torque set points of variable-speed turbines. The models and some results are reviewed in (Wefelmeier, 1993) and details appear in (Rossen, 1990). A windfarm of 3 variable-speed turbines and 6 constant-speed turbines in a 3-by-3 layout is simulated. A mixture of variable-speed and constant-speed turbines is specifically chosen because the variable-speed turbines, although more expensive to construct, may be used to store energy in their rotors in the short term and hence assist in smoothing the windfarm's power output. The windfarm controller's input is the power generated by the farm and its outputs are the electrical torque set points of the variable-speed turbines. The constant-speed machines are not controlled by the windfarm controller. In low wind-speeds, the simulation results indicate reduced power variations of 40% as well as slightly higher energy capture.

The only practical experience with a windfarm controller on a commercial windfarm that would appear to have been reported is in (Mutone et al., 1986). The site is the windfarm of Westinghouse turbines at Kahuku Point on Hawaii. This windfarm is reported as having a sufficiently weak connection to the grid that it was expected that significant perturbations of the local grid voltage would be caused as load conditions varied. In addition, the reactive power usage at the windfarm is controlled and subject to fees. The windfarm control objectives are clearly stated: firstly, the power generated by the windfarm is to be maximised; secondly, the local voltage levels are to be maintained within prescribed limits. The turbines have synchronous generators and the set point for the windfarm is chosen from the following: voltage; reactive power; power factor. Either open loop or closed loop control is then selected. Under open loop control, the windfarm controller transmits a set point to each turbine and does not monitor or adjust the set points thereafter. Under closed loop control, the windfarm controller transmits set points based upon the performance of both the individual turbines and the windfarm as a whole. The turbines are then periodically monitored and new set points are sent out on the basis of the current performance. It is concluded in (Mutone et al., 1986) that expectations are high for trouble-free integration of the windfarm into the utility system.

1.4.5 On/off control

"On/off control" is used here to mean the strategic control of windfarms whereby individual turbines are taken off-line under certain conditions to improve the overall performance of the windfarm. For turbines with variable set points, on/off control is a sub-set of windfarm set point control. The strategies that have been implemented are generally qualitative and relatively unsophisticated, using open-loop algorithms which are functions of the wind-direction only.

One example of on/off control employed in practice is that used by the world's largest windfarm operator, US Windpower (Pacific Gas & Electric Co./US Windpower, 1992) (Smith, 1993). In the Altamont Pass in California, the winds are usually either south-west to north-east or vice versa, and consequently the turbines are normally arranged in rows running north-west to south-east, with large gaps between each row but small inter-turbine spacings along the rows. The US Windpower windfarms have inter-turbine spacings along the rows of 1.3 rotor diameters (D) and turbines of the '56-100' design. Figure 1.1 illustrates the turbines and the layout. The turbines are constant speed, pitch-controlled machines, with a rotor which operates downwind of the tower. The control strategy consists of switching off either 2 in every 3 turbines or 3 in every 4 when the wind blows along a row, and it applies to all US Windpower's 3500 turbines in Altamont Pass, representing an installed capacity of 350 MW. This strategy reduces the powerful fatigue loads which would otherwise be caused in long lines of turbines with 1.3 D inter-turbine spacings by the strong wake-induced turbulence.

A commercial European windfarm where on/off control is in use is the IJsselmeerdike windfarm in the Netherlands (Sisouw de Zilwa and Bakema, 1991). There are 50 constant speed, pitch-controlled turbines in this windfarm, in a single row running due north–south with 5D inter-turbine spacings. Although the wind blows north–south or south–north ($\pm 22.5°$) for a relatively large fraction (30%) of the time, various geographic factors make this the optimal layout (Paes and Nauta, 1988). The control strategy consists of switching off 1 turbine in every 2 when the wind blows north–south or south–north ($\pm 22.5°$), and it runs automatically on a central computer (Landman, 1991). The objectives of the control strategy are the maximisation of energy capture and the prevention of fatigue damage (Paes and Nauta, 1988) (Landman, 1991). Off-line calculations of the energy capture for winds blowing north–south or south–north have been carried out, and it has been found that the average power output obtained by running the windfarm with all the turbines switched on is lower than with every second turbine switched off (Landman, 1991). The objective of the prevention of fatigue damage is based on a specific hypothesised case which can occur in low wind-speeds at a turbine operating in the wakes of other turbines. Under certain unusual wind conditions, with a high velocity deficit and large variations in the wind-direction in the combination of wakes that the downstream turbine experiences, the downstream turbine could gradually yaw its rotor around until it was operating with its rotor downwind of its tower. (Landman, 1991). The operation of an upwind turbine in the downwind configuration would cause unforeseen and potentially very damaging loadings. However, the supervisory control strategy effectively doubles the wake lengths and therefore lessens the wake effects significantly, so the probability of turbines yawing around to face downwind is strongly reduced.

Another example of a European windfarm that utilises on/off control is the Zeebrugge windfarm in Belgium. This windfarm consists of 23 constant-speed, pitch-controlled turbines in three groups (De Wilde et al., 1988). Wake effects are possible along rows of turbines in various directions and a study has found that, with the wind blowing along a row of turbines separated by approximately 6D, power losses of 45% are experienced at the second and subsequent turbines (Van Leuven et al., 1989). The supervisory control strategy runs on a central computer and under certain conditions it will switch off up to three turbines (Dekens, 1990). The sole objective of the control strategy is the maximisation of energy capture.

Some research work in this field is reported in (Föllings and Pfeiffer, 1988). For turbines standing in a line configuration with uniform inter-turbine spacings, the effect on windfarm energy capture of switching off 1 in every 2 or 2 in every 3 turbines is investigated using a computer model. The results indicate that the total energy capture is increased by both types of scheme, for any inter-turbine spacings between 3D and 8D.

1.5 Aims of the work

The overall aim of the work is the investigation, development and testing of windfarm supervisory control algorithms, where the control objectives are the maximisation of the financial income from the electricity generated and the minimisation of the fatigue damage sustained by the turbines. This aim is sub-divided into the following areas of work:

1. develop a computer simulation of a windfarm that is suitable for the purposes of testing supervisory control algorithms
2. develop a suitable cost function for use in the windfarm controller
3. develop suitable models of the wind turbine fatigue damage
4. examine the feasibility of an on-line version of the controller

1.6 Structure of the thesis

The general structure of the thesis is listed below. A more detailed introduction precedes the individual chapters.

Chapter 2 describes in detail the development of steps in a methodology for the development of windfarm supervisory control strategies. The concept of 'ideal' power curves is introduced, and a method for their generation using 'ideal' lifetime performance functions is presented whereby ensemble averaging using long-term wind-speed probability distributions is employed. Simple cost functions are then developed which are based upon time-averaging and are therefore suitable for practical use in an on-line windfarm controller, and a method for the evaluation of the cost function's weighting parameters is described.

Chapter 3 begins with a description of the specific windfarm characteristics that are used in the test-bed simulation of a small (3-turbine) windfarm. This is followed by a detailed description of the wind models that are employed and a wake model that was developed specifically for this application. Results are then presented which illustrate the satisfactory operation of the turbine and wake models, including the convergence of estimates made using the simulation.

Chapter 4 briefly describes the methods that are used for estimating turbine fatigue damage rates and then presents a broad range of fatigue damage results. The results are for widely varying wind conditions and for both single turbines and the 3-turbine windfarm. Particular attention is paid to the variation of fatigue damage with power set point.

Chapter 5 presents all the results from the application of the methodology for the development of windfarm supervisory control strategies that is described in Chapter 2, using the simulation described in Chapter 3 and the fatigue damage models presented in Chapter 4. Following the results of the application of the methodology, results are presented for a simple on-line controller which illustrate the feasibility of the on-line implementation of the control strategy.

Chapter 6 draws conclusions from the results presented throughout the thesis and presents suggestions for further work.

Appendix A provides a list of the principal symbols used in equations throughout the text. Appendix B describes the wind turbine model and Appendix C gives the values of some parameters used in the model. Appendix D describes the application of confidence intervals to, and a procedure for, the determination of the convergence of estimated means. Appendix E presents an overview of the genetic search algorithm that is used to determine the 'ideal' power curves. Appendix F lists the simple extremum control algorithm that is used for on-line control.

Chapter 2

Development of windfarm supervisory control strategies

2.1 Introduction

The chapter begins with a description of the windfarm controller structure that is employed, in Section 2.2. Thereafter, the steps in a methodology for the development of windfarm supervisory control algorithms are explained. These steps are as follows:

1. Produce 'ideal' power curves for an individual turbine from an 'ideal' lifetime performance function (Section 2.3). The lifetime financial income and fatigue damage are estimated via the use of both ensemble averaging and time averaging.

2. Produce a cost function for isolated, non-interacting turbines within a windfarm (Section 2.4). The 'ideal' lifetime performance function is not suitable for use in an on-line controller, so simple cost functions are described that can give identical power curves to the 'ideal' curves. The cost functions use time averaging only and no ensemble averaging.

3. Produce a cost function for interacting turbines within a windfarm (Section 2.5). The cost function for non-interacting turbines is extended to the case of interacting turbines.

4. Develop an on-line version of the control strategy. The principal options for on-line implementation are described in Section 2.6, together with the basic operational characteristics that a steady-state windfarm supervisory controller must possess.

2.2 Structure of windfarm supervisory controller

To achieve the windfarm control objectives of maximising the financial income from the electricity generated while minimising the fatigue damage of the turbines, the turbines are split into two separate groups for control purposes:

1. **Control of non-interacting turbines** These are turbines which, dynamically, are neither in the near or intermediate wakes of any other turbines nor are creating a wake in which other turbines are standing. These turbines therefore have negligible interaction with any other turbines in the windfarm and are dynamically isolated*. There can be no performance benefit from controlling these turbines as anything other than individual machines under single-input single-output supervisory control. If controlled as part of a group of turbines, no higher performance can be attained but the increased dimensionality of the problem makes the control problem more complex and thereby reduces the controller's performance. Therefore non-interacting turbines should be controlled as isolated units, with the same cost function and weighting parameter

*Note that this is not the case if the control of power quality is an objective of the windfarm control, and the turbines should then not be divided up into non-interacting and interacting groups.

for each turbine.

2. **Control of interacting turbines** These are turbines which, dynamically, are either in the near or intermediate wakes of one or more other turbines or are creating a wake in which one or more other turbines are standing. Improvements in performance can be gained by controlling these turbines as groups of interacting machines under multi-input multi-output supervisory control. If there are groups of turbines dynamically linked by wakes within a windfarm, and these groups are separate from one another, then each group should be controlled separately. As in the case of non-interacting turbines, the overall control of dynamically separate units cannot be of any advantage and leads to reduced controller performance.

As the wind-direction incident upon a windfarm slowly changes, the turbines within the farm will move from the 'non-interacting' classification into the 'interacting' classification, and vice versa. Therefore the windfarm controller must include a coordinating algorithm which takes a running average of the wind-direction, classifies the turbines as either interacting or non-interacting and, when necessary, oversees the smooth transition from individual turbine control to group control or vice versa. Figure 2.1 illustrates the resulting structure of the windfarm controller, with the coordination carried out in the blocks marked "decision logic".

2.3 'Ideal' power curves

This section examines how a turbine would be ideally operated, assuming that the shape of the power curve can be varied but that upper and lower wind-speed limits remain.

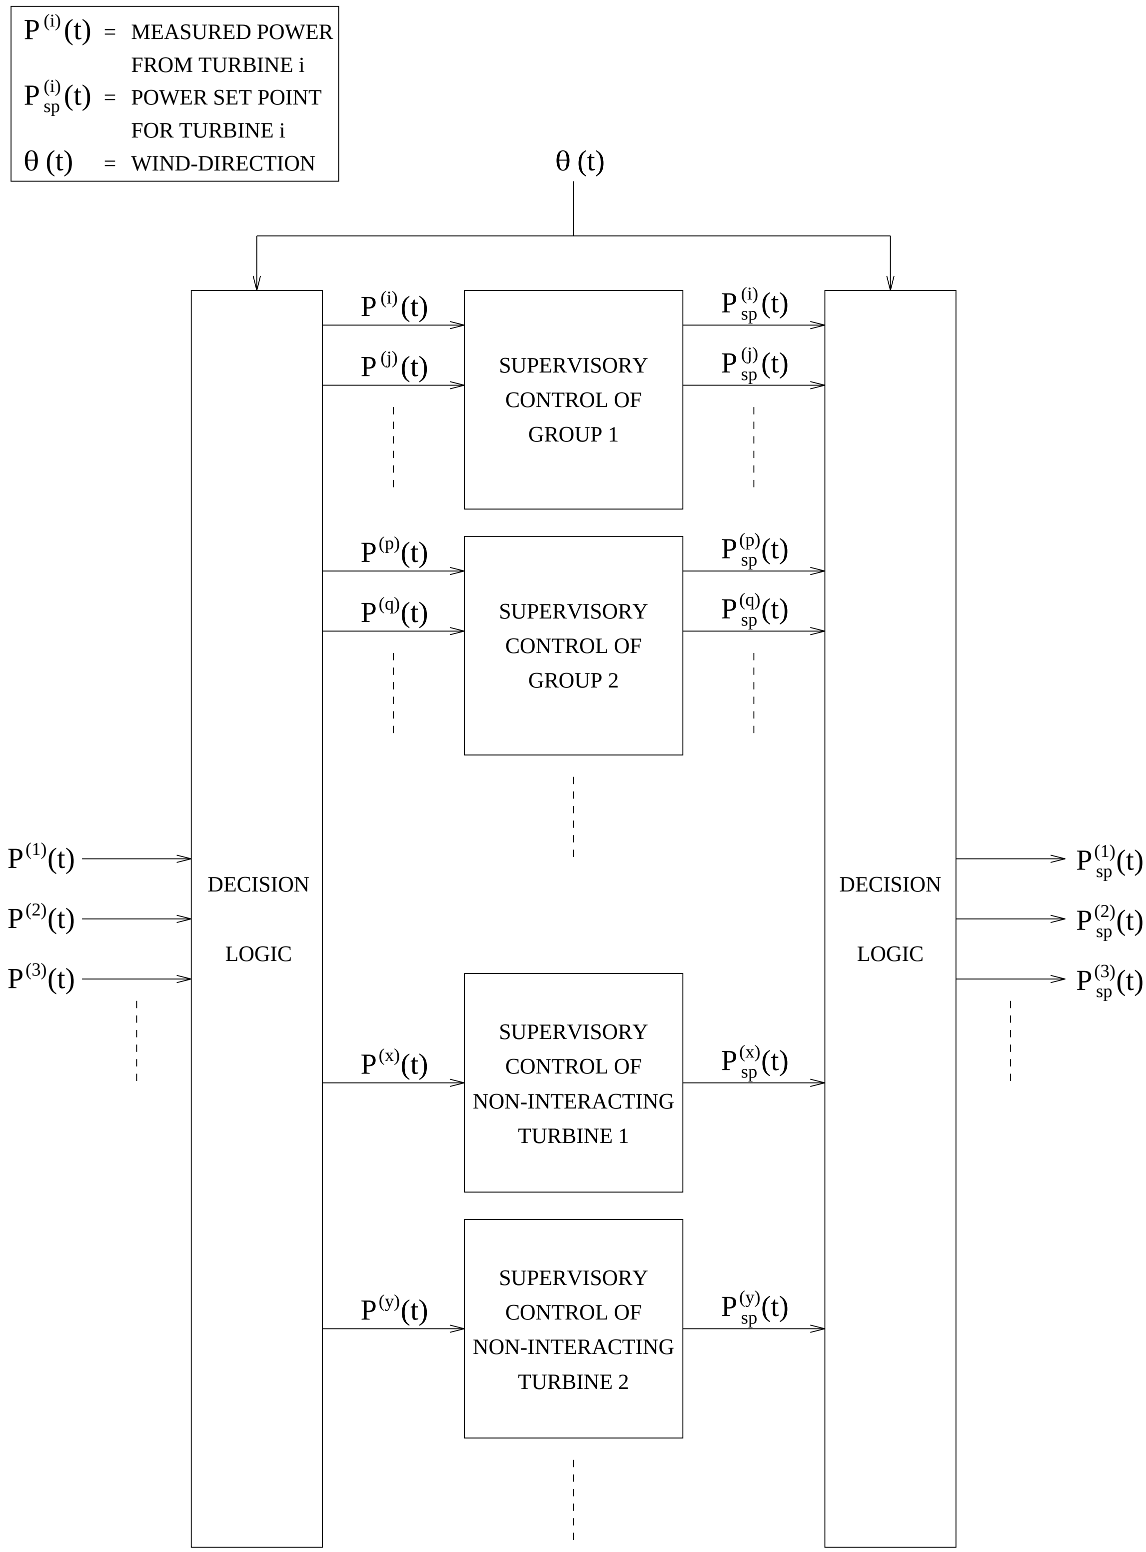

Figure 2.1: Coordination of windfarm control

2.3.1 'Ideal' lifetime performance function

The development of ideal power curves involves the maximisation of a lifetime performance function. A complete lifetime performance function must contain elements to represent:

1. income from generated electricity
2. capital and replacement costs of the components in the turbine
3. expenditure on maintenance

The second and third items involve functions that estimate fatigue damage over the turbine's lifetime and then compute the minimal lifetime cost required to purchase and maintain the necessary components. The capital, replacement, and maintenance costs of the components can be minimised by sizing the components correctly, by evaluating whether it would be cheaper to replace the components one or more times during the lifetime of the wind turbine, and by optimising the turbine's operating strategy. The operating strategy further affects the income from the generated electricity, and therefore all elements in the performance function are interlinked. This gives a complex function which encompasses both turbine design and operation.

The complexity of the performance function is reduced by decoupling the design and operation issues. Fatigue damage is included as a constraint upon the lifetime profit that can be achieved, rather than as a term within the performance function. As a result the ideal power curves are optimised for existing designs of components used in the turbine, and the resulting supervisory control strategies then give either:

- improved reliability of an existing design of turbine

  or

- the ability to operate a particular design of turbine on a site with wind conditions outside the turbine's design envelope

The lifetime performance function is then to:

1. maximise the financial income from the electricity generated by the turbine
2. constrain the turbine's fatigue damage to be within a specified amount

The rate of financial income ($\dot{m}(t)$) and the rate of accumulation of fatigue damage ($\dot{d}(t)$) are non-stationary processes, and so a complete ideal performance function would evaluate the total financial income and fatigue damage by integrating $\dot{m}(t)$ and $\dot{d}(t)$ over the lifetime of the turbine. However, as the design life of a typical wind turbine is 20 years, the application of a complete ideal performance function is infeasible, and an alternative approach is necessary.

In general it is recommended that non-stationary processes should be divided into piecewise stationary segments, if practicable, to simplify analysis (Bendat and Piersol, 1986)[p. 425]. In the case of wind energy such a division is indeed possible, facilitated by the 'spectral gap' in the spectrum of wind-speed (see Figure 1.2 on page 4). Consequently time-averages of the fluctuations in the rate of energy capture ($\bar{\dot{e}}$) and fatigue damage ($\bar{\dot{d}}$) that are caused by atmospheric turbulence can be evaluated with low estimation errors.

As tariffs for generated electricity can vary with time of day, season and year, the lifetime of the turbine is sub-divided into multiple periods in each day, multiple seasons in each year, and multiple years within the lifetime. Ensemble averages of the rate of income and fatigue damage for period number i_p in season i_s in year i_y are then calculated, using values of $\overline{\dot{e}}$ and $\overline{\dot{d}}$ for wind-speeds across the turbine's operating range together with a discrete long-term wind-speed probability distribution $p_{\{i_s,i_p\}}(\overline{v}_i)$:

$$\hat{\overline{\dot{m}}}_p(i_y,i_s,i_p) = \sum_{i=1}^{i_{hi}} \left\{ p_{\{i_s,i_p\}}(\overline{v}_i) \,.\, K_m(i_y,i_s,i_p) \,.\, \overline{\dot{e}}\left[\overline{v}_i, P_{sp}(i_y,i_s,i_p,\overline{v}_i)\right] \right\} \quad (2.1)$$

$$\hat{\overline{\dot{d}}}_p(i_y,i_s,i_p) = \sum_{i=1}^{i_{hi}} \left\{ p_{\{i_s,i_p\}}(\overline{v}_i) \,.\, \overline{\dot{d}}\left[\overline{v}_i, P_{sp}(i_y,i_s,i_p,\overline{v}_i)\right] \right\} \quad (2.2)$$

$\hat{\overline{\dot{m}}}_p$ = ensemble average of financial income for period i_p

i_y = year number

i_s = season number

i_p = period number

i = wind-speed interval number

i_{hi} = wind-speed interval number for upper cut-out wind-speed

$\overline{v}_i$ = average wind-speed for interval i

$p_{\{i_s,i_p\}}(\overline{v}_i)$ = probability distribution of wind-speed $\overline{v}_i$ for period i_p in season i_s

K_m = tariff for electricity generated

$\overline{\dot{e}}$ = short-term time average of rate of energy capture

P_{sp} = power set point

$\hat{\overline{\dot{d}}}_p$ = ensemble average of fatigue damage for period i_p

$\overline{\dot{d}}$ = short-term time average of fatigue damage

For period i_p, the total estimated income $\hat{\overline{\dot{m}}}_p(i_y,i_s,i_p)$ and fatigue damage $\hat{\overline{\dot{d}}}_p(i_y,i_s,i_p)$ are:

$$\hat{\overline{m}}_p(i_y, i_s, i_p) \;=\; T_p(i_p) \,.\, \hat{\overline{\dot{m}}}_p(i_y, i_s, i_p) \tag{2.3}$$

$$\hat{\overline{d}}_p(i_y, i_s, i_p) \;=\; T_p(i_p) \,.\, \hat{\overline{\dot{d}}}_p(i_y, i_s, i_p) \tag{2.4}$$

where

$$\sum_{i_p=1}^{n_p} T_p(i_p) = 24 \text{ hours}$$

$T_p(i_p)$ = length of period i_p (hours)

n_p = number of periods in day

The estimated lifetime income $\hat{m}_{life}$ and damage $\hat{d}_{life}$ are then evaluated by summing $\hat{\overline{m}}_p(i_y, i_s, i_p)$ and $\hat{\overline{d}}_p(i_y, i_s, i_p)$ over the lifetime of the turbine, that is, over all periods i_p in all seasons i_s in all years i_y:

$$\hat{m}_{life} \;=\; \sum_{i_y=1}^{n_y} \sum_{i_s=1}^{n_s} n_d(i_s) \sum_{i_p=1}^{n_p} \hat{\overline{m}}_p(i_y, i_s, i_p) \tag{2.5}$$

$$\hat{d}_{life} \;=\; \sum_{i_y=1}^{n_y} \sum_{i_s=1}^{n_s} n_d(i_s) \sum_{i_p=1}^{n_p} \hat{\overline{d}}_p(i_y, i_s, i_p) \tag{2.6}$$

where

$$\sum_{i_s=1}^{n_s} n_d(i_s) = 365 \text{ days}$$

n_y = number of years in design lifetime of turbine

n_s = number of seasons in year

$n_d(i_s)$ = number of days in season i_s

For the case where daily and seasonal changes are not included, that is, $n_p = n_s = 1$, the above approach is commonly applied for the purposes of estimating annual energy capture (Eggleston and Stoddard, 1987)[pp. 67*ff*] (Justus, 1985). It has also been proposed or applied for the purposes of estimating annual fatigue damage (Madsen and Frandsen, 1984b) (Thomsen et al., 1992) (WEG 33-4115, 1991) (Eggleston and Stoddard, 1987)[pp. 311*ff*].

The complete 'ideal' performance function is formed by combining Equations 2.1, 2.3 and 2.5 to give the lifetime income and by combining Equations 2.2, 2.4 and 2.6 to give the lifetime fatigue damage. The complete function is summarised in Equations 2.7–2.9. In Equation 2.7 the lifetime financial income is maximised, and in Equations 2.8 and 2.9 constraints of the form $A.\hat{d}_{max}$ are applied to the lifetime fatigue damage of components within the turbine. $\hat{d}_{max}$ is the maximum estimated fatigue damage that a component can experience, and is equal to the fatigue damage that a component experiences for no set point de-rating (current industrial practice). The parameter A is the fatigue damage safety factor that is applied to $\hat{d}_{max}$. For example, to constrain the gearbox's lifetime fatigue damage to 80% times the damage experienced when P_{sp} is equal to rated power, A_{gear} is set to 0.8.

'Ideal' lifetime performance function

$$\hat{m}_{life} =$$

$$\max_{P_{sp}(i_y,i_s,i_p,\overline{v}_i)} \sum_{i_y=1}^{n_y} \sum_{i_s=1}^{n_s} n_d(i_s) \underbrace{\sum_{i_p=1}^{n_p} T_p(i_p) \underbrace{\sum_{i=1}^{i_{hi}} \{ \underbrace{p_{\{i_s,i_p\}}(\overline{v}_i)}_{\text{probability}} \, . \, \underbrace{K_m(i_y,i_s,i_p)}_{\text{tariff(p/kWh)}} \, . \, \underbrace{\overline{\overline{e}}\,[\overline{v}_i, P_{sp}(i_y,i_s,i_p,\overline{v}_i)]}_{\text{average power (kW)}} \}}_{\text{income in the period } i_p \ (\pounds)}}_{\text{income in one day } (\pounds)} \tag{2.7}$$

such that

$$\sum_{i_y=1}^{n_y} \sum_{i_s=1}^{n_s} n_d(i_s) \sum_{i_p=1}^{n_p} T_p(i_p) \sum_{i=1}^{i_{hi}} \left\{ p_{\{i_s,i_p\}}(\overline{v}_i) . \, \overline{\dot{d}}_{gear}\,[\overline{v}_i, P_{sp}(i_y,i_s,i_p,\overline{v}_i)] \right\} \;\; \leq \;\; A_{gear}\hat{d}_{gear_{max}} \tag{2.8}$$

$$\sum_{i_y=1}^{n_y} \sum_{i_s=1}^{n_s} n_d(i_s) \sum_{i_p=1}^{n_p} T_p(i_p) \sum_{i=1}^{i_{hi}} \left\{ p_{\{i_s,i_p\}}(\overline{v}_i) . \, \overline{\dot{d}}_{blades}\,[\overline{v}_i, P_{sp}(i_y,i_s,i_p,\overline{v}_i)] \right\} \;\; \leq \;\; A_{blades}\hat{d}_{blades_{max}} \tag{2.9}$$

$$\vdots$$

A = fatigue damage safety factor ($0 \leq A \leq 1$)

Subscripts:

max = maximum fatigue damage that component can experience

= fatigue damage for no de-rating of P_{sp}

$gear$ = gearbox

2.3.2 Simplified 'ideal' performance function for two tariff seasons per annum

For practical implementation, it is assumed that the electricity tariffs and the value of money are both constant from year to year. Hourly variations in tariff are neglected, and the following seasons are used:

- Winter — December to February
- Rest Of Year — March to November

Details of the probability distributions that are used for each season are given in section 3.4.2. To calculate the tariffs for the two seasons, averages are taken of the full South Western Electricity Board tariffs (see Table 1.1 on page 10). The results are presented in Table 2.1. The normalised values show that, with the diurnal

	Winter	Rest Of Year
Season number i_s	1	2
Tariff $K_m(i_s)$ (1988 prices)	3.82 p/kWh	2.22 p/kWh
Normalised tariff $\frac{K_m(i_s)}{K_m(2)}$	1.72	1

Table 2.1: Average electricity tariffs for Winter and Rest Of Year

variations in tariff removed and with only two seasons per annum, the ratio of tariffs between Winter and the Rest Of Year remains relatively large.

The average rate of energy capture and average rate of accumulation of fatigue damage are both functions of the turbulence intensity I $(= \sigma_v/\overline{v})$ as well as the average wind-speed $\overline{v}$. However, although annual probability distributions are available for $\overline{v}$, the same is not the case for I. As a consequence, the 'ideal' performance function is limited to the region $\overline{v} \geq 8$ m/s (for $\overline{v}$ at 25 m above ground level), because then near-neutral meteorological conditions predominate (Pal Arya, 1988)[p. 57] (Justus, 1985), and under near-neutral conditions the turbulence intensity is typically in the range 10% to 20% (Pal Arya, 1988)[p. 118]. It is then assumed that I is constant and equal to 15%.

One fatigue damage constraint only is applied, for the most expensive component in the turbine, the gearbox. The resulting performance function, that is used in the searches for ideal power curves, is then:

Simplified 'ideal' performance function for two seasons per annum

$$\hat{m}_{year} = \max_{P_{sp}(i_s,\overline{v}_i)} \underbrace{\sum_{i_s=1}^{2} K_m(i_s)\, \underbrace{T_s(i_s) \sum_{i=1}^{i_{hi}} \{ \underbrace{p_{\{i_s\}}(\overline{v}_i)}_{\text{probability}} \;.\; \underbrace{\hat{\overline{\dot{e}}}\,[\overline{v}_i, P_{sp}(i_s,\overline{v}_i)]}_{\text{average power (kW)}} \}}_{\text{energy capture in season } i_s \text{ (kWh)}}}_{\text{income in season } i_s\ (\pounds)} \qquad (2.10)$$

such that

$$\sum_{i_s=1}^{2} T_s(i_s) \sum_{i=1}^{i_{hi}} \left\{ p_{\{i_s\}}(\overline{v}_i) \,.\, \widehat{\overline{\dot{d}_{gear}}}\,[\overline{v}_i, P_{sp}(i_s,\overline{v}_i)] \right\} \leq \frac{1}{20} A_{gear} \hat{d}_{gear_{max}} \qquad (2.11)$$

and

$$I = 15\%$$

$\hat{m}_{year}$ = estimated financial income in one year of operation

i_s = 1 for Winter; 2 for Rest Of Year

$T_s(i_s)$ = length of season i_s (hours)

$\hat{\overline{\dot{e}}}$ = estimated mean rate of energy capture (time-average)

$\widehat{\overline{\dot{d}_{gear}}}$ = estimated mean rate of gearbox fatigue damage (time-average)

The fraction $\frac{1}{20}$ in Equation 2.11 represents the scaling down of the function from over the design lifetime of the turbine, 20 years, to over the period of 1 year.

The estimates $\hat{\dot{e}}$ and $\widehat{\overline{\dot{d}_{gear}}}$ are evaluated by carrying out time-averaging in simulation, for all combinations of mean wind-speeds and power set points at discrete intervals across the ranges 8 m/s $\leq \overline{v} \leq v_{hi}$ and $0 \leq P_{sp} \leq P_{rated}$ (P_{rated} = rated power of the turbine). Having obtained this data across the complete operating region, the performance of any possible Winter/Rest Of Year pair of power curves can be easily calculated.

2.3.3 Simplified 'ideal' performance function for one tariff season per annum

'Ideal' power curves are also found for the case of one tariff season per annum, partly for purposes of comparison with the two-tariff results, and partly because some utilities operate with a constant tariff. An example of the latter is the current UK situation for renewable energy power stations, which receive subsidies at a flat rate regardless of the demand on the network.

The assumptions that are made in the previous section are all valid in this case as well, and the tariff is again calculated by averaging the South Western Electricity Board tariffs, giving 2.62 p/kWh. The performance function in this case is:

Simplified 'ideal' performance function for one season per annum

$$\hat{m}_{year} = \max_{P_{sp}(\overline{v}_i)} K_m T_y \underbrace{\sum_{i=1}^{i_{hi}} \{ \underbrace{p(\overline{v}_i)}_{\text{probability}} . \underbrace{\hat{\dot{e}}[\overline{v}_i, P_{sp}(\overline{v}_i)]}_{\text{average power (kW)}} \}}_{\text{energy capture per annum (kWh)}} \tag{2.12}$$

such that

$$T_y \sum_{i=1}^{i_{hi}} \left\{ p(\overline{v}_i) . \widehat{\overline{\dot{d}_{gear}}} [\overline{v}_i, P_{sp}(\overline{v}_i)] \right\} \leq \frac{1}{20} A_{gear} \hat{d}_{gear_{max}} \tag{2.13}$$

and

$$I = 15\%$$

T_y = length of year (hours)

The estimates $\hat{\dot{e}}$ and $\widehat{\overline{\dot{d}_{gear}}}$ that are evaluated via simulation for use in the case of two tariff seasons per annum are also used for one season per annum.

2.4 Simple cost function for non-interacting turbines

A simple cost function is described for those turbines within a windfarm that, at any particular time, are non-interacting, and that would therefore be under single-input single-output supervisory control.

2.4.1 Formulation of the cost function

An on-line controller cannot use the 'ideal' performance functions, because these functions require the estimates $\hat{\overline{\dot{e}}}$ and $\hat{\overline{\dot{d}}}$ to be obtained at average wind-speeds across the whole operating range from lower cut-in to upper cut-out. The controller must calculate the new set point at any given time-step on the basis of estimates $\hat{\overline{\dot{e}}}$ and $\hat{\overline{\dot{d}}}$ at only one value of $\overline{v}$. Therefore only time-averaging is required in the cost function and the ensemble averaging of the lifetime performance function is not involved.

For ease of practical application, a cost function is required that only needs a measurement of electrical power. Also, to ease both analysis and on-site tuning, the function must have a minimal number of weighting parameters. However, when used in an optimising controller, the power set point curves that result from the use of the cost function should approximate to the 'ideal' power curves.

In the two-tariff case, the cost function that is used requires just one weighting parameter for each tariff season into which the year is split, and hence the function is different in the Winter season compared to in the Rest Of Year season. The cost functions for two and one tariffs per annum are presented in Equations 2.14 and 2.15 respectively.

Simple cost function for a non-interacting turbine and two tariffs p.a.

$$J(i_s) = \int_{t-T_c}^{t} \left[\lambda_w(i_s)\,.\,\dot{d}(t) - \dot{m}(t)\right] dt \tag{2.14}$$

Simple cost function for a non-interacting turbine and one tariff p.a.

$$J = \int_{t-T_c}^{t} \left[\lambda_w\, \dot{d}(t) - \dot{m}(t)\right] dt \tag{2.15}$$

J = cost to be minimised

i_s = season number

T_c = controller time-step

λ_w = weighting parameter or parameters

$\dot{d}(t)$ = rate of accumulation of fatigue damage

$\dot{m}(t)$ = rate of income

= $K_m(i_s)\,\dot{e}(t)$ for two tariff seasons per annum

= $K_{m_{year}}\dot{e}(t)$ for one tariff season per annum

The effect of the weighting parameter λ_w is as follows. Larger values of λ_w increase the weight of the fatigue damage term within the function relative to the income term and, in a controller, will drive down the power set point during periods when the levels of fatigue damage are high relative to the levels of income. Smaller values of λ_w, conversely, decrease the weight of the fatigue damage term within the function relative to the income term. The same trade-off between financial income and fatigue damage can be achieved at all non-interacting turbines within a windfarm by using the same value of weighting parameter throughout.

2.4.2 Evaluation of weighting parameters

A method for the evaluation of the cost function's weighting parameters is described in this section, for the more general case of two, rather than one, tariff seasons per annum.

The weighting parameters are evaluated using the same optimisation criteria as in the case of the 'ideal' power curves. The lifetime financial income is maximised such that the lifetime fatigue damage is constrained to be within a specified fraction of the maximum possible damage. The estimates $\hat{\dot{e}}$ and $\widehat{\dot{d}_{gear}}$ that are evaluated via simulation are again used to calculate the lifetime income and fatigue damage.

Algebraically, the optimisation is the same as in equations 2.10 and 2.11 (or 2.12 and 2.13 in the 1-tariff case), except that in this case the power set point function $P_{sp}(i_s, \overline{v})$ is not itself the independent function. Instead, $P_{sp}(i_s, \overline{v})$ is generated from the new independent variables, the weighting parameters $\lambda_{w_{win}}$ ($= \lambda_w$ in Winter) and $\lambda_{w_{roy}}$ ($= \lambda_w$ in Rest Of Year):

$$P_{sp}(i_s, \overline{v}) = f(\lambda_{w_{win}}, \lambda_{w_{roy}})$$

The objective of the evaluation of the weighting parameters is to achieve a level of performance by the power curves $P_{sp}(i_s, \overline{v})$ that is similar to that of the 'ideal' power curves. However, the number of available independent variables is reduced substantially in the weighting parameter optimisations. For example, if the operational range of wind-speeds is discretised into 20 intervals in the 'ideal' optimisations, the number of independent variables falls from 40 to 2.

To carry out the optimisation that yields the weighting parameters, it is necessary to be able to calculate the lifetime financial income and fatigue damage that results from the use of specific values of $\lambda_{w_{win}}$ and $\lambda_{w_{roy}}$. This is achieved by firstly evaluating the two power curves that result from the use of $\lambda_{w_{win}}$ and $\lambda_{w_{roy}}$, and by then using these curves to calculate the lifetime income and fatigue damage.

To evaluate the power curves that a certain pair of weighting parameters will give, consider the operation of an optimising controller running in perfectly steady-state conditions at a specific wind-speed. Under these circumstances, the controller would simply vary the set point until it found the point of minimum cost. This same procedure can therefore be carried out to evaluate $P_{sp}(i_s, \overline{v})$. For a given average wind-speed $\overline{v}'$ in season i_s', the optimal power set point $P_{sp\_opt}(i_s', \overline{v}')$ that gives the minimum cost of operation $J_{min}\left[P_{sp\_opt}\left(i_s', \overline{v}'\right)\right]$ is found:

$$J_{min}\left[P_{sp\_opt}\left(i_s', \overline{v}'\right)\right] = \min_{P_{sp}(i_s', \overline{v}')} \left\{J\left[\lambda_w(i_s'), P_{sp}\left(i_s', \overline{v}'\right)\right]\right\}$$

$$\text{for} \qquad 0 \leq P_{sp}(i_s', \overline{v}') \leq P_{rated}$$

where $\lambda_w(i_s')$ is the weighting parameter for season i_s', and $J\left[\lambda_w(i_s'), P_{sp}\left(i_s', \overline{v}'\right)\right]$ is evaluated using Equation 2.14 (or Equation 2.15 for one season per annum). The values of P_{sp_opt} are then calculated for all wind-speeds $\overline{v}$ in both seasons i_s, and the result is the required pair of power curves $P_{sp}(i_s, \overline{v})$.

2.5 Simple cost function for interacting turbines

A simple cost function is presented for those turbines within a windfarm that, at any particular time, are interacting with one another, and that would therefore be under multi-input multi-output supervisory control.

The cost function for non-interacting turbines is extended directly to groups of interacting turbines, using sums for the damage and energy capture terms and using the same weighting parameter as for non-interacting turbines:

Simple cost function for a group of interacting turbines and two tariffs p.a.

$$J(i_s) = \sum_{j=1}^{n_t} \int_{t-T_c}^{t} \left[\lambda_w(i_s) \, . \, \dot{d}^{(j)}(t) - \dot{m}^{(j)}(t)\right] dt \qquad (2.16)$$

Simple cost function for a group of interacting turbines and one tariff p.a.

$$J = \sum_{j=1}^{n_t} \int_{t-T_c}^{t} \left[\lambda_w \, \dot{d}^{(j)}(t) - \dot{m}^{(j)}(t)\right] dt \qquad (2.17)$$

j = turbine number

n_t = total number of turbines in windfarm

$\dot{d}^{(j)}(t)$ = rate of accumulation of fatigue damage for turbine j

$\dot{m}^{(j)}(t)$ = rate of income for turbine j

The optimal weighting parameters that are developed for the supervisory control of non-interacting turbines may not be optimal for groups of turbines. The optimal weighting parameters for groups of turbines will depend upon:

- the numbers of turbines standing in a row
- the inter-turbine spacings
- the topography of the land upon which the turbines stand

Therefore the optimal weighting parameters will be heavily site-dependent. However to evaluate the optimal weighting parameters off-line, it is necessary to simulate all combinations of interacting turbines in every windfarm to which a supervisory controller is to be applied, and simple models of wake and topographic effects that can give sufficient accuracy in simulation are not yet available. Consequently a universally applicable weighting parameter is considerably more practical, even though it will generally be sub-optimal.

Any universally applicable weighting parameter that is proposed must be tested to find out whether it can significantly improve performance. To test a weighting parameter, a similar procedure is followed to that for non-interacting turbines. Firstly, the power curves that result from the use of the weighting parameter in an on-line controller are evaluated, and then the total lifetime income and damage that arise from the application of the power curves are calculated.

There are two possible methods for finding the power curves that result from the use of the weighting parameter. In the first method, data for the time averages $\overline{\dot{d}}$ and $\overline{\dot{e}}$ is generated for all combinations of power set points for the interacting turbines at wind-speeds across the operating range. This is the same approach as for the non-interacting turbines, except that in this case the number of dimensions of the region for which data must be found is equal to the number of turbines in the group. As a result, even for a small group of turbines the number of data points that are required for $\overline{\dot{d}}$ and $\overline{\dot{e}}$ is so large that the computation time required is prohibitive. In the second method, an optimisation technique is used in conjunction with the simulation. For each wind-speed across the operating range in turn, successive simulation sweeps are carried out. The optimisation algorithm varies the power set points at each simulation sweep and searches for the set of power set points that give the minimum cost of operation J (J is evaluated from $\overline{\dot{e}}$ and $\overline{\dot{d}}$). Once the cost of operation converges to within a specified tolerance, the search terminates and the optimum set of power set points are recorded. This procedure is followed for all wind-speeds across the operating range, thereby building up the optimal power set point curves. This optimisation technique requires a considerably smaller number of evaluations of $\overline{d}$ and $\overline{e}$ to be carried out than the first method.

2.6 On-line implementation

The final step in the methodology for the development of windfarm controllers is the on-line implementation of the control strategy, for which steady-state optimising control techniques are utilised. Steady-state control can be readily applied to windfarms by exploiting the 'spectral gap' in the wind-speed spectrum (see Figure 1.2 on page 4). The spectral gap allows the wind to be considered as consisting of two components: a high frequency component (atmospheric turbulence), and low frequency components. The supervisory controller operates with such a time-step that the turbulent wind is apparently stationary, and its control action compensates for the lower frequency disturbances. The principal lower frequency disturbances are caused by the passage of frontal systems and by diurnal warming and cooling cycles. The latter cause diurnal variations in the stability of the planetary boundary layer and, consequently, diurnal fluctuations in the wind-speeds that are experienced by wind turbines.

An optimising controller must:

1. allow the system to settle to sufficiently close to steady-state following set point changes before reliable estimates of the elements in the cost function can be made;

2. have a time-step which is:

   (a) long enough so that estimation errors caused by fluctuations in wind-speed due to atmospheric turbulence are small;

(b) short enough so that the wind-speed is approximately stationary, that is, so that the lower frequency fluctuations in wind-speed are negligible;

3. have a time-step which is short enough so that, over the number of time-steps required to approach the optimum, the variations in the cost which are caused by low frequency fluctuations in the wind-speed are small.

To satisfy item 2(a), reference to the wind-speed spectrum in Figure 1.2 on page 4 indicates that the time-step of the optimising controller (T_c) must be greater than approximately 10 minutes. As long as the frequency $\frac{1}{T_c}$ remains within the spectral gap, then the larger the value of T_c, the more reliable will be the controller's estimates. Item 3 gives rise to a more restrictive lower bound upon the time-step than item 2(b), with the result that the frequency $\frac{1}{T_c}$ must be significantly larger than the frequencies on the lower edge of the 'spectral gap'. Hence T_c must be significantly smaller than approximately 2 hours. The smaller the value of T_c, the more time-steps the controller will have in which to achieve convergence before low frequency effects are felt. Consequently, when choosing T_c a trade-off has to be made between the reliability of the controller's estimates and the severity of the disturbances experienced by the controller.

There are two principal options for the type of supervisory control techniques that can be used:

- Model-based optimising control
- Extremum (or hill-climbing) control

In theory, model-based control should be able to offer higher controller performance than extremum control. For a review of some techniques for steady-state model-based optimising control, see (Roberts and Lin, 1991). However in this application it is required that the windfarm controller be applicable to any site, regardless of the topography or the layout of the turbines. Therefore the physical processes that have to be represented in a controller's model include: wind flow across complex topography; wake and wake interaction effects; wake and wake interaction effects in complex topography. Further, both the topographic effects and the wake effects change with wind-direction. Simple models for topographic and wake effects that provide sufficiently accurate results are not currently available, and models which could achieve results of sufficient accuracy would demand such extensive hardware on which to run in real-time that there would be no industrial interest in practical application.

As a result, the application of the simpler extremum control techniques is examined initially. The relatively wide nature of the spectral gap makes the use of such techniques a realistic option. Extremum control was popular in the 1950s and 1960s following the seminal work described in (Draper and Li, 1951), and in recent years there has been a revival in interest in the field, with cross-fertilisation of ideas from the areas of self-tuning control (Wellstead and Scotson, 1990) (Zarrop and Rommens, 1993) and neural networks (Nascimento et al., 1992).

2.7 Conclusions

This chapter has described a suitable structure for windfarm supervisory controllers and has presented a methodology for the development of such controllers. A lifetime performance function has been developed in which the design and operational issues are decoupled by the inclusion of fatigue damage not as a term within the performance function but rather as a constraint upon the lifetime profit. A similar approach has been adopted for the testing of controllers, whereby controller performance is measured by examining the fraction of the maximum possible financial income that the use of the controller generates together with the fraction of the maximum possible fatigue damage that the turbine or turbines experience.

To implement the methodology described in this chapter, a windfarm simulation is required, together with models of fatigue damage. Chapter 3 gives a description of, and presents validation results from, the windfarm simulation.

Chapter 3

Windfarm simulation for testing supervisory controllers

3.1 Introduction

This chapter describes those factors that must be considered when simulating windfarms for the purposes of designing and testing windfarm supervisory controllers, when the objectives of the supervisory control are the maximisation of financial income and the minimisation of fatigue damage.

There are two possible methods of modelling a windfarm:

1. Using a small-scale model in a wind tunnel;
2. Using a computer simulation.

An example of the former is an investigation that examines the mean and turbulent flow within windfarms of horizontal axis turbines, as well as the turbines' performance and their blade loads (Hassan et al., 1990).

Computer simulation provides a much greater degree of flexibility than the use of scale models, however. Firstly, any parameters within the model can be easily changed. Secondly, more controlled experiments can be carried out because precisely the same time series can repeatedly be input to the model. There has been more work carried out on the computer simulation of windfarms than on the use of scale models, and two types of computer models have been developed to date:

1. Steady-state models for the prediction of windfarm energy output, for the purposes of evaluating optimal site layouts;
2. Dynamic models for the prediction of power fluctuations on grid systems.

Time-domain simulations have been developed for examining the effect of significant penetrations of windpower on large utility grids, on the time-scale of minutes to hours (Bouwmeester et al., 1985) (Peltola et al., 1992) (Schlueter et al., 1983b) (Schlueter et al., 1984) (Simburger and Cretcher, 1983). Fewer complete models have been developed for the detailed examination of fluctuations from windfarms on the time-scale of seconds, that is, incorporating models of wake effects and other higher-frequency dynamics (Estanqueiro et al., 1993) (van Oort, 1988). These models have also been developed for the purposes of examining windfarm power quality and the stability of the grid. Significant research effort is ongoing in the areas of wake and topographic modelling, and a review of developments in these fields is presented in (van der Snoek, 1988).

Although no models appear to have been reported for the testing of windfarm control algorithms for which the minimisation of fatigue damage is an objective, a variety of models have been produced for the testing of controllers of individual wind turbines for which the minimisation of fatigue damage is an objective. A review of models of wind turbines for control system design is given in (Leithead et al., 1992) and the simulation of wind turbines for control system design is discussed in (Leithead et al., 1989).

In this thesis, a computer simulation of a small windfarm is used as the test-bed for supervisory control algorithms. An overview of the principal characteristics of the simulation is given in Section 3.2, and a wind model that generates the wind-speed signal from a requested average wind-speed and turbulence intensity is described in Section 3.3. Probability distributions for the wind-speed over the long term are presented in Section 3.4, and models of the turbines' wakes, and of the interactions between overlapping wakes, are described in Section 3.5. Section 3.6 presents results which illustrate the validity of the turbine model, and Section 3.7 presents results for the mean power and long-term energy capture of a single turbine as functions of the power set point and mean wind-speed. Section 3.8 presents results which illustrate the validity of the wake and wake superposition models.

3.2 A test-bed windfarm simulation

An overview of the windfarm simulation that was used as a test-bed for supervisory control algorithms is presented in this section.

The following two site-dependent characteristics have a significant effect upon the performance of turbines within a windfarm (see Section 1.1 for further details):

1. The operation of turbines within the wakes of one or more other turbines;
2. The topography of the windfarm site.

Both of these characteristics affect the mean wind-speeds ($\overline{v}$) and the turbulence intensities (I) that are experienced by the individual turbines. The wake effects additionally introduce dynamic linkage between turbines, including time delays in the airflow between the turbines. The effects upon $\overline{v}$ and I plus the dynamic interconnection between turbines all have an impact upon the design of windfarm supervisory controllers and must be represented in any simulation that is to be used for the design and testing of such controllers.

The layout that is used in the test-bed simulation includes wake effects but not topographic effects. A row of three turbines standing on flat land is modelled, with the wind passing directly along the row of turbines. Therefore one of the two important effects on average wind-speed and turbulence associated with grouping turbines together in windfarms is represented, as well as dynamic linkage between the turbines. Almost all windfarms contain considerably more than three turbines, but the run-times involved with larger simulations would have been prohibitive. The models currently available for wakes and wake interaction are still relatively crude, however as long as the results of the simulation credibly represents the operation of three turbines within a typical windfarm, the simulation is suitable as a test-bed. As the effect of irregular topography is to give rise to variations in $\overline{v}$ and I, and both of these variables are also affected by wakes, then if performance improvements can be obtained on sites where only wake effects are present, it is likely that performance improvements will also be achieved on sites with significant topographic effects.

Two sets of inter-turbine spacings are investigated, as illustrated by Figures 3.1(a) and 3.1(b). The configuration with the larger spacings in Figure 3.1(a) is taken from the layout of a 25-turbine windfarm that was planned for a site in Wales (Gamble, 1990). The 7.9D/7.0D spacings illustrated in Figure 3.1(a) represent the smallest inter-turbine spacings of any turbines standing in a row within the windfarm. This windfarm was, however, planned as a demonstration project and consequently the layout is conservative, with larger spacings between the turbines than would typically be found in a commercial windfarm. The 3D/3D spacings illustrated in Figure 3.1(b) are more typical of the closest spacings that are found in commercial windfarms.

The particular design of wind turbine that is modelled in the windfarm simulation

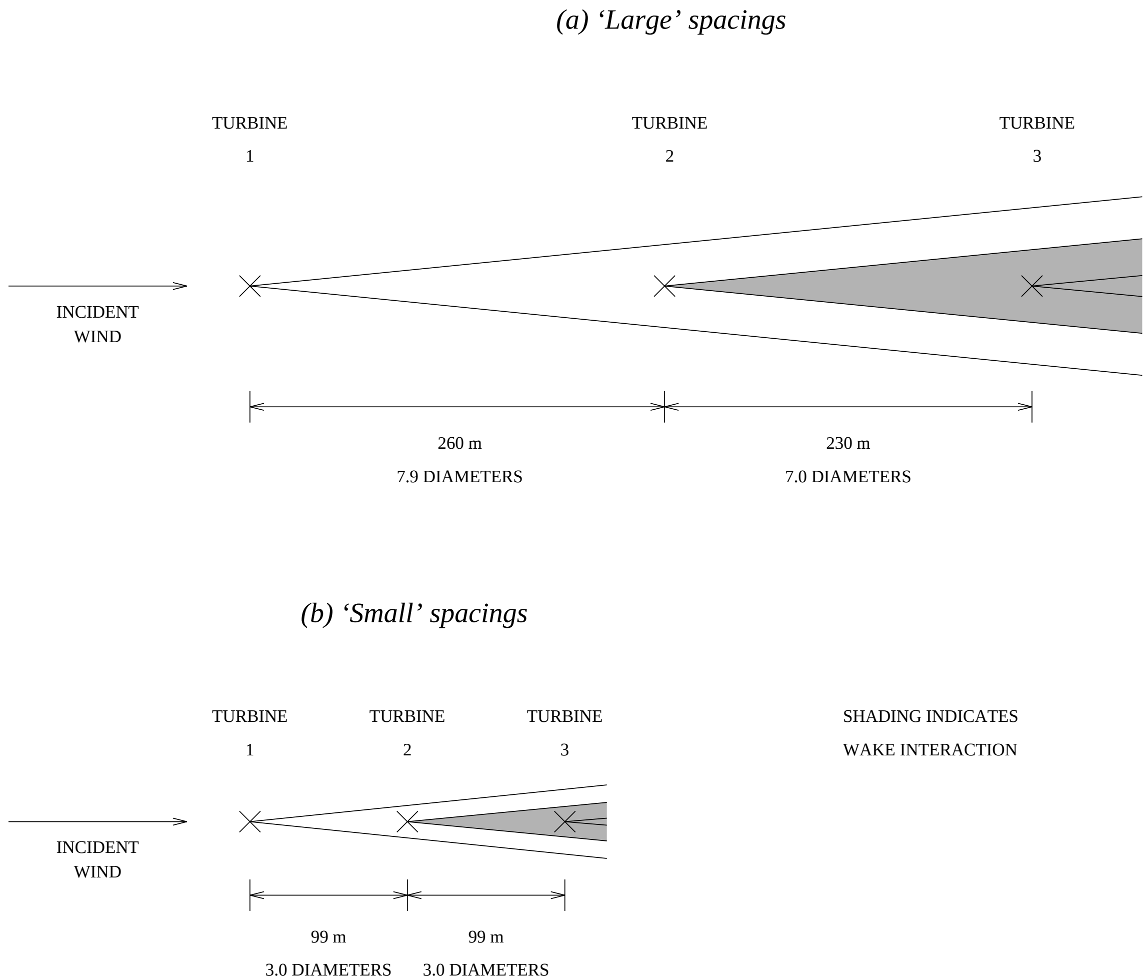

Figure 3.1: Layouts of turbines modelled in the windfarm simulation

is the MS–3, a commercial design that is dynamically similar to most other commercially available designs. Its principal characteristics are listed in Table 3.1 (Spruce, 1988b), the model of the turbine is described in detail in Appendix B (pp. 195*ff*), and further parameters are given in Appendix C (pp. 204*ff*). In addition, a line diagram of the turbine is presented in Figure 1.3, its power set point and achieved power curves are illustrated in Figure 1.4, and its ex-works component costs are presented in Table 1.3.

Rotor axis	horizontal
Rotor position	upwind
Number of blades	2
Power regulation	full-span blade pitch control
Hub type	teetered
Rated power output	300 kW
Rotor diameter	33 m
Rotational speed of low-speed shaft	48 rpm
Rotational speed of high-speed shaft	1500 rpm
Cut-in wind-speed	5 m/s
Rated wind-speed	11.5 m/s
Cut-out wind-speed	25 m/s

Table 3.1: Principal characteristics of the wind turbine modelled in the windfarm simulation

The slowest dynamics in the simulation are the time averages of the rates of energy capture and fatigue damage. The estimation is carried out over long enough periods for the estimation errors caused by fluctuations in wind-speed due to atmospheric turbulence to be small. In practice, this gives minimum convergence times of approximately 10 minutes. The fastest dynamics in the simulation, by contrast, are those of the turbines' pitch controllers, the time-step of which is 60 ms. As the difference in the time-scales of the fastest and slowest dynamics is several orders of magnitude, the system is dynamically 'stiff'.

The 'stiff' nature of the system and the stochastic nature of the wind necessitate the use of a variable-step integration algorithm (Leithead et al., 1989). The Adams-Moulton algorithm was selected for this application, because the simulations run approximately four times faster when using this algorithm than when using the Gear's Stiff algorithm. The maximum integration time-step that is used is 0.003 seconds. To implement the simulation, a proprietary simulation software package is employed (ACSL, 1991). The suitability of this package for the simulation of windfarms is described in (Spruce and Dexter, 1991).

3.3 Model of the turbulent wind

The mean wind-speed ($\overline{v}$) and turbulence intensity ($I = \sigma_v/\overline{v}$) of the wind that is incident upon the turbine which stands in the free stream must be specified at the start of each simulation run. $\overline{v}$ and I at the downstream turbines are then generated by the wake and wake interaction models.

For a desired mean wind-speed and turbulence intensity, the output from the wind model is a stochastic process with a Dryden spectrum and the required mean and standard deviation. This signal is generated as follows:

1. Gaussian white noise of zero mean and unity standard deviation is produced using the simulation package's random number generator.

2. This white noise is filtered so that it has a 'Dryden' spectrum* (Leithead et al., 1992)[Appendix H]. The Dryden spectrum can be defined as follows (Fordham, 1985)[p. 114]:

$$S_v(\omega) = \frac{K_d A_d \sigma_i^2}{1 + B_d \omega^2}$$

where

$$A_d = \frac{4L}{\overline{v}}$$

*As a variable time-step integration algorithm is employed in the simulation, interpolation of the resulting signal is carried out at each time-step (Leithead et al., 1992)[Appendix H].

$$B_d = \left(\frac{L}{\overline{v}}\right)^2$$

K_d = constant of multiplication

L = turbulence integral length scale. Set to 100 m throughout.

$S_v(\omega)$ = power spectrum of output

ω = frequency

σ_i = standard deviation of white noise input = 1

At the start of each simulation run, K_d is adjusted as necessary to achieve the required value of σ_v and hence the required value of I.

3. Finally, the resulting signal of zero mean is off-set by $\overline{v}$ to give the correct mean.

3.4 Long-term wind-speed probability distributions

The Weibull distribution is used to represent long-term wind-speed probability distributions. Its probability density is (Eggleston and Stoddard, 1987)[p. 68]:

$$p\left(v' < v < v' + dv'\right) = \frac{k_w}{c_w}\left(\frac{v'}{c_w}\right)^{k_w - 1} \exp\left[-\left(\frac{v'}{c_w}\right)^{k_w}\right] dv'$$

where v, v' = wind-speed

$p\left(v' < v < v' + dv'\right)$ = probability of v between v' and $v' + dv'$

k_w = Weibull shape factor

c_w = Weibull wind-speed normalising factor

The relationship with the annual mean wind-speed $\overline{v}_{year}$ is:

$$\overline{v}_{year} = c_w \Gamma \left(1 + \frac{1}{k_w}\right)$$

where $\Gamma(.)$ is the gamma function.

All measured wind-speed parameters that are referred to in this section had been measured at, or are scaled up to, 25 m above ground level, which is the hub height of the turbine that is modelled.

3.4.1 One season per annum

The Weibull parameters that are used for one tariff season per annum are as follows:

$$\begin{aligned} \overline{v}_{year} &= 7.64 \text{ m/s} \\ k_w &= 1.92 \\ c_w &= 8.60 \text{ m/s} \end{aligned}$$

These parameters are average values of measured site data for six commercial wind-farm sites in Wales. The site data was supplied by the manufacturer of the turbine that is modelled (Valpy, 1992). The probability density associated with the above Weibull parameters is plotted in Figure 3.2.

3.4.2 Two seasons per annum

Seasonal wind-speed parameters are required for the following two seasons:

- Winter December to February
- Rest Of Year March to November

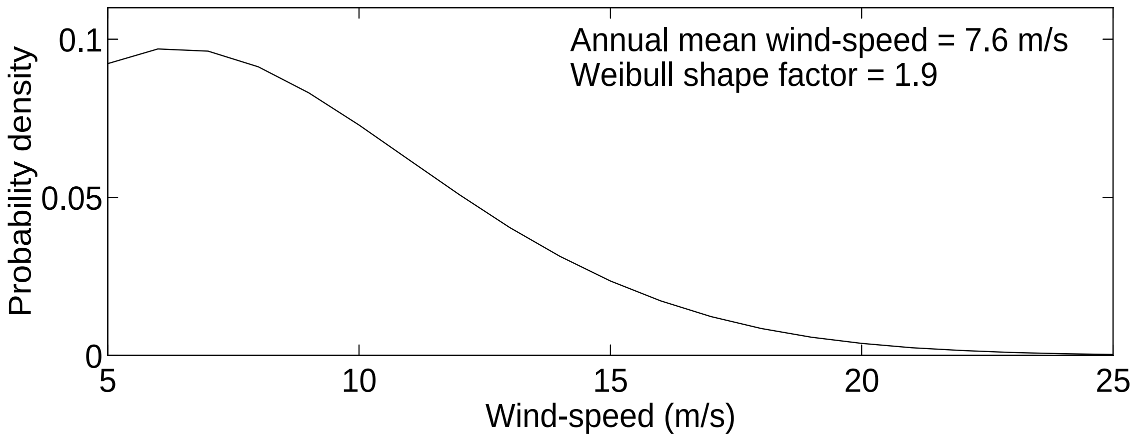

Figure 3.2: Annual wind-speed probability density

It is assumed that the Weibull shape factor is the same in both seasons. This assumption is based upon the results of a comprehensive analysis of long-term measured data from fourteen meteorological stations at dispersed sites around the U.K. (Halliday, 1983). The results of the study indicate that seasonal variations in Weibull shape factor are typically small and contain no obvious trend.

The averaged Weibull parameters for the windfarm site data are similar to the specific data for the Scilly meteorological station in Halliday's report. For the Scilly station, $\overline{v}_{year} = 7.3$ m/s and $k_w = 2.0$. The averaged annual mean wind-speed for the windfarm sites is split between the Winter and Rest Of Year seasons such that the ratios of the seasonal mean wind-speeds to the annual mean wind-speed are the same as those for the Scilly data.

The resulting Weibull parameters for the two seasons are as follows:

	Winter	Rest Of Year
$\overline{v}$	9.4 m/s	7.1 m/s
k_w	1.92	1.92
c_w	11 m/s	8.0 m/s

The associated probability densities are plotted in Figure 3.3.

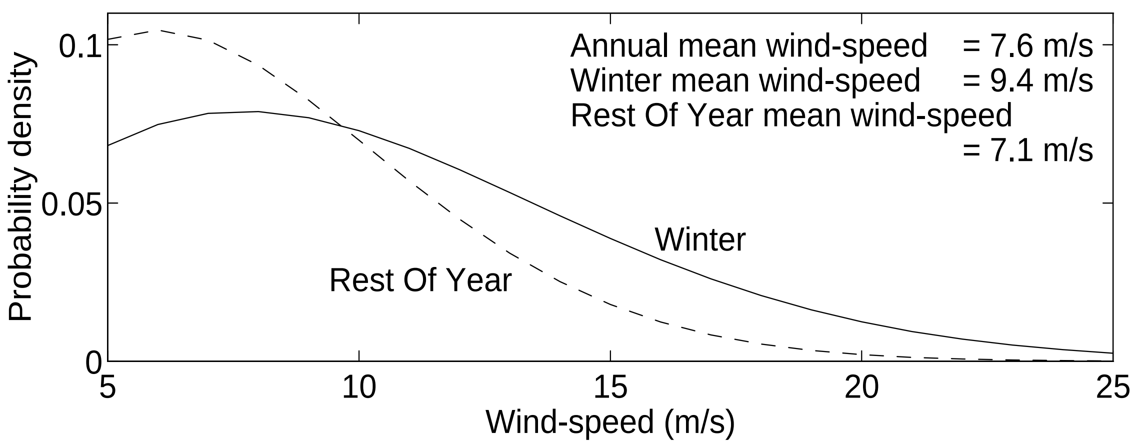

Figure 3.3: Wind-speed probability densities for Winter and Rest Of Year

3.5 Wake and wake interaction models

The wake and wake interaction models exploit the lack of coherence of high frequency changes in wind-speed in the longitudinal direction. The wind-speed in the wake is split into two components, one for low frequency fluctuations in wind-speed and the other for high frequency fluctuations. The result is a simple model which represents the changes in $\overline{v}$ and I downstream of an operational turbine as well as the time-delay associated with changes in wind-speed as the air moves between turbines.

In this section, the basic wake and wake interaction models are described. A refinement to the basic models that corrects the standard deviation of the wind-speed at the downstream turbines is described later, in Section 3.8.3, following results for the basic models.

3.5.1 Longitudinal coherence of wind-speed

Coherence γ_{xy} can be used to represent the success that fluctuations in wind-speed have in moving from turbine to turbine. Coherence is a cross-spectral frequency function which is defined as (Bendat and Piersol, 1986)[p. 535]:

$$\gamma_{xy}(f) = \frac{|S_{xy}(f)|}{\sqrt{S_{xx}(f)S_{yy}(f)}}$$

where f is frequency (Hz) and $S_{xy}(f)$ is the cross-spectral density of signals $x(t)$ and $y(t)$. For two points separated by distance Δx, in average wind-speed $\overline{v}$, coherence can be expressed as (Frost et al., 1978)[p. 4.87]:

$$\gamma_{xy}(f) = \exp\left(-\frac{a\,\Delta x}{\overline{v}} f\right)$$

where a is the decay coefficient. The value used for a is 4.5, as given in the literature for the longitudinal case. The variation of the longitudinal coherence of fluctuations of wind-speed with a normalised frequency η is presented in Figure 3.4, with η plotted on a logarithmic scale. η is given by:

$$\eta = \frac{\Delta x}{\overline{v}} f \tag{3.1}$$

Figure 3.4 shows that higher frequency changes in wind-speed are incoherent between the two points, and only lower frequency, 'bulk' changes reach the downstream point.

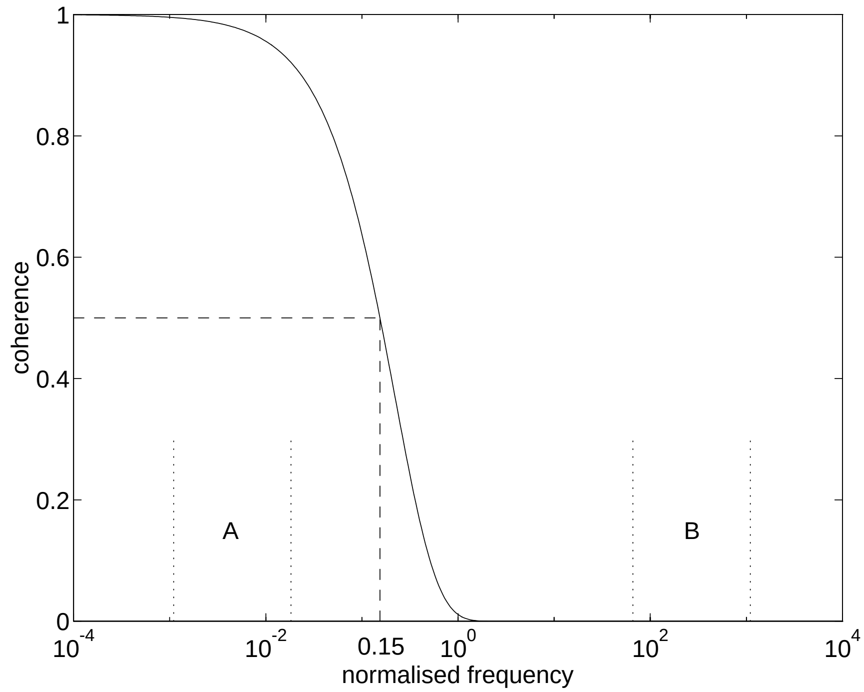

Figure 3.4: Longitudinal coherence of wind-speed vs. normalised frequency

The wake model is split into two components, one for low-frequency fluctuations in wind-speed and the other for high frequency fluctuations. The dividing frequency is defined as that for which $\gamma_{xy} = 0.5$ and, as shown by Figure 3.4, is situated at $\eta = 0.15$.

The model is developed for the 'intermediate' and 'far' sections of the wake, which is the region beyond approximately three rotor diameters (D) downstream of the turbine generating the wake (Garrad Hassan, 1989). In many windfarms, those turbines that operate in the wakes of other turbines do so within the intermediate or far wake region.

The wide range of frequencies over which the simulation must be capable of operation can be illustrated by considering the regions A and B in Figure 3.4. The following ranges for Δx and $\overline{v}$ are used when evaluating the ranges of frequencies indicated by regions A and B:

$$3D < \Delta x < 10D \qquad\qquad 5 \text{ m/s} < \overline{v} < 25 \text{ m/s}$$

Region A then illustrates the range of values of η that correspond to the absolute frequency $f = 1/(1 \text{ hour})$:

$$1.1 \times 10^{-3} < \eta < 1.8 \times 10^{-2}$$

And region B illustrates the range of values of η that correspond to $f = 1/T_{pi}$, where T_{pi} (= 60 ms) is the time-step of the turbine's pitch controller:

$$6.6 \times 10 < \eta < 1.1 \times 10^{3}$$

Both of the frequencies $1/(1 \text{ hour})$ and $1/T_{pi}$ are within the resolution of the simulation. The fluctuations of wind-speed with frequencies in region A are almost completely coherent whilst the frequencies in region B have negligible coherence.

A block diagram of the wake model is presented in Figure 3.5. This diagram illustrates the division of the wake model into low and high frequency signals. The filters F1 to F3 remove the frequencies which are assumed to be incoherent from turbine to turbine, that is, those frequencies for which $\eta > 0.15$. Using Equation 3.1, the cut-off frequency of filters F1 to F3 is then given by:

$$f_0 = 0.15 \frac{\overline{v}}{\Delta x} \tag{3.2}$$

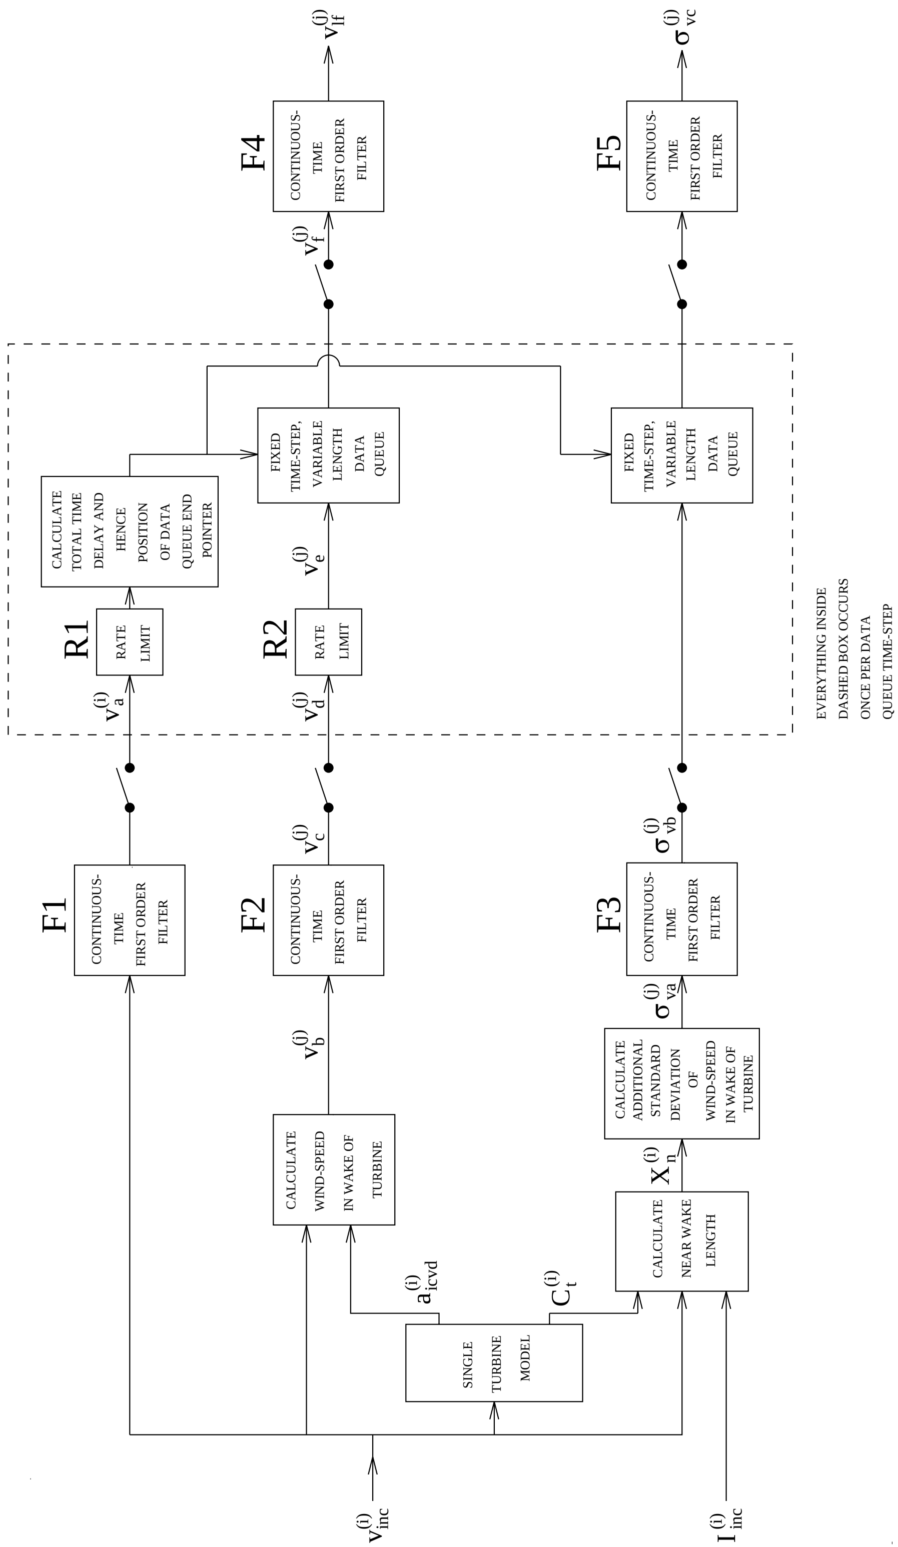

Figure 3.5: Block diagram of model of wake from turbine i to turbine j

3.5.2 Low-frequency fluctuations in wind-speed

The model that is used to calculate the low frequency component of wind-speed at the downstream turbine assumes jet-like flow, in which the wind-speed is constant across the wake at any given distance downstream (Katić et al., 1986):

$$v_b^{(j)} = v_{inc}^{(i)} \left[1 - \frac{2a_{icvd}^{(i)}}{1 - \frac{2k\Delta x^{(i)(j)}}{D^{(i)}}} \right]$$

where $v_b^{(j)}$ = wind-speed that will occur at downstream turbine j
$v_{inc}^{(i)}$ = wind-speed incident upon upstream turbine i
$a_{icvd}^{(i)}$ = initial centre-line velocity deficit behind rotor of turbine i
k = gradient of the spread of the wake, taken from the literature as 10%

$v_b^{(j)}$, $v_{inc}^{(i)}$ and $a_{icvd}^{(i)}$ are marked on Figure 3.5.

3.5.3 High-frequency fluctuations in wind-speed

The model that is used to calculate the high frequency component of wind-speed is based upon a general empirical description of the additional turbulence intensity caused in the wake of a wind turbine (Garrad Hassan, 1989). The model is valid for the intermediate and far wake regions, and has the following form:

$$I_{add}^{(j)} = 4{\cdot}8 \left(C_t^{(i)}\right)^{0{\cdot}7} \left(I_{inc}^{(i)}\right)^{0{\cdot}68} \left(\frac{\Delta x^{(i)(j)}}{X_n^{(i)}}\right)^{-0{\cdot}57}$$

where $I_{add}^{(j)}$ is the additional turbulence intensity that will be caused at downstream turbine j; $C_t^{(i)}$ is the thrust coefficient of turbine i (see Appendix B); $I_{inc}^{(i)}$ is the turbulence intensity of the wind incident upon upstream turbine i; $X_n^{(i)}$ is the length of the near wake of turbine i. $X_n^{(i)}$ is a function of $C_t^{(i)}$, $v_{inc}^{(i)}$, $I_{inc}^{(i)}$ and machine parameters (Garrad Hassan, 1989). All of the above variables are marked on Figure 3.5.

$I_{add}^{(j)}$ is used to calculate $\sigma_{va}^{(j)}$, which is filtered in the same way as $v_b^{(j)}$ to give $\sigma_{vb}^{(j)}$. $\sigma_{vb}^{(j)}$ is then passed through the delay model in the same manner as $v_c^{(j)}$.

3.5.4 Transport delay

The time-delay of the low and high frequency components of wind-speed is modelled by a data 'queue' plus reconstruction filters. Figure 3.6 illustrates the correct operation of the transport delay model for the low frequency wind-speed signal. All

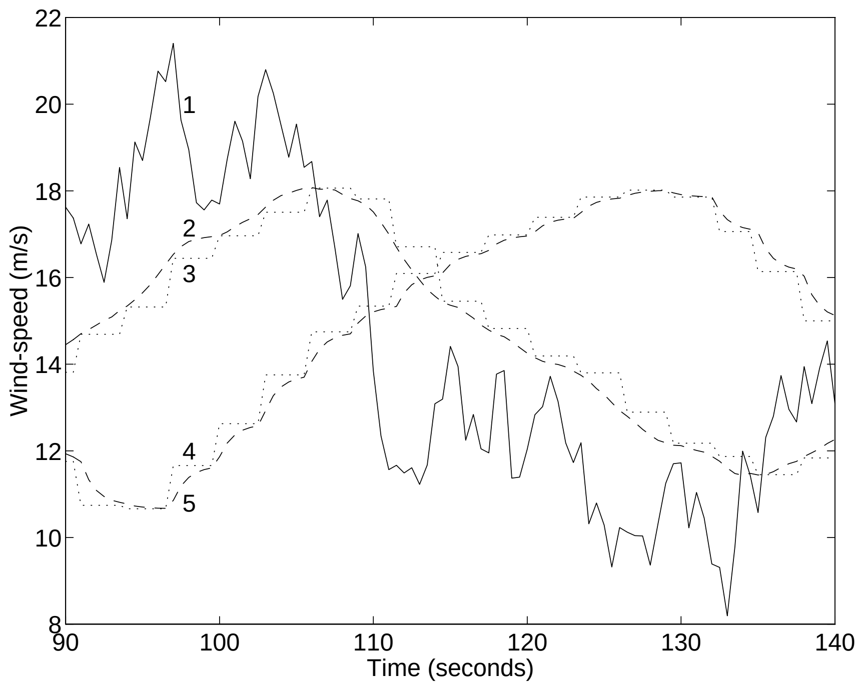

Figure 3.6: Wind-speeds at various points within the low frequency delay model

variables plotted in Figure 3.6 are marked on Figure 3.5. Curve 1 is $v_b^{(j)}$ and curve 2 is $v_c^{(j)}$, the wind-speed signal following the initial filtering. Curve 3 is $v_e^{(j)}$, the sampled signal before the time-delay, and curve 4 is $v_f^{(j)}$, the signal at the output of the time-delay. Curve 5 is $v_{lf}^{(j)}$, the filtered, delayed version of $v_b^{(j)}$ that forms the low frequency component of the wind-speed that is incident upon the downstream turbine. (No rate limiting takes place during the period of operation illustrated in Figure 3.6.)

The cut-off frequency of the filters F1 to F3, f_0, is given in Equation 3.2. It is assumed that the frequency responses of these three filters are reduced to a negligible level by $10f_0$. The Nyquist interval corresponding to a frequency of $10f_0$ is $1/(20f_0)$ and hence the signals $v_c^{(j)}$ and $\sigma_{vb}^{(j)}$ are sampled at an interval of $1/(20f_0)$, to allow their recovery without aliasing (Poularikas and Seely, 1991)[p. 448]. The output from the data queue is passed through filters F4 and F5, which have cut-off frequencies of $10f_0$, to reconstruct the signals $v_c^{(j)}$ and $\sigma_{vb}^{(j)}$ (Poularikas and Seely, 1991)[p. 450].

The data queue is updated at a constant rate, and is an array with a variable number of elements in use at any one time. The current value of the time-delay is used to determine the position of the end pointer of the queue. The rate limits R1 and R2 are inserted to ensure that the behaviour of the queue does not become anti-causal. This would have been possible if $v_d^{(j)}$ decreased fast enough for the end pointer to move down the array so quickly that it pointed to an element that had already been output. For the operation of the data queue to be causal:

$$|\Delta T_{delay}| < T_{step}$$

where ΔT_{delay} is the change in the length of the time delay represented by the data queue over one data queue time-step (in seconds), and T_{step} is the length of the data queue's time-step (in seconds). This relationship yields the following upper limit for changes in $v_a^{(i)}$ ($v_a^{(i)}$ is marked on Figure 3.5):

$$\Delta v_{max} = K_{rl} \frac{\left(v_a^{(i)}\right)^2}{20\ \eta_0\ v_{lfmax}}$$

where Δv_{max} is the magnitude of the maximum change in $v_a^{(i)}$ that can occur in one data queue time-step; K_{rl} is a constant that is used for tightening the rate limit ($0 \leq K_{rl} \leq 1$); η_0 is the maximum value of η for which frequencies are assumed to be coherent from turbine to turbine ($= 0.15$); v_{lfmax} is the maximum value of the low frequency component of v that it is required to pass through the time-delay (set to 30 m/s). The maximum allowable change Δv_{max} is applied to both $v_a^{(i)}$ and $v_d^{(j)}$. The position of the end pointer of the data queue is continuously monitored and, if the queue ever becomes anti-causal, the simulation terminates. Results showed that without the constant K_{rl} the data queue was still anti-causal under certain circumstances. K_{rl} was therefore applied, a range of values were tested, and a value of 0.5 has been found to prevent any anti-causality across the whole operating range of wind-speeds, at all turbulence intensities that have been investigated ($10\% < I < 20\%$), and for both the 3D/3D and 7.9D/7D inter-turbine spacings.

3.5.5 Synthesis of wind-speed at a downstream turbine

Turbine within a single wake

When there is no wake interaction, as at turbine 2 in Figure 3.1, the wind-speed incident upon the downstream turbine ($v_{inc}^{(j)}$) is calculated as shown in Figure 3.7. The total standard deviation of the incident wind-speed ($\sigma_{vinc}^{(j)}$) is calculated from the standard deviations of the free wind ($\sigma_{v\infty}$) and the additional standard deviation caused by operation within the wake:

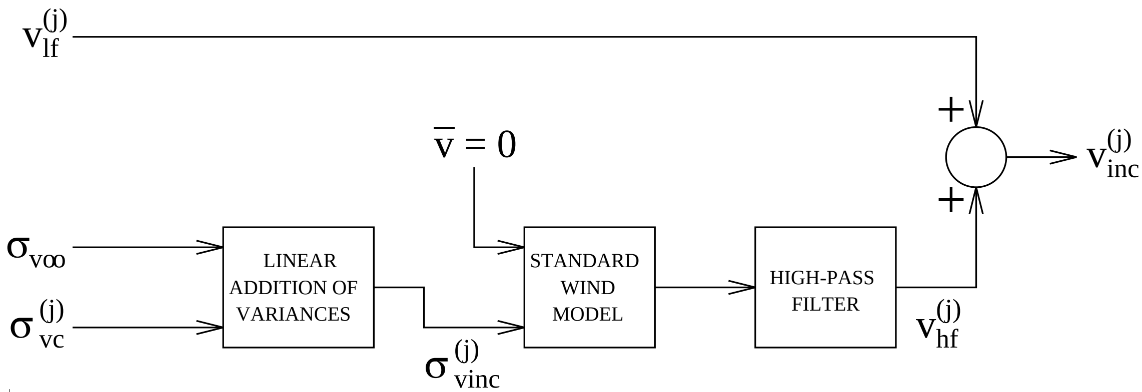

Figure 3.7: Synthesis of wind-speed at a turbine within a single wake

$$\sigma_{vinc}^{(j)} = \sqrt{\sigma_{v\infty}^2 + \left(\sigma_{vc}^{(j)}\right)^2}$$

A wind-speed of zero mean and standard deviation $\sigma_{vinc}^{(j)}$ is then generated using the standard wind model (see Section 3.3). Although the wind-speed generated by the model has zero mean, the constants A_d and B_d which are used in the model (step 2 in Section 3.3) are functions of $\overline{v}$. The value for $\overline{v}$ that is used for calculating A_d and B_d is equal to that of the free stream wind incident upon the upstream turbine. The signal which is generated by step 2 of the wind model receives an off-set of 0 m/s in step 3, and hence the signal produced by the wind model has zero mean and standard deviation $\sigma_{vinc}^{(j)}$.

The modelling of the wind-speed in the wake of an operational wind turbine with a Dryden spectrum is supported by results which are presented in (Quarton and Morgan, 1989). In this study, the wind-speed was measured just upstream and 5D downstream of an operational turbine at Nibe in Denmark. The resulting spectra of the wind-speed at the two points both have approximately the same shape, thereby providing some evidence that the shape of the spectrum of wind-speed does not alter within turbines' wakes.

The signal that is generated by the standard wind model is passed through a high-pass filter to remove the low frequency part of the wind-speed at the downstream turbine which is represented by $v_{lf}^{(j)}$. The signal at the output of the high-pass filter ($v_{hf}^{(j)}$) is then added to $v_{lf}^{(j)}$ to give the wind-speed incident upon the downstream turbine ($v_{inc}^{(j)}$).

First order filters are used for both the high-pass filter in Figure 3.7 and the low-pass filter F2 in Figure 3.5. It is these two filters that separate the wind-speed in the wake into low and high frequency components, prior to the summation of the low and high frequency components to give $v_{inc}^{(j)}$, and consequently the use of first order filters must introduce errors in the spectra of $v_{inc}^{(j)}$. However, if the resulting spectra for $v_{inc}^{(j)}$ are still typical of those found in the wakes of turbines operating within a windfarm, then the wake model will serve its purpose of providing a suitable test-bed for supervisory controllers.

Turbine within two interacting wakes

When two or more wakes overlap, as at turbine 3 in Figure 3.1, the low frequency component of wind-speed $v_{lf}^{(j)}$ is calculated by subtracting the velocity deficits caused by the upstream machines ($\Delta v_{lf}^{(i)}$) from the free stream wind-speed ($v_{lf\infty}$) (Vermeulen et al., 1981):

$$v_{lf}^{(j)} = v_{lf\infty} - \Delta v_{lf}^{(1)} - \Delta v_{lf}^{(2)} - \dots \tag{3.3}$$

In this case the velocity deficits and $v_{lf\infty}$ are passed through the delay model instead of $v_b^{(j)}$.

The high frequency component of wind-speed is calculated in the same way as for the previous case of a turbine operating within only one wake. The only difference is that the standard deviation is now calculated as:

$$\sigma_{vinc}^{(j)} = \sqrt{\sigma_{v\infty}^2 + \left(\sigma_{vc}^{(1)}\right)^2 + \left(\sigma_{vc}^{(2)}\right)^2 + \ldots} \tag{3.4}$$

As in the case of the single wake, a high frequency signal of zero mean is generated and passed through a high-pass filter before being added to the low frequency signal to give $v_{inc}^{(j)}$.

3.6 Validation of simulation results for single wind turbine

3.6.1 Variation with mean wind-speed

Mean power curve

In Figure 3.8, the turbine's simulated power curve is compared with a power curve from measured site data (Bossanyi et al., 1992)(Phase 2), for the standard operating practice of $P_{sp} = P_{rated}$ across the operating range. For the range of wind-speeds for which site data is available, the simulation results fit the measured data well. (Measured data is not available for $\overline{v} > 19$ m/s.) The site data is for an MS-3 turbine at Myers Hill near Glasgow and is offset by +2 m/s to compensate for a miscalibration of the wind-speed signal which is described in (Bossanyi et al., 1992)(Phase 2).

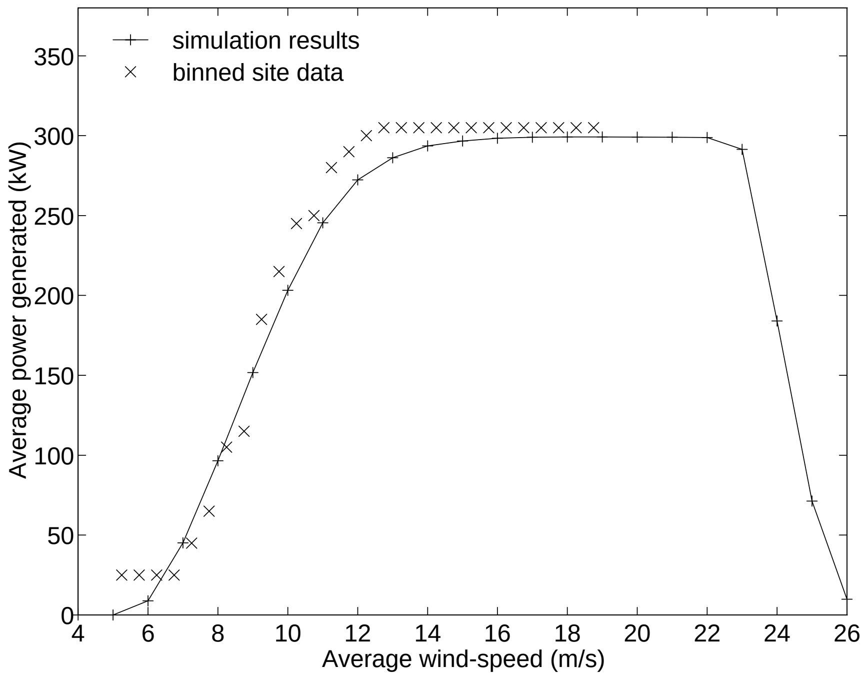

Figure 3.8: Comparison of simulated and actual power curves for 1 turbine. ($P_{sp} = P_{rated} = 300$ kW)

The fall in the simulated mean power in the region 23 m/s $\leq \overline{v} \leq$ 26 m/s is due to the high wind-speed cut-out/in algorithm in the turbine's controller. The mean power level does not fall suddenly because of hysteresis in the simulated cut-out/in algorithm. (See page 203 in Appendix B for details of this algorithm.)

Standard deviation of power

Figure 3.9 compares the simulated standard deviation of generated power with measured data from two sites, as a function of mean wind-speed and for $P_{sp} = P_{rated}$ throughout.

The measured data for site 1 is from the same study as the measured power curve data that was presented in the previous section (Bossanyi et al., 1992)(Phase 2),

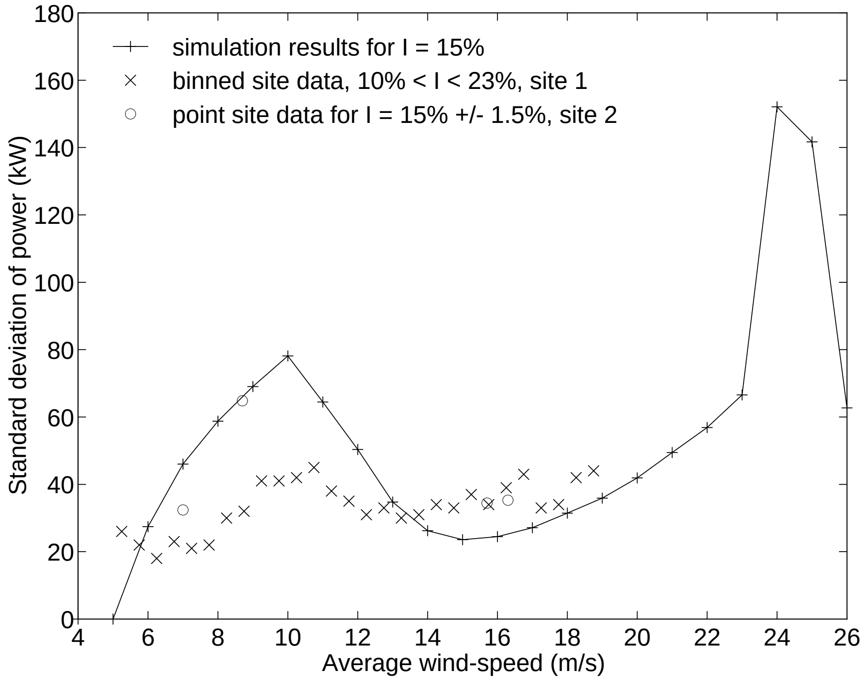

Figure 3.9: Comparison of simulated and actual standard deviation of power vs. wind-speed for 1 turbine. ($P_{sp} = P_{rated} = 300$ kW)

and is again offset by +2 m/s. A comparison of the simulated results with those for site 1 indicates reasonable agreement between the two data sets. As $\overline{v}$ increases, the standard deviation σ_p for site 1 rises to a maximum at approximately 10 m/s, then falls until a minimum is reached at approximately 13 m/s, and then starts to rise again. The shape of the simulated standard deviation is similar, although with a larger peak at 10 m/s. Other than modelling errors, two factors could account for the differences between the results: firstly, σ_p is sensitive to the tuning of the turbine's controller, and the proportional and integral gains that were used at site 1 are not known; secondly, σ_p is sensitive to the turbulence intensity I, and the site data is for a relatively wide range of turbulence intensities of between 10% and 23%.

The measured data for site 2 is for an MS-3 turbine at Carmarthen Bay in Wales and consists of individual 10-minute averages (Warren, 1990). The proportional and integral gains used in the simulated controller are the same as those for the

Carmarthen Bay turbine, and the data points that are plotted are all for I within $\pm 1.5\%$ of the 15% used in the simulation. A comparison of the simulated results with the four data points for site 2 indicates very close agreement. In particular, support for the magnitude of the peak in the region around 10 m/s is provided by the measured data point at 8.7 m/s which falls close to the simulated curve.

The simulated σ_p exhibits a large peak in the region where the turbine's upper cut-out/in algorithm operates, 23 m/s $< v <$ 25 m/s. Such a peak is to be expected because the turbine is off-line for large proportions of the time, leading to lower average power levels and therefore higher standard deviations of power.

A measured curve of σ_p versus $\overline{v}$ for the Nibe 'B' turbine (Madsen, 1988) provides further evidence to support the shape of the simulated σ_p–$\overline{v}$ curve. The Nibe 'B' is a medium-sized turbine which employs pitch control with a PI control algorithm to regulate the power, and hence it is similar, in broad terms, to the MS-3 turbine. The curve of σ_p versus $\overline{v}$ for the Nibe turbine, which is not plotted in Figure 3.9, covers the range 5 m/s $< \overline{v} <$ 18 m/s and has almost exactly the same shape as the simulated data across this range of wind-speeds.

3.6.2 Transient behaviour

In this section, nine time series are presented, one illustrating a typical section of simulated free stream wind and the other eight illustrating the simulated power output in various average wind-speeds and turbulence intensities.

The typical time series of simulated wind-speed is presented in Figure 3.10(a). A 50-second initialisation period preceded these results, to allow initialisation transients to decay. Initialisation periods of suitable duration preceded all simulations for which results are presented in this thesis.

With regard to the variation of power with average wind-speed, the operating curve (power curve) of achieved mean power versus mean wind-speed that was presented in Figure 3.8 shows that there are five distinct operating regions:

1. Around the lower cut-in/out.
2. Between the lower cut-in/out and the rated wind-speed.
3. On the 'knee' of the curve at around the rated wind-speed.
4. Between the rated wind-speed and the upper cut-out/in.
5. Around the upper cut-out/in.

Figures 3.10(b–c) and 3.11(a–c) present time responses for the simulated power output in regions 1–5 respectively. The dotted line on each plot illustrates the rated power of the turbine, 300 kW.

Figure 3.10(b), which corresponds to the wind-speed series in Figure 3.10(a), shows that in region 1 the turbine is off-line for relatively long periods of time, and negative power levels are generated for significant periods. (Details of the cut-in/out algorithm are presented on page 203 in section B.3.) The standard deviation of the power is therefore low, as illustrated by Figure 3.9.

Figure 3.10(c) shows that in region 2 there are wide fluctuations in power output, with occasional dips beneath 0 kW but very few excursions above 300 kW. This is because the pitch angle† is set to 0° for the majority of the time to maximise energy capture and, as $P_{aero} \propto v^3$, the fluctuations in wind-speed consequently cause large changes in power output (see Figure 3.9).

†The convention for pitch angle ϕ is as follows:

$\phi = 0°$	blades at 'fine'	maximum aerodynamic torque generated
$\phi = 90°$	blades at 'feather'	minimum aerodynamic torque generated

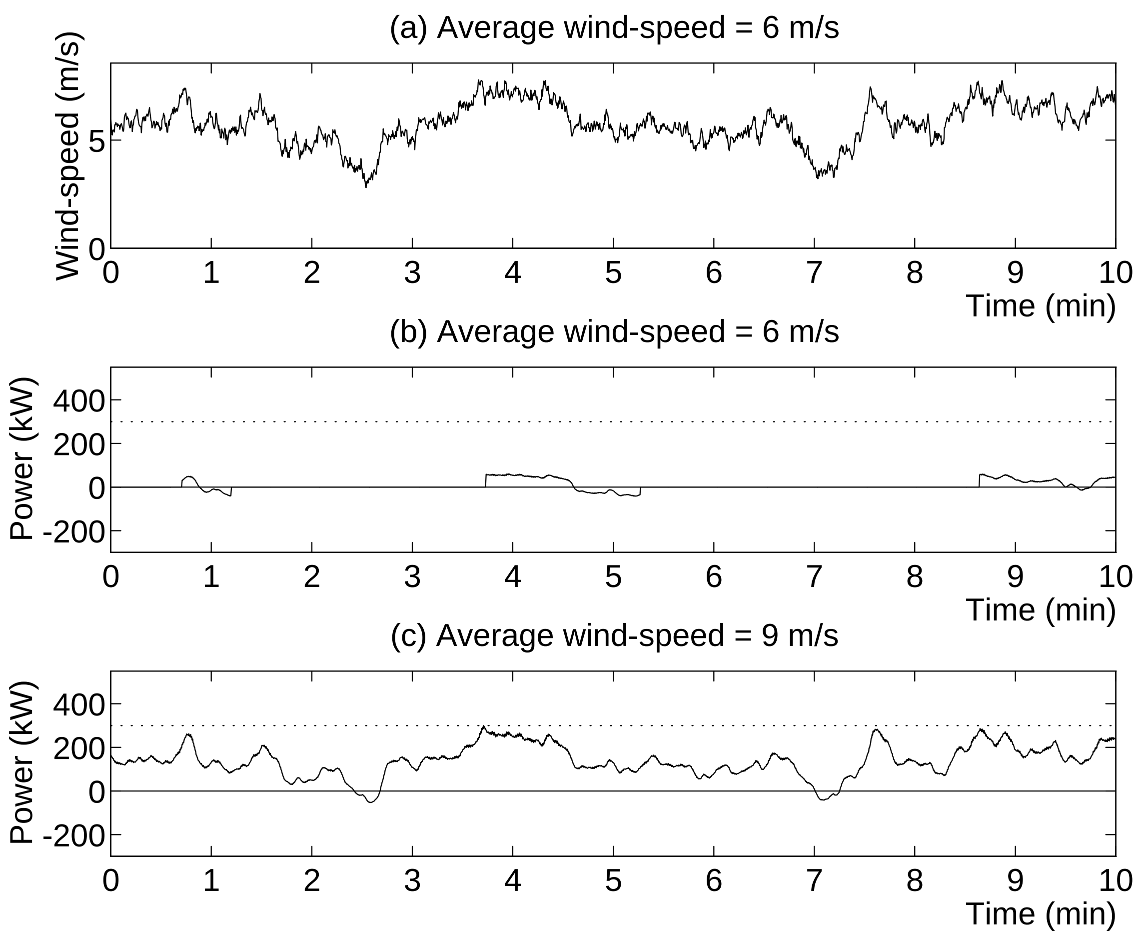

Figure 3.10: Wind-speed and power against time for $I = 15\%$. (1 turbine; $P_{sp} = P_{rated} = 300\ kW$.)

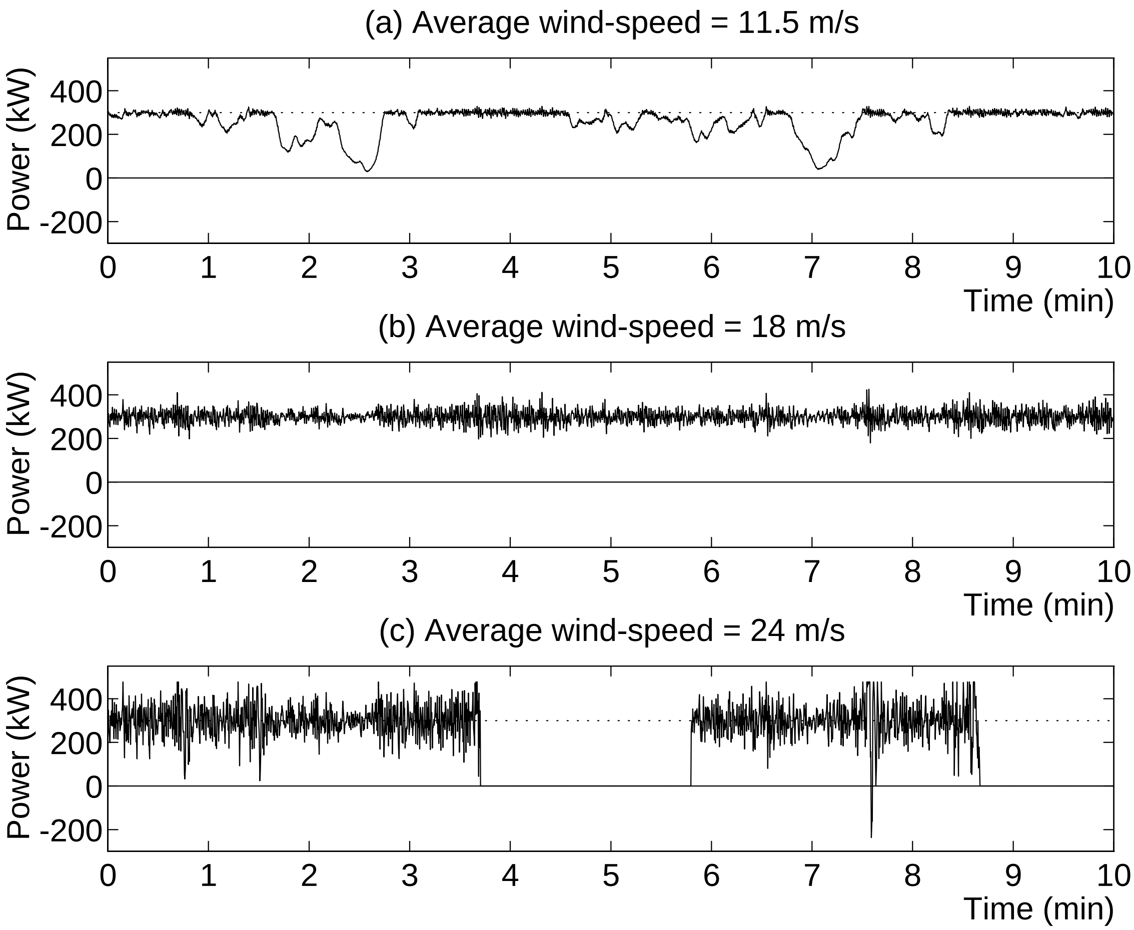

Figure 3.11: Power against time for various wind-speeds and $I = 15\%$. (1 turbine; $P_{sp} = P_{rated} = 300$ kW.)

Figure 3.11(a) shows that in region 3 there is good regulation to 300 kW for much of the time, but reductions in wind-speed beneath v_{rated} over anything other than short periods of time can cause large reductions in power output. As a result the mean power level is somewhat lower than 300 kW and the standard deviation of power is high in this region.

Figure 3.11(b) shows that in region 4 the dips in v below v_{rated} are negligible, so $\overline{P} \approx P_{rated}$. Also, σ_p is not particularly large in this region because the turbine's controller is able to regulate the power relatively effectively.

Figure 3.11(c) shows that in region 5 the operating regime is quite severe. σ_p is clearly very large, due to the difficulty of carrying out power regulation in this region, and yet $\overline{P}$ is significantly below P_{rated} because of the time spent off-line. (Details of the cut-in/out algorithm are presented on page 203 in section B.3.) Negative power excursions are even experienced occasionally. As described in section B.2, the turbine's high-speed shaft contains a slipping clutch which slips when $P_{hss} > 1.7\, P_{rated}$ and thereby causes P_{hss} to be limited to 1.7 P_{rated}. In Figure 3.11(c) the slipping action of this clutch can be seen to give a flat top to the highest power excursions in the plot, for example between $t \approx 7.5$ minutes and $t \approx 8.7$ minutes.

In general, Figures 3.10 and 3.11 show that, for average wind-speeds across the whole operating range, the model of the single turbine behaves sensibly and exhibits the correct basic operational characteristics in the time domain.

As well as considering the response of power as a function of $\overline{v}$, the variation of $P(t)$ as a function of the other variable that characterises the turbulent wind, the turbulence intensity, should be considered. Figure 3.12 presents time series of $P(t)$ for $I = 10\%$, 15% and 20%, at a constant average wind-speed of 11.5 m/s. (The $I = 15\%$ plot of Figure 3.11(a) is reproduced in Figure 3.12(b) for ease of comparison.) It can be seen that as I increases, the dips below the rated power become more

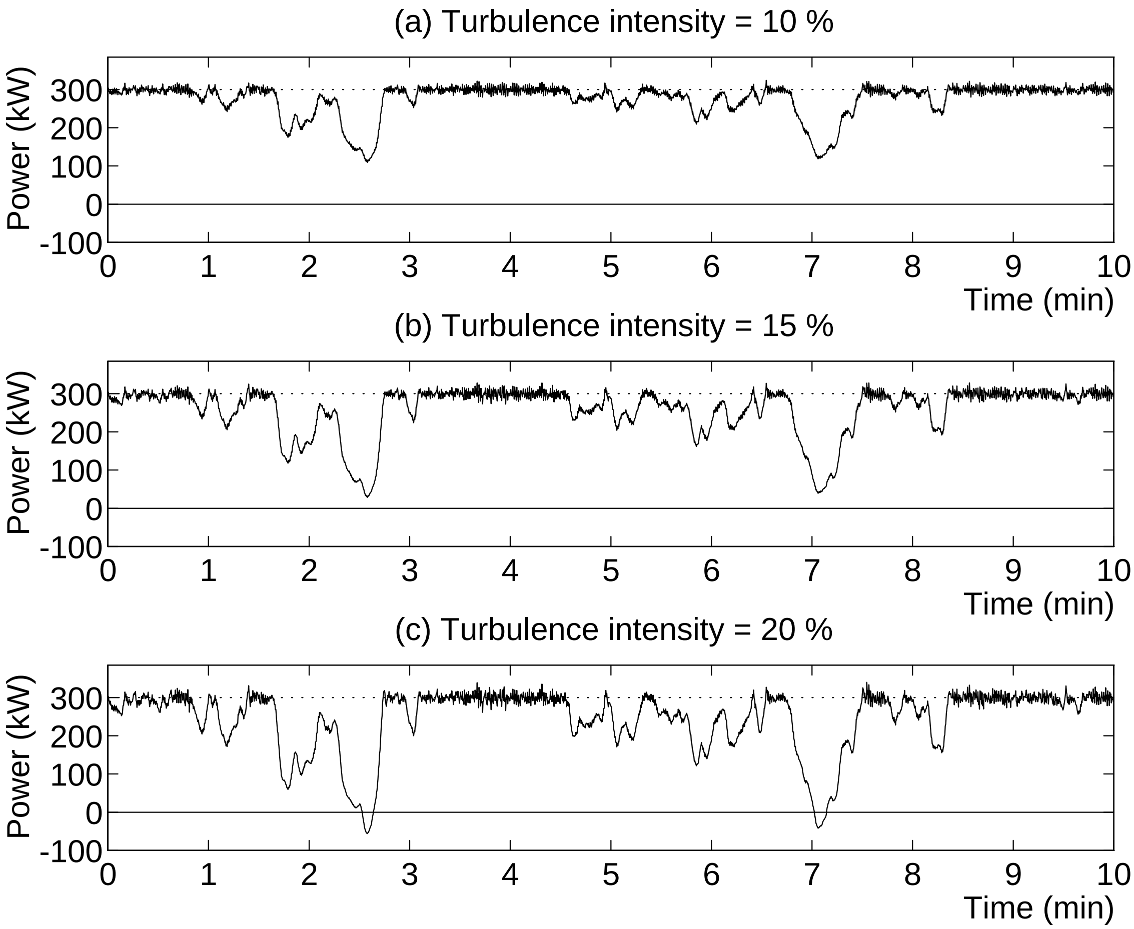

Figure 3.12: Power against time for various turbulence intensities and $\overline{v} = 11.5\ m/s$. *(1 turbine;* $P_{sp} = P_{rated} = 300\ kW$.*)*

severe. Also, the regulation at around rated power degrades somewhat. It is the effect of the dips beneath rated power that is particularly important, as this causes $\overline{P}$ to decrease and σ_p to increase. Figure 3.12 shows that, for turbulence intensities in the range of interest, the model of the single turbine again behaves sensibly and exhibits the correct basic operational characteristics in the time domain.

3.6.3 Convergence of estimated mean power

The simulation runs all provide estimates of the rate of energy capture with a specific minimum estimated degree of confidence. This is achieved via the application of confidence intervals, whereby the confidence level is monitored continuously during each simulation run and then when the confidence attains a specified threshold the estimate is considered to have converged. In addition to the confidence intervals, two heuristic rules are applied:

1. Every simulation is forced to run for a minimum of 10 minutes. It is required that the estimates should include the effects of the full range of frequencies in the atmospheric turbulence and, in the spectrum of wind-speed, atmospheric turbulence contains significant power at frequencies down to approximately 1/(10 mins) (see Figure 1.2).

2. If an estimated mean is less than 10^{-3} times the maximum possible value of the estimate across the operating range, and if more than 10 minutes have elapsed, the run is terminated. This protects against unacceptably long simulation runs under conditions where the mean is negligible.

Details of the estimation of the confidence intervals and of the heuristic rules are presented in Appendix D.

An example which shows the satisfactory convergence of the estimated mean rate of energy capture is illustrated by Figure 3.13, for $\overline{v} = 11.5$ m/s and $I = 15\%$. Figure 3.13(a) presents the instantaneous rate of energy capture and Figure 3.13(b) gives the estimated mean rate ($\hat{\overline{\dot{e}}}$). The estimated standard deviation of the estimated mean ($\hat{\sigma}_{\overline{\dot{e}}}$) is presented in Figure 3.13(c), indicating that the long-term trend is towards a reduction in the variability of the estimated mean. The convergence

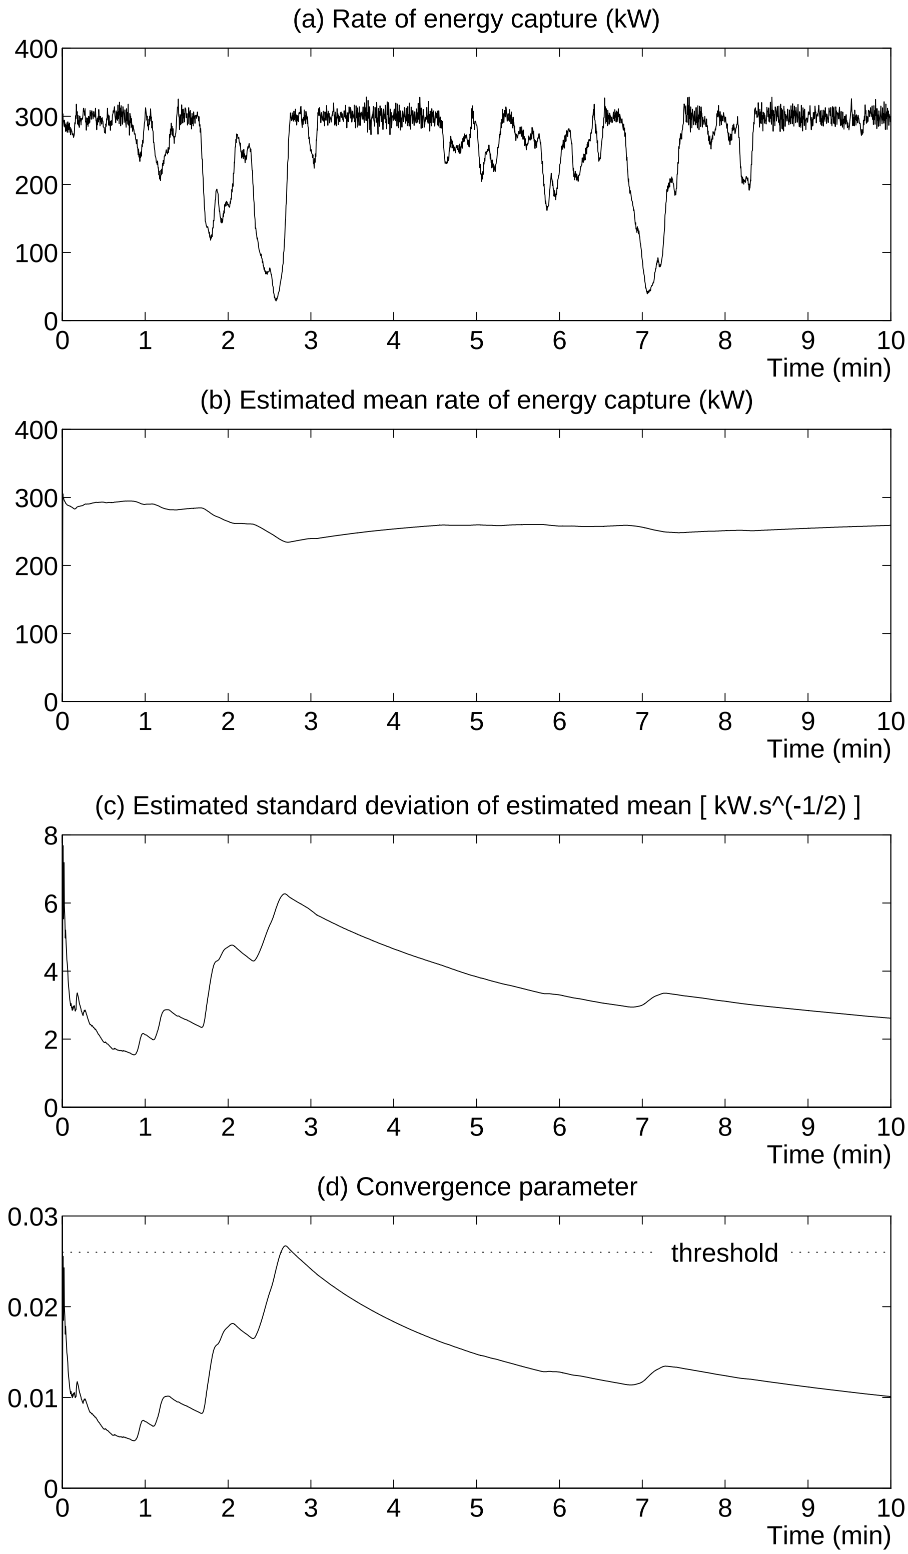

Figure 3.13: Convergence of the estimated mean rate of energy capture. (1 turbine; $P_{sp} = P_{rated} = 300$ kW; $\overline{v} = 11.5$ m/s; $I = 15\%$.)

parameter a_{conv} is defined as $\hat{\sigma}_{\hat{\bar{\dot{e}}}} / \hat{\bar{\dot{e}}}$ (see Equations D.2 and D.3) and a_{conv} is plotted in Figure 3.13(d). The confidence in the estimate increases rapidly in the first minute, then decreases significantly during the following 2 minutes, and finally increases steadily thereafter. The drop in confidence for 1 min $< t <$ 3 mins is due to the movement of the mean away from a relatively stable plateau, as Figure 3.13(b) shows. This change in the mean is itself caused by large dips in power beneath the rated power for 1 min $< t <$ 3 mins, as Figure 3.13(a) clearly indicates.

The importance of the first heuristic rule is illustrated by Figure 3.13(d). The threshold marked on the graph represents an estimated 95% confidence of the estimated mean being within +/- 5% of the true mean. (This threshold is applied to the mean rate of energy capture throughout all experiments.) The confidence interval is satisfied when a_{conv} falls beneath the threshold. This occurs after a very short period of time in this simulation, as Figure 3.13(d) indicates. The estimated mean at the start of the simulation is close to 300 kW, whereas after 10 minutes it is 259 kW, or approximately 14% lower. The performance of any windfarm is highly sensitive to the total energy capture, and errors in estimates of the mean power of larger than the order of 1% would be unacceptable for use in a supervisory controller. The initial satisfaction of the estimated confidence interval after a short period of time shows that $\dot{e}(t)$ can be apparently stationary over the short term. However, the estimate $\hat{\bar{\dot{e}}}$ that is actually required is the longer-term mean that is caused by fluctuations due to atmospheric turbulence. This longer-term mean can be found by forcing the estimation to continue for at least 10 minutes.

The calculation of the confidence intervals, which are used for the evaluation of the threshold, makes the assumption that the samples used to estimate the mean have a normal distribution. The samples of $\dot{e}$ are not always normally distributed, as can be observed in Figure 3.13(a). However, the first heuristic rule forces the estimation to continue for long enough such that a large number of independent samples are used in the estimation. The samples of $\dot{e}$ can be considered as a number of random processes $X_1, \cdots, X_n$, each with the same probability distribution. The mean of the random process $X = \frac{1}{n}(X_1 + \cdots + X_n)$ has the same value as the mean of $\dot{e}$ and, as n is large, the Central Limit Theorem states that the distribution of X is asymptotically normal.

The samples of $\dot{e}(t)$ which are used in the estimation are taken at every integration time-step of the simulation. The maximum time-step is 0.003 seconds, while the dominant time constant in the forward path of the turbine's dynamics is 9 seconds. The filtering effect of the turbine's dynamics means that the samples of $\dot{e}(t)$ are not independent of one another. An investigation was carried out to check whether, over a 10 minute period, a large number of independent samples are used in the estimation. 100 samples of $\dot{e}(t)$ were recorded at 6 second intervals for $\overline{v} = 11.5$ m/s and $I = 15\%$, and a run test (Bendat and Piersol, 1986)(pp. 95-7) was applied to the samples. The run test (with a typical level of significance of 0.02) indicates that the 100 samples are independent. Therefore a sufficiently large number of independent samples are used in the estimation for the samples to be assumed to be normally distributed.

3.7 Variation with set point for single wind turbine

3.7.1 Mean power

Figure 3.14 presents results for the mean rate of energy capture versus wind-speed and power set point, for a turbulence intensity of 15%. Site data is not available for the validation of these results. However, it can be seen that the mean power varies

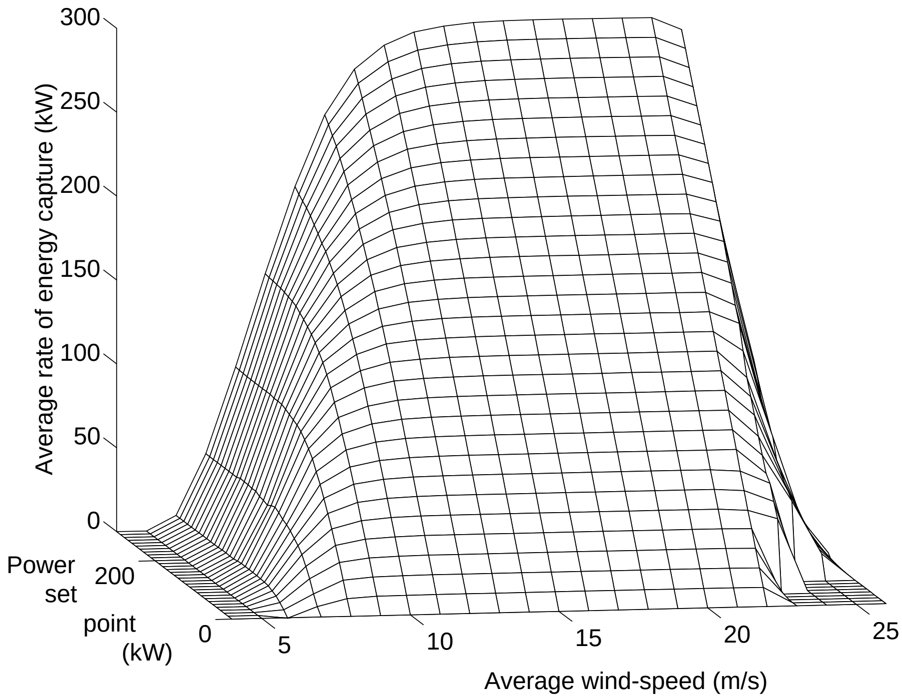

Figure 3.14: Mean rate of energy capture for 1 turbine. (I = 15%)

as a function of power set point in a logical manner.

The curve along the back of the plot, for which $P_{sp} = P_{rated}$, is the power curve for standard industrial practice (see Figure 3.8). From this standard power curve it can be seen that, in the wind-speed range 14 m/s $< \overline{v} <$ 22 m/s, there is sufficient power in the wind to maintain the average power output at rated power. Therefore if the set point is set to a level below rated power, the average power output must remain equal to the set point. Figure 3.14 indicates that this is indeed the case, as the average power output falls linearly to zero as a function of set point in the range 14 m/s $< \overline{v} <$ 22 m/s.

For similar reasons, as the power set point is increased from 0 kW when the mean wind-speed is 6 m/s or 7 m/s, the curve of mean power output versus set point is linear. In these low wind-speeds, however, there is insufficient power in the wind to achieve rated power. Therefore when the power set point is close to rated power, the pitch angle of the blades is constant at 0° to maximise power output, and changes in the set point cause no change in mean power output. This effect is seen at 6 m/s and 7 m/s in Figure 3.14, when reductions in set point from the rated power initially cause no change in the mean power output.

As the wind-speed increases from 7 m/s to 14 m/s, the curves of mean power versus power set point move through a transitional phase, until they become linear functions by approximately 14 m/s.

It can be seen in Figure 3.14 that the mean power output is equal to zero at some points in the region 23 m/s $\leq \overline{v} \leq$ 26 m/s and 0 kW $\leq P_{sp} \leq$ 140 kW. Simulation runs were not carried out at these points because of the prohibitive run-times that would have been necessary to satisfy the confidence interval, and the mean power output is instead set to zero.

3.7.2 Cumulative energy capture

A clearer picture of the levels of energy capture that a turbine is capable of generating as a function of wind-speed and power set point can be seen by considering the function $p(\overline{v}) \times \hat{\overline{e}}(\overline{v}, P_{sp})$. This function is plotted in Figure 3.15, and represents the cumulative average rate of energy capture over the time period represented by the Weibull distribution, that is, over one year. The results are normalised to the value of energy capture obtained for $\overline{v} = v_{rated} = 11.5$ m/s and $P_{sp} = P_{rated} = 300$ kW. The turbulence intensity is again equal to 15% and the Weibull parameters of $\overline{v}_{year} = 7.6$ m/s and $k_w = 1.9$ are those that are given in Section 3.4 for one tariff season per annum. From Figure 3.15 it can be seen that, operationally, mid-range wind-speeds are of prime importance, with low and high wind-speeds contributing relatively small amounts to the total energy capture. In addition, it can be seen that

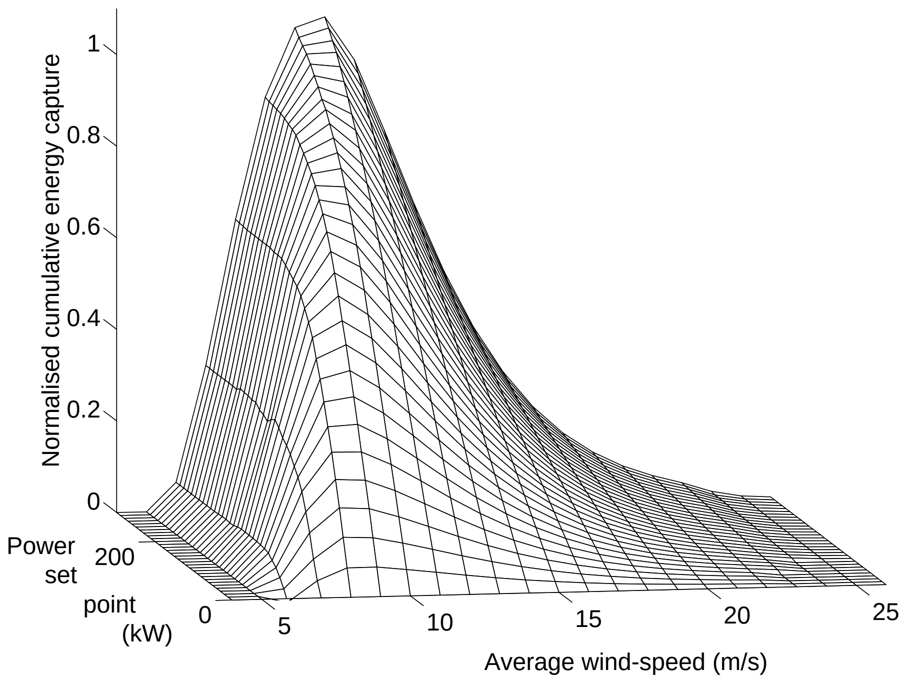

Figure 3.15: Normalised annual energy capture for 1 turbine. ($\overline{v}_{year}$ = *7.6 m/s;* k_w = *1.9;* I = *15%)*

the points in the region 23 m/s $\leq \overline{v} \leq$ 26 m/s and 0 kW $\leq P_{sp} \leq$ 140 kW make a negligible contribution to the annual energy capture, and therefore the effect of setting $\hat{\dot{e}}$ to zero in this region is also negligible.

3.8 Validation of simulation results for wake and wake interactions

The results presented in this section illustrate the means and standard deviations of the wind-speeds at turbines 1, 2 and 3, for both the 'large' (7.9D/7D) and 'small' (3D/3D) inter-turbine spacings. See Figure 3.1 (page 58) for the numbering and layouts of the turbines. For ease of comparison, all the results are for $\overline{v}$ = 16 m/s and I = 15%. All results are for 1 hour of simulated operation.

Inter-turbine spacings (1→2; 2→3):			7.9D 7D
Average free wind-speed requested:			16 m/s
Free wind turbulence intensity requested:			15%
Resulting $\sigma_{v\infty}$:			2.40 m/s
Length of simulation:			1 hour
	Turbine 1	Turbine 2	Turbine 3
$\overline{v}$ (m/s)	15.9	16.0	16.0
σ_v (m/s)	2.43	2.43	2.16
I (%)	15.2	15.2	13.6
σ_{vlf} (m/s)	—	1.44	1.08
σ_{vhf} (m/s)	—	1.89	1.85

Table 3.2: Wind-speed statistics for all turbines off-line and 'large' inter-turbine spacings

3.8.1 Results for 'large' inter-turbine spacings

Table 3.2 presents results for the 'large' (7.9D/7D) inter-turbine spacings. $\sigma_{v\infty}$ is the standard deviation of the free stream wind-speed; σ_{vlf} is the standard deviation of the low frequency component of the wind-speed ($v_{lf}^{(j)}$) that is output by the wake model; σ_{vhf} is the standard deviation of the high frequency component of the wind-speed ($v_{hf}^{(j)}$). (See Figure 3.5 for $v_{lf}^{(j)}$ and Figure 3.7 for $v_{hf}^{(j)}$.) The turbines are all off-line in Table 3.2 and therefore the results illustrate the model's representation of the flow of the free stream wind from point to point.

$\overline{v}^{(2)}$ and $\overline{v}^{(3)}$ are both satisfactory, falling within 1% of $\overline{v}^{(1)}$, and $\sigma_v^{(2)}$ is equal to $\sigma_v^{(2)}$ (to 3 s.f.). Therefore, for this layout and for these incident wind conditions, the

Inter-turbine spacings (1→2; 2→3):			7.9D 7D
Average free wind-speed requested:			16 m/s
Free wind turbulence intensity requested:			15%
Resulting $\sigma_{v\infty}$:			2.40 m/s
Length of simulation:			1 hour
	Turbine 1	Turbine 2	Turbine 3
$\overline{v}$ (m/s)	15.9 (=)	15.8 (-0.3)	15.5 (-0.5)
σ_v (m/s)	2.43 (=)	2.52 (+0.09)	2.31 (+0.15)
I (%)	15.2 (=)	15.8 (+0.6)	14.5 (+0.9)
σ_{vlf} (m/s)	—	1.48 (+0.04)	1.14 (+0.06)
σ_{vhf} (m/s)	—	1.98 (+0.09)	1.99 (+0.14)
I_{target} (%)	—	15.7 (+0.7)	16.1 (+1.1)

(The values in brackets give the change relative to the values obtained for all turbines off-line.)

Table 3.3: Wind-speed statistics for all turbines on-line and 'large' inter-turbine spacings. ($P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated}$)

model of the wake of a single turbine operates correctly. However, $\sigma_v^{(3)}$ is 11% lower than $\sigma_v^{(1)}$, which shows that significant errors are introduced in the variance of the wind-speed by the model of overlapping wakes.

Table 3.3 presents results for the same inter-turbine spacings but for all turbines on-line, with the power set points all equal to rated power. The values in brackets in the table give the change in each variable relative to the values that are obtained for all turbines off-line. As before, the values for $\overline{v}^{(2)}$ and $\overline{v}^{(3)}$ are reasonable. $\overline{v}$ falls at turbine 2 due to the wake of turbine 1 and falls further at turbine 3 due to the combined wakes of turbines 1 and 2.

I_{target} is introduced in Table 3.3 to assist in the validation of $I^{(2)}$ and $I^{(3)}$. It provides a reliable indicator of the turbulence intensity that the model should give for the downstream turbines. I_{target} is calculated by taking the mean value of the turbulence intensity that is predicted will be experienced in the wake of an operational turbine by the general empirical model introduced in Section 3.5.3. Although I_{target} is used in the course of the production of the high-frequency component of wind-speed, this high-frequency component is passed through a high-pass filter and is then added to the low-frequency component. Consequently the turbulence intensity of the simulated wind-speed at the downstream turbines may not equal I_{target}.

$I^{(2)}$ differs from $I^{(2)}_{target}$ by only 1%, which indicates that the model of a single wake is reliable for the turbines on-line, as was the case previously for the turbines off-line. $I^{(3)}$ is 10% lower than $I^{(3)}_{target}$, however, thereby again indicating that the wake superposition model introduces significant errors in $\sigma_v^{(3)}$.

The changes in the variables from the off-line to the on-line case, as indicated by the values in brackets in Table 3.3, are all reasonable. This includes $I^{(3)}$, which increases by 0.9% while I_{target} increases by 1.1%.

In general, the operation of the wake model with the 'large' spacings is satisfactory, with the important exception of the standard deviation of wind-speed at the turbine which stands within overlapping wakes, turbine number 3. A correction for the error in $\sigma_v^{(3)}$ is introduced later, in Section 3.8.3.

Inter-turbine spacings (1→2; 2→3):			3D 3D
Average free wind-speed requested:			16 m/s
Free wind turbulence intensity requested:			15%
Resulting $\sigma_{v\infty}$:			2.40 m/s
Length of simulation:			1 hour
	Turbine 1	Turbine 2	Turbine 3
$\overline{v}$ (m/s)	15.9	16.1	16.1
σ_v (m/s)	2.43	2.43	2.13
I (%)	15.2	15.2	13.3
σ_{vlf} (m/s)	—	1.91	1.60
σ_{vhf} (m/s)	—	1.46	1.46

Table 3.4: Wind-speed statistics for all turbines off-line and 'small' inter-turbine spacings

3.8.2 Results for 'small' inter-turbine spacings

Table 3.4 presents results for all turbines off-line and for the 'small' (3D/3D) inter-turbine spacings. All values for $\overline{v}^{(i)}$, ${\sigma_v}^{(i)}$ and $I^{(i)}$ are similar to those for the larger spacings, as can be seen by comparison with Table 3.2. However, the values of $\sigma_{vlf}^{(i)}$ and $\sigma_{vhf}^{(i)}$ are significantly different in this case, illustrating the effect that the change in the inter-turbine spacings has, via the normalised frequency η, on the frequency ranges of the low and high frequency signals in the wake model. $\sigma_v^{(3)}$ (= 2.13 m/s) has a similar value to that obtained for the large spacings with all turbines off-line (2.16 m/s), thereby indicating that the error in $\sigma_v^{(3)}$ is approximately independent of the inter-turbine spacings, for the range of spacings over which the simulation is used.

The cause of the underestimation of $\sigma_v^{(3)}$ has been investigated by examining the statistics of the simulated low and high frequency components of wind-speed at turbine 3, for all turbines off-line and with $\overline{v} = 16$ m/s and $I = 15\%$.

For the low frequency component of wind-speed, the velocity deficits due to turbines 1 and 2 ($\Delta v_{lf}^{(1)}$ and $\Delta v_{lf}^{(2)}$ in Equation 3.3) were both found to be zero, as would be expected when both turbines are switched off. Consequently $\sigma_{vlf}^{(3)}$ is equal to the standard deviation of the filtered, delayed, free stream wind-speed ($v_{lf\infty}$ in Equation 3.3). The cut-off frequency (f_{lp}) of the low-pass filter that produces $v_{lf\infty}$ is calculated using Equation 3.2, with Δx equal to the spacing between turbines 1 and 3 (6 rotor diameters).

For the high frequency component of wind-speed, the standard deviations of wind-speed due to turbines 1 and 2 ($\sigma_{vc}^{(1)}$ and $\sigma_{vc}^{(2)}$ in Equation 3.4) were both found to be zero, which is again as would be expected when both turbines are switched off. Consequently $\sigma_{vinc}^{(3)} = \sigma_{v\infty} = 2.40$ m/s. The output of the standard wind model was found to have a mean of 0.04 m/s and a standard deviation of 2.42 m/s, both of which are satisfactory. However, the cut-off frequency (f_{hp}) of the high-pass filter that produces $v_{hf}^{(3)}$ is calculated on the basis that turbine 2, when on-line, dominates the turbulence experienced by turbine 3. Hence, in Equation 3.2, Δx is equal to the spacing between turbines 2 and 3 (3 rotor diameters). Consequently $f_{hp} = 2f_{lp}$, and the high frequency component of wind-speed produced at turbine 3 due to turbine 1 does not include fluctuations with frequencies in the range $f_{lp} < f < f_{hp}$. As a result the standard deviation of the high-frequency component of wind-speed at turbine 3 is underestimated, and hence $\sigma_v^{(3)}$ is underestimated.

Table 3.5 presents results for all turbines on-line. These results can be compared with those for the larger spacings in Table 3.3. It can then be seen that the wake effects

Inter-turbine spacings (1→2; 2→3):		3D 3D	
Average free wind-speed requested:		16 m/s	
Free wind turbulence intensity requested:		15%	
Resulting $\sigma_{v\infty}$:		2.40 m/s	
Length of simulation:		1 hour	
	Turbine 1	Turbine 2	Turbine 3
$\overline{v}$ (m/s)	15.9 (=)	15.3 (-0.7)	14.8 (-1.2)
σ_v (m/s)	2.43 (=)	2.69 (+0.26)	2.57 (+0.43)
I (%)	15.2 (=)	16.8 (+1.6)	16.0 (+2.7)
σ_{vlf} (m/s)	—	2.05 (+0.14)	1.84 (+0.25)
σ_{vhf} (m/s)	—	1.68 (+0.22)	1.84 (+0.38)
I_{target} (%)	—	17.1 (+2.1)	18.6 (+3.6)

(The values in brackets give the change relative to the values obtained for all turbines off-line.)

Table 3.5: Wind-speed statistics for all turbines on-line and 'small' inter-turbine spacings. ($P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated}$)

for the smaller spacings are, correctly, accentuated, with larger velocity deficits and higher levels of turbulence intensity at the downstream turbines. $\sigma_v^{(3)}$ is still in error, again giving rise to a considerably lower value for $I^{(3)}$ (16.0%) than I_{target} indicates should be found (18.6%). This effect is consistently seen in simulations carried out in other wind conditions as well.

3.8.3 Correction of $\sigma_v^{(2)}$ and $\sigma_v^{(3)}$

The underestimation of $\sigma_v^{(3)}$ means that, without any alterations, the test-bed simulation would yield conservative improvements in performance for a supervisory controller, because the simulated operating conditions at turbine number 3 underestimate the true effect of the operation of the upstream turbines. Corrections are calculated off-line for $\sigma_v^{(3)}$ and, as the error in $\sigma_v^{(3)}$ changes little with the inter-turbine spacings, the corrections are calculated as functions of $\overline{v}$ and I only. Corrections are also calculated for $\sigma_v^{(2)}$, because over simulated periods of more than 1 hour, and under various wind conditions, $\sigma_v^{(2)}$ also moves away from its target value slightly.

The errors in $I^{(2)}$ and $I^{(3)}$ are not simple biases because, as Table 3.5 illustrates, the increases in $I^{(2)}$ and $I^{(3)}$ when the turbines change from the off-line to the on-line state are not as large as the increases in $I^{(2)}_{target}$ and $I^{(3)}_{target}$. Consequently a bias term is not used for the purposes of compensating for the error and a multiplicative constant is investigated instead. This multiplier is applied to the σ_v demand that is input to the wind model at the downstream turbine. The correction is calculated for specific pairs of $\overline{v}$ and I and for all turbines turned off. When calculating the correction term, it is assumed that the low frequency signal is fixed. For the upstream turbine or turbines switched off, the target value of the variance of the high-frequency signal at downstream turbine i $[(\sigma_0^{(i)})^2]$ is then calculated as follows:

$$\left(\sigma_0^{(i)}\right)^2 = \sigma_{v\infty}^2 - \left(\sigma_{vlfoff}^{(i)}\right)^2$$

where $\sigma_{vlfoff}^{(i)}$ is the standard deviation of the low frequency component of the windspeed at downstream turbine i, for the upstream turbine or turbines switched off. If $a_\sigma^{(i)}$ is the parameter of multiplication for $\sigma_v^{(i)}$, then:

$$\left(a_{\sigma}^{(i)}\sigma_{vhfoff}^{(i)}\right)^2 = \sigma_{v\infty}^2 - \left(\sigma_{vlfoff}^{(i)}\right)^2$$

Therefore:

$$a_{\sigma}^{(i)} = \sqrt{\frac{\sigma_{v\infty}^2 - \left(\sigma_{vlfoff}^{(i)}\right)^2}{\left(\sigma_{vhfoff}^{(i)}\right)^2}}$$

Values for $a_{\sigma}^{(i)}$ were calculated for all combinations of the following values of $\overline{v}$ and I:

$$\overline{v} = 8 \text{ m/s}, 10 \text{ m/s}, 12 \text{ m/s}, \cdots, 24 \text{ m/s} \qquad \text{and} \qquad I = 10\%, 15\%, 20\%$$

Table 3.6 presents the resulting wind-speed statistics when the corrections are applied to $\sigma_v^{(2)}$ and $\sigma_v^{(3)}$, for all turbines off-line. These results can be compared with those for the same conditions but with no corrections in Table 3.4. It can be seen that $\sigma_{vhf}^{(2)}$ increases slightly, from 1.46 m/s to 1.50 m/s, and so $\sigma_v^{(2)}$ moves away from its ideal value ($\sigma_v^{(1)}$) of 2.43 m/s, to 2.46 m/s. The reason for the change in $\sigma_v^{(2)}$ is that the corrections are calculated using more than one hour of simulated operation, and hence are slightly less accurate for 1 hour of operation. $\sigma_{vhf}^{(3)}$ increases substantially, from 1.46 m/s to 1.81 m/s, and $\sigma_v^{(3)}$ also increases, from 2.13 m/s to 2.38 m/s. The latter figure is relatively close to the target value for $\sigma_v^{(3)}$ of 2.43 m/s and it can be concluded that, for the purposes of increasing $\sigma_v^{(3)}$, the application of the correction parameter is successful for the case of all turbines off-line.

Table 3.7 presents results for all the turbines on-line. These results can be compared with those for the same conditions but with no corrections applied to $\sigma_v^{(2)}$ and $\sigma_v^{(3)}$ in Table 3.5. $\sigma_{vhf}^{(2)}$ and $\sigma_v^{(2)}$ increase slightly and consequently $I^{(2)}$ moves closer towards its target value $I_{target}^{(2)}$ of 17.1%. $\sigma_{vhf}^{(3)}$ and $\sigma_v^{(3)}$ increase substantially and $I^{(3)}$ rises to 18.0%, which is much closer to its target value of 18.6%. Therefore the application

Inter-turbine spacings (1→2; 2→3):			3D 3D
Average free wind-speed requested:			16 m/s
Free wind turbulence intensity requested:			15%
Resulting $\sigma_{v\infty}$:			2.40 m/s
Length of simulation:			1 hour
	Turbine 1	Turbine 2	Turbine 3
$\overline{v}$ (m/s)	15.9	16.1	16.1
σ_v (m/s)	2.43	2.46	2.38
I (%)	15.2	15.3	14.9
σ_{vlf} (m/s)	—	1.91	1.60
σ_{vhf} (m/s)	—	1.50	1.81

Table 3.6: Wind-speed statistics for all turbines off-line with corrections applied to $\sigma_v^{(2)}$ and $\sigma_v^{(3)}$

of the correction parameter is also successful for all turbines on-line.

As a function of power set point, the additional turbulence in a wake is at its maximum when the power set point is also at its maximum. Therefore the results in Table 3.7 show that the use of a multiplier correction term which is based upon values of σ_v for $P_{sp} = 0$ does not cause $I^{(3)}$ to exceed $I^{(3)}_{target}$ when P_{sp} is at its maximum value of P_{rated}. Checks were carried out which verified that $I^{(3)}$ does not exceed $I^{(3)}_{target}$ at any of the extremities of the range of wind conditions for which the corrections were calculated, that is, for $\overline{v} = 8$ m/s and 24 m/s and for $I = 10\%$ and 20%. Therefore it is likely that the use of a multiplicative constant for the correction of $\sigma_v^{(2)}$ and $\sigma_v^{(3)}$ does not cause excessive increases in turbulence intensity at turbine number 3 under any wind conditions that the simulation is required to be capable of modelling.

Inter-turbine spacings (1→2; 2→3):			3D 3D
Average free wind-speed requested:			16 m/s
Free wind turbulence intensity requested:			15%
Resulting $\sigma_{v\infty}$:			2.40 m/s
Length of simulation:			1 hour
	Turbine 1	Turbine 2	Turbine 3
$\overline{v}$ (m/s)	15.9 (=)	15.3 (-0.7)	14.8 (-1.2)
σ_v (m/s)	2.43 (=)	2.71 (+0.26)	2.89 (+0.51)
I (%)	15.2 (=)	16.9 (+1.6)	18.0 (+3.1)
σ_{vlf} (m/s)	—	2.05 (+0.14)	1.84 (+0.25)
σ_{vhf} (m/s)	—	1.72 (+0.22)	2.27 (+0.47)
I_{target} (%)	—	17.1 (+2.1)	18.6 (+3.6)

(The values in brackets give the change relative to the values obtained with all turbines off-line.)

Table 3.7: Wind-speed statistics for all turbines on-line and corrections applied to $\sigma_v^{(2)}$ and $\sigma_v^{(3)}$. ($P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated}$)

3.9 Conclusions

This chapter has described those factors that must be considered when simulating windfarms for the purposes of testing supervisory controllers. A test-bed simulation of three turbines in a row has been presented and it has been shown that the windfarm model is dynamically 'stiff'. Models have been described for the turbulent wind and for the long-term probability distribution of wind-speed.

Results have been presented for the simulated mean rate of energy capture and for the long-term cumulative energy capture, as functions of mean wind-speed and power set point. The results for the long-term cumulative energy capture show that the mid-range wind-speeds are operationally of prime importance, with the low and high wind-speeds contributing relatively small amounts to the total energy capture.

It is concluded that wake effects must be included in the test-bed windfarm simulation. This is because, as well as affecting mean wind-speeds and turbulence intensities at the downstream turbines, wakes create dependency between the dynamics of the turbines and therefore affect the estimation and control that is carried out by a supervisory controller. Topographic effects, however, can be neglected and the simulation remains a suitable test-bed for supervisory control algorithms.

It has been shown that the wake and wake interaction models can be split up into two components, a low frequency part and a high frequency part. Results have been presented which illustrate that the resulting model of a single wake operates satisfactorily. The wake interaction model, however, underestimates the turbulence intensity at the turbine that stands within overlapping wakes, and a correction parameter has been introduced to compensate for this error. The use of the correction parameter only amplifies the part of the spectrum of the wind-speed that is represented by the high frequency component, and does not affect the low frequency component. It has been shown that the correction parameter achieves the desired objective of increasing the turbulence intensity at the turbine that stands within overlapping wakes so that it is satisfactorily close to its target value, where the target value is evaluated using an empirical model.

It can be concluded that the resulting windfarm simulation provides a suitable test-bed for supervisory controllers.

Chapter 4

Models of wind turbine fatigue damage

4.1 Introduction

This chapter describes the models of wind turbine fatigue damage that are used in the windfarm simulation, and presents results from the various models.

In Section 4.2, descriptions of the models are given. The gearbox is the single most expensive component in the design of turbine that is modelled in the windfarm simulation, and models of fatigue damage for the gearbox bearings and teeth are described in Section 4.2.1. Further models of fatigue damage are described in Section 4.2.2 for providing preliminary indications of the rates of fatigue damage experienced by structural components within the turbine. The use of the standard deviation of torque or power as a generic term for modelling fatigue damage is described in Section 4.2.3.

Simulation results of the estimated damage of a single turbine are presented in Section 4.3, and results for the 3-turbine windfarm are presented in Section 4.4 for both 'large' (7.9D/7D) and 'small' (3D/3D) inter-turbine spacings.

4.2 Estimation of fatigue damage

Fatigue damage is estimated using Miner's rule in all of the models (Miner, 1945), with the exception of the method described in Section 4.2.3 where the standard deviation of power is used to represent fatigue damage.

To apply Miner's rule to a particular material, the material's curve of stress amplitude versus cycles to failure (S–N curve) is required. Two typical S–N curves, both approximated with straight line segments and plotted on log-log axes, are illustrated in Figure 4.1 (Merritt, 1971). Curve (a) is for a non-ferrous metal and curve (b)

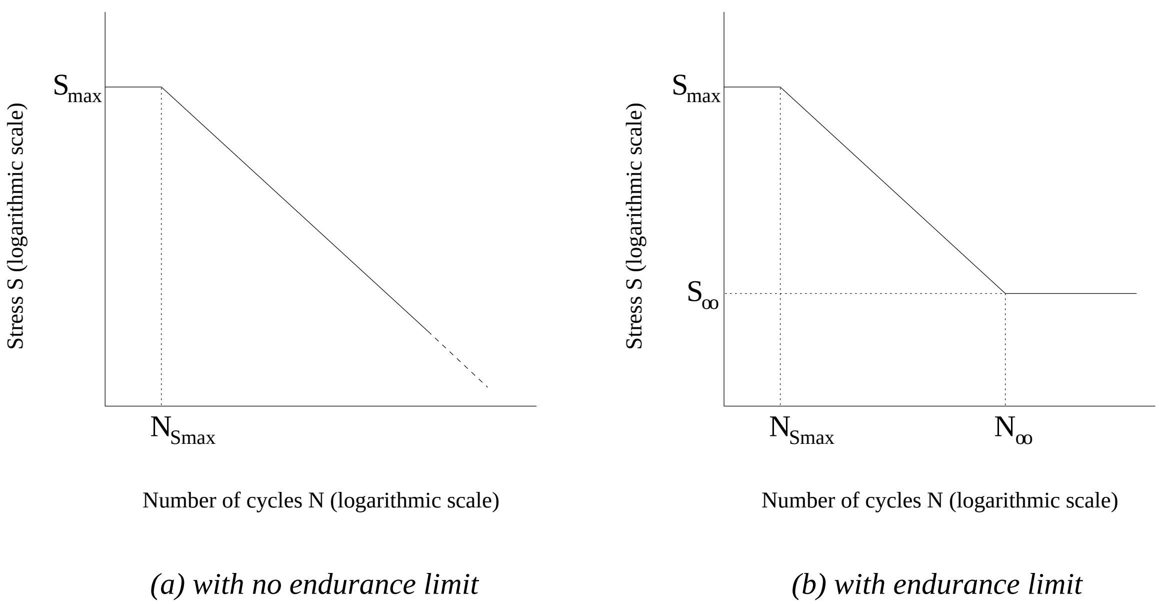

Figure 4.1: Typical S–N curves

is for a steel. Curve (b) illustrates that the curve for steels has a lower 'knee', or endurance limit, at S_∞. Stress amplitudes which are below this limit do not cause failure, and are not counted in the Miner's sum.

For the purposes of designing and testing supervisory controllers, time averages of the rate of fatigue damage $\bar{\hat{\dot{d}}}$ are required. The combination of Miner's sum with a linear representation of the $logS$–$logN$ curve can be used to give the following relationship (Merritt, 1971):

$$\bar{\hat{\dot{d}}} = \frac{1}{T}\sum_{i=1}^{n_c} S_i^b \tag{4.1}$$

where n_c is the number of stress cycles that occur in time T; S_i is the amplitude of stress cycle i; b is the slope index of the material. The slope index is equal to the negative reciprocal of the slope of the $logS$–$logN$ curve.

4.2.1 Gearbox

The design of gearbox upon which the fatigue damage model is based is typical of designs used in wind turbines. It steps up the rotational speed from 48 rpm to 1500 rpm, and contains two planetary stages followed by a parallel stage.

Bearings

It is assumed that the fatigue life of the bearings is more critical than the fatigue life of the gear teeth (Hicks, 1992). In the gearbox under consideration, roller bearings are used. These have no endurance limit and hence the shape of the S–N curve is as shown in Figure 4.1(a) (Merritt, 1971). The slope index b for roller bearings is 3.33.

The fatigue damage of the bearing at the rotor end of the low-speed shaft (the 'front bearing') is modelled, because this bearing has the shortest predicted fatigue life of all the bearings in the design of gearbox under consideration (Compact Orbital Gears, 1991). Only the fatigue loads caused by the aerodynamic rotor thrust are considered. An analysis has been carried out of the relative magnitude of the thrust loadings on the low-speed shaft's bearings due to aerodynamic thrust and to the weight of the turbine's rotor. The aerodynamic loads on the front bearing account for 40% or more of the total loadings throughout the range 8 m/s $< \overline{v} <$ 24 m/s, for the case where the power set point is equal to rated power. At wind-speeds around $\overline{v} = 10$ m/s, the aerodynamic loads rise to 60% of the total loadings.

The fatigue loads due to aerodynamic thrust on the front bearing are proportional to the thrust, for all operation except when the aerodynamic thrust is small. Further, aerodynamic thrust is approximately proportional to low-speed shaft torque, as illustrated by Figure 4.2. The data presented in Figure 4.2 is generated using

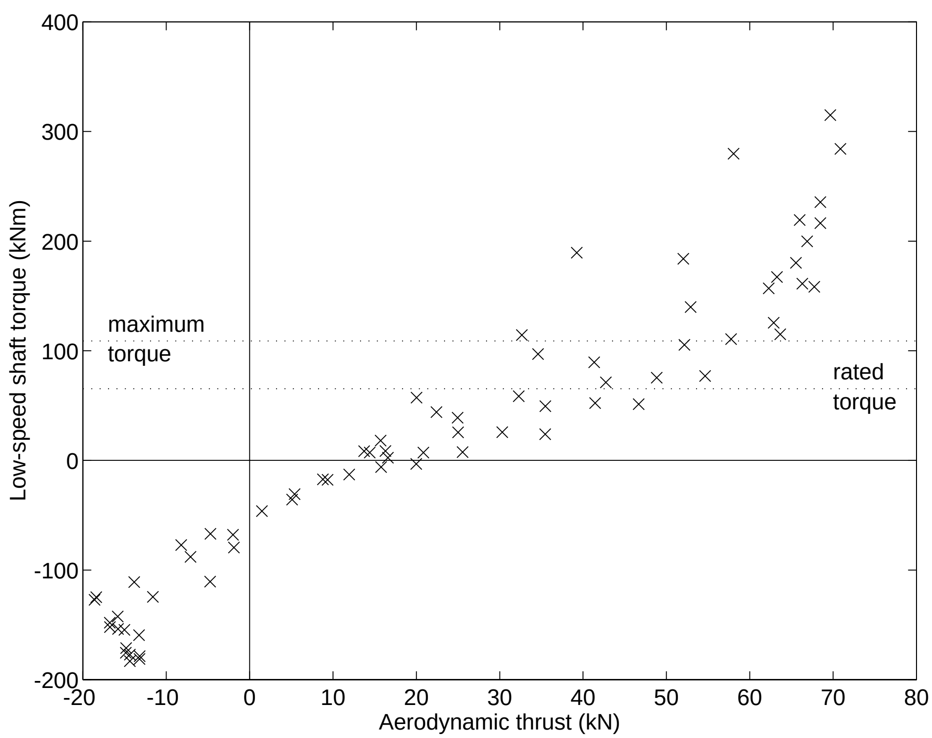

Figure 4.2: Correlation between aerodynamic thrust and low-speed shaft torque. (1 turbine; 5 m/s $< v <$ 28 m/s; $-2^o < \phi < 25^o$.)

steady-state values, given in Appendix C, for the power coefficient (C_p) and thrust coefficient (C_t). The C_p and C_t data had been generated using a proprietary software package (BLADES, 1988). Equations for aerodynamic power and thrust as functions of C_p and C_t are given in Appendix B.

All the data points from Appendix C are used in Figure 4.2, so operating points across the ranges 5 m/s $< v <$ 28 m/s and $-2^o < \phi < 25^o$ are plotted. This gives rise to some operating points which are rarely encountered in practice. High pitch angles in low wind-speeds give large negative thrusts and large negative torques. The group of points at negative torques with magnitudes greater than the rated torque will almost never be experienced in practice. Low pitch angles in high wind-speeds give torques which are considerably larger than the rated torque. A slipping clutch in the turbine's high-speed shaft limits the torque to the level marked "maximum torque" in Figure 4.2. (See Appendix B for further details of the slipping clutch.) Therefore the group of points above the maximum torque in Figure 4.2 are never experienced in practice, on the design of wind turbine that is modelled.

In this thesis, it is assumed that the bearing fatigue loads due to aerodynamic thrust are approximately proportional to the low-speed shaft torque, so that the supervisory controller can estimate the fatigue damage using the measurement of generated power alone (see Section 2.4.1). The turbine runs at constant rotational speed and hence the number of stress cycles experienced by the rollers in the bearing is proportional to time. From Equation 4.1, the mean rate of damage can be estimated as follows:

$$\hat{\bar{\dot{d}}}_b = \frac{1}{T} \int_0^T q_{lss}^{3 \cdot 33}(t)\; dt \tag{4.2}$$

where

$\hat{\dot{d}}_b$ = estimated mean rate of bearing fatigue damage

$q_{lss}(t)$ = low-speed shaft torque

Teeth

In contrast to bearings, gearbox teeth have an endurance limit (Merritt, 1971). Hence the approximate shape of the S–N curve for gearbox teeth is as shown in Figure 4.1(b).

The fatigue damage of the annulus teeth in the first planetary stage is modelled. This is because these teeth have the lowest design factor of all the teeth in the design of gearbox under consideration (Compact Orbital Gears, 1991). The low-speed shaft torque acts upon these teeth and is therefore used in the estimation of fatigue damage. As gearbox teeth have an endurance limit, torques below the limit cause no fatigue damage, and the estimated rate of teeth fatigue damage $\hat{\dot{d}}_t$ is as follows:

$$\hat{\dot{d}}_t = \frac{1}{T} \int_0^T u\left(q_{lss}(t) - q_{lss_\infty}\right)\ q_{lss}^b(t)\ dt$$

where $u(x - x_0)$ is the following step function:

$$u(x - x_0) = \begin{cases} 1 & \text{if } x > x_0 \\ 0 & \text{otherwise} \end{cases}$$

and q_{lss_∞} is the torque that corresponds to the stress S_∞ at the lower knee of the S–N curve. For the particular design of gearbox that is modelled, $q_{lss_\infty} = 67.6$ kNm, which corresponds to 340 kW (Spruce, 1988b).

Many gearboxes have nitrided steel teeth, including the design that is modelled (Compact Orbital Gears, 1991). The slope index b for bath nitrided steel is 16 (BS 436:Part 3, 1986) and so the estimated rate of damage is given by:

$$\hat{\bar{d}}_{t16} = \frac{1}{T}\int_0^T u(q_{lss}(t) - q_{lss_\infty})\; q_{lss}^{16}(t)\; dt \tag{4.3}$$

Steels which have not experienced nitriding are also commonly used in gear teeth, however, and this class of teeth is also modelled. Slope indices for four typical types of non-nitrided steel teeth are given in (BS 436:Part 3, 1986). Their values are between 5.0 and 8.5. Merritt recommends the use of a slope index of 5 for steels, as there are many sources of variability present in the fatigue actions being modelled (Merritt, 1971). For non-nitrided steel teeth, $b = 5$ is used:

$$\hat{\bar{d}}_{t5} = \frac{1}{T}\int_0^T u(q_{lss}(t) - q_{lss_\infty})\; q_{lss}^{5}(t)\; dt$$

4.2.2 Structural components

Transfer functions can, in principal, be used to relate the effective wind-speed experienced by the turbine to the structural modes, where the effective wind-speed is the free stream turbulent wind modified by the rotational sampling of the wind by the rotor. The resulting fatigue damage rates can then be estimated and all the significant structural modes must be considered (Leithead et al., 1991). However, the computational load associated with this model of structural damage will be high, and in this thesis a simple measure of structural fatigue damage is used to minimise the computational load.

It is assumed that the turbine consists of a single generalised 'component', whose loading is proportional to the low-speed shaft torque. This 'component' is assumed to have no endurance limit, as is the case in, for example, a general methodology for the evaluation of wind turbine structural damage rates which is described in (Holley and Madsen, 1985). The torque cycles that cause the fatigue damage are identified using rainflow cycle counting (Downing and Socie, 1982) and Miner's sum is then applied (see Equation 4.1):

$$\hat{\bar{d}}_s = \frac{1}{T}\sum_{i=1}^{n_c}(\Delta q_{lss_i})^b$$

where

i	= cycle number;
n_c	= number of cycles in time T;
Δq_{lss_i}	= low-speed shaft torque amplitude for the ith cycle.

Values for b of 3, 5 and 7 are used, again as in the analysis reported in (Holley and Madsen, 1985).

The generalised model of structural damage only gives a crude indication of the blade fatigue damage. Harmful loads can be absorbed by the blades and will then not appear in the low-speed shaft torque (Leithead et al., 1992). In particular, the pitching of the blades introduces low frequency fluctuations in the blade loads which make significant contributions to the fatigue damage (Quarton, 1990).

The model of generalised structural fatigue damage should provide a better indication of the fatigue loads experienced by the tower and by components in the hub, as these components experience loadings which are more likely to be proportional to the low-speed shaft torque.

4.2.3 Standard deviation of torque/power

In the literature on the control of wind turbines, there is general agreement that one of the control objectives is to minimise the variability of the drive-train torque or, equivalently for constant speed turbines, minimise the variability of the output power (de la Salle et al., 1990). The reasoning is that reduced fluctuations will give rise to lower fatigue loads and improved power quality. Results for the standard deviation of torque or power that a controller can achieve are sometimes presented, and it is assumed that the rate of fatigue damage of the drive-train is a monotonic function of the standard deviation of torque or power (de Boer et al., 1991) (Ekelund and Schmidtbauer, 1993) (Leithead et al., 1992). The standard deviation of power is also presented in this thesis, for the purposes of investigating its possible use as a generic indicator of the fatigue damage of components within the turbine.

4.3 Results for simulation of 1 turbine

4.3.1 Convergence of damage estimates

The simulation runs provide estimates of all measures of the rate of fatigue damage with a specific minimum estimated degree of confidence. The technique is the same as that applied to energy capture, using confidence intervals plus two heuristic rules. See section 3.6.3 and Appendix D for further details.

The convergence of the estimated mean rate of gearbox bearing fatigue damage is illustrated by Figure 4.3, for which $\overline{v} = 11.5$ m/s and $I = 15\%$. Plots (a)–(d) in this figure are directly comparable with plots (a)–(d) in Figure 3.13 (page 85), where the latter illustrate the convergence of the estimated mean rate of energy capture. The

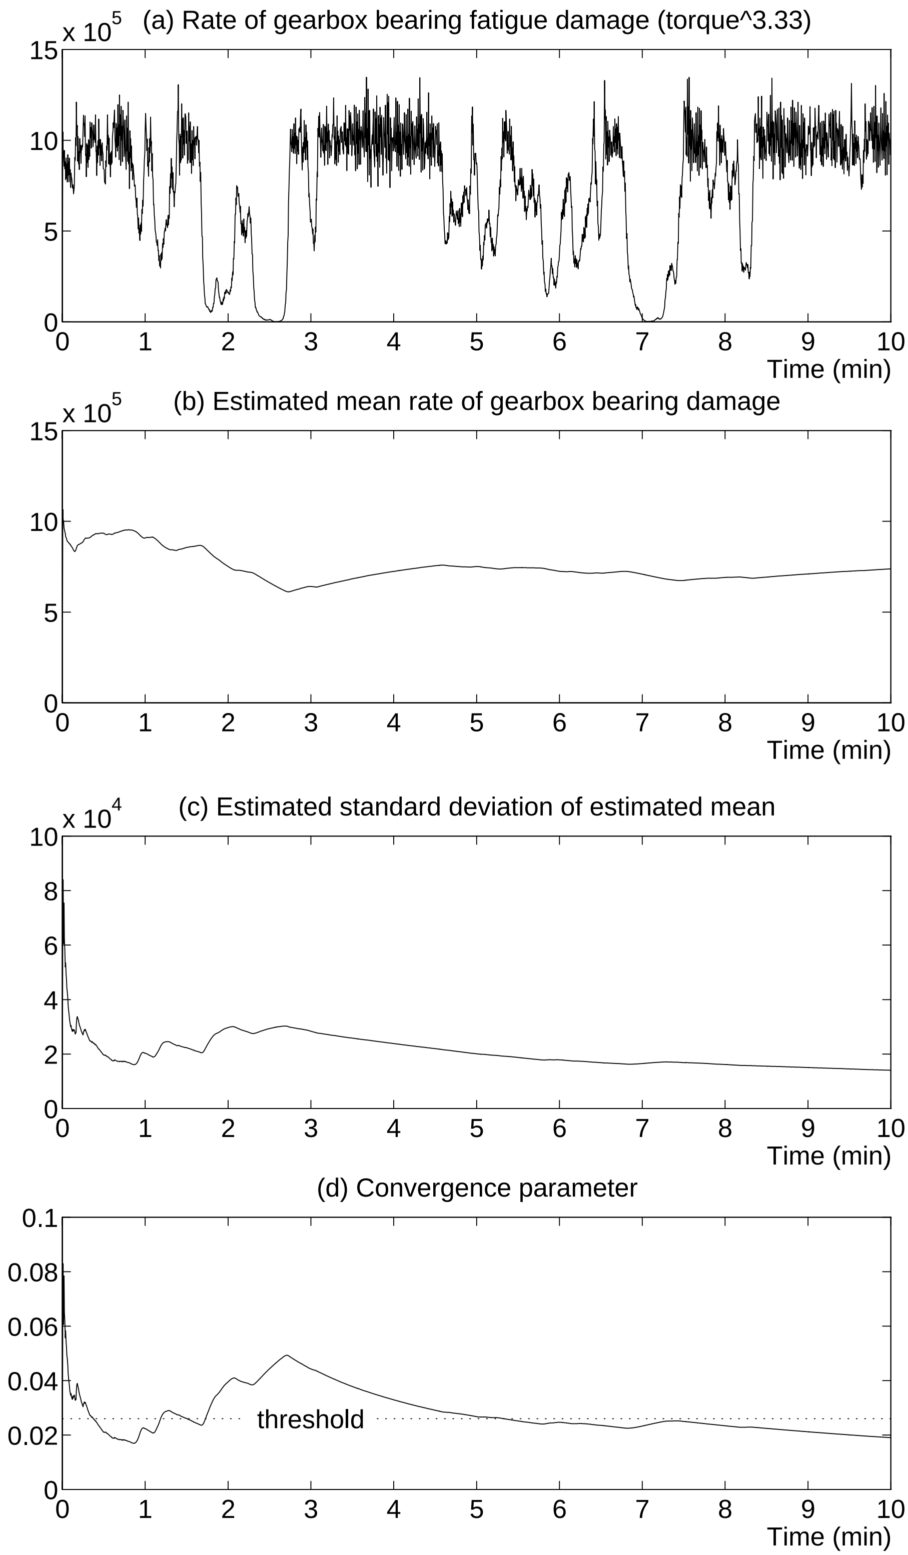

Figure 4.3: Convergence of the estimated mean rate of gearbox bearing fatigue damage. (1 turbine; $P_{sp} = P_{rated} = 300$ kW; $\overline{v} = 11.5$ m/s; $I = 15\%$.)

wind series is the same as that used to generate the results shown in Figure 3.13. As before, the threshold in Figure 4.3(d) is for a 95% confidence of the estimated mean being within +/- 5% of the true mean. Satisfactory convergence of the estimated mean rate of bearing damage is achieved, but once again it is important to force all simulations to run for a minimum period to prevent premature convergence.

The principal difference between the convergence of the rates of gearbox bearing damage $\hat{\bar{\dot{d}}}_b$ and energy capture $\hat{\bar{\dot{e}}}$ is the higher degree of variability in the bearing damage and the consequent slower convergence of $\hat{\bar{\dot{d}}}_b$. From Equation 4.2:

$$\hat{\bar{\dot{d}}}_b \propto \int_0^T q_{lss}^{3 \cdot 33}(t)\, dt$$

For energy capture, $P(t) \approx \omega_{lss}\, q_{lss}(t)$, where ω_{lss} is the low-speed shaft speed and is constant. Hence:

$$\hat{\bar{\dot{e}}} \propto \int_0^T q_{lss}(t)\, dt$$

The larger exponent of $q_{lss}(t)$ in the estimate of $\hat{\bar{\dot{d}}}_b$ causes a higher degree of variability in this estimate than in $\hat{\bar{\dot{e}}}$.

The convergence properties of the other estimates of fatigue damage are illustrated by Table 4.1. This table presents the times that are required to achieve convergence for the estimated mean rates of the following variables: gearbox bearing damage; gearbox teeth damage (with $b = 5$); general structural damage (with $b = 7$); energy capture. The convergence times for the mean rate of energy capture are included so that a complete comparison of the speed of convergence of all the estimates can be made. The wind-speeds given in the table lie in each of the five key regions across the operating range (as defined in section 3.6.2), with two wind-speeds (24 m/s and

$\overline{v}$ (m/s)	TIME TO CONVERGENCE OF ESTIMATED MEAN (minutes)			
	Rate of fatigue damage			Rate of energy capture
	Gearbox bearings ($b = 3.33$)	Gearbox teeth ($b = 5$)	General structural ($b = 7$)	
6	25	$10^{\S}$	$10^{\S}$	42
9	$10^{\dagger}$	$10^{\S}$	17	$10^{\dagger}$
11.5	$10^{\dagger}$	190	17	$10^{\dagger}$
18	$10^{\dagger}$	17	$10^{\dagger}$	$10^{\dagger}$
24	$10^{\dagger}$	21	$10^{\dagger}$	$10^{\dagger}$
26	120	2700	16	68

† run forced to continue for 10 minutes even though confidence had already achieved target.

§ run terminated because estimated mean was less than 10^{-3} times the maximum possible.

Table 4.1: Times to convergence for estimates of rates of fatigue damage and of energy capture. (1 turbine; $P_{sp} = P_{rated} = 300$ kW; $I = 15\%$.)

26 m/s) in the upper cut-out/in region.

Convergence times for the estimated rate of fatigue damage of nitrided gearbox teeth are not included in the table, because it is not possible for the confidence threshold to be attained within acceptable simulation times. The reason for the slow convergence of $\hat{\dot{d}}_{t16}$ is the high slope index (16) for nitrided teeth, which gives a highly non-linear term $[q_{lss}^{16}(t)]$ in the integrand of the estimate, as Equation 4.3 illustrates. One interpretation of the model for nitrided teeth is that, because the exponent of 16 is so large, all stress cycles do negligible damage until a single large spike is experienced of sufficient magnitude that it causes a tooth to fail completely. The unacceptably slow convergence of $\hat{\dot{d}}_{t16}$ led to results for this measure of fatigue damage not being collected.

The confidence level that is used to generate the convergence times shown in the table is a 90% confidence of the estimated mean being within +/- 10% of the true mean. This level is used throughout all experiments for the estimates of the gearbox teeth damage and the general structural damage, both of which are only used for the evaluation of the performance of the supervisory controllers. The bearing damage and energy capture estimates, which are used in the controller design, are estimated to a 95% confidence of the estimated mean being within +/- 5% of the true mean.

Table 4.1 indicates that the estimates of gearbox bearing damage, structural damage and energy capture have the best convergence properties and require similar times to achieve convergence, while the speed of convergence of the estimate of gearbox teeth damage is substantially worse.

The table also shows that the more data with a non-zero numerical value is received on average in unit time, the faster the estimate converges. The convergence at mid-range wind-speeds can be compared with the convergence at 26 m/s, for example. In the former case the turbine is on-line for all of the time and the convergence is relatively fast. In the latter case, however, the turbine is off-line for much of the time and the convergence is relatively slow. When the turbine is off-line for part of the time, there are two important effects on the estimation: firstly, significant asymmetry in the probability density of $\dot{d}$ (or $\dot{e}$) is introduced, because a spike appears in the probability density at $\dot{d} = 0$; secondly, the average amount of data available from on-line operation is reduced, and hence there is greater uncertainty about the damage during on-line operation. One result of these two factors is that the convergence worsens when the fraction of the time spent off-line increases, as the results in Table 4.1 for 24 m/s and 26 m/s indicate. This lengthening of the convergence time is not seen around the lower cut-in/out wind-speeds because of the effect of the second heuristic rule (see the times marked § in the table).

The convergence is particularly sensitive to whether the component for which the fatigue damage is to be estimated has an endurance limit or not. This sensitivity is indicated by the relatively slow convergence of the estimated rate of gearbox teeth damage under certain conditions, even with the second heuristic rule applied. Estimation is more difficult for components which have an endurance limit for the same reasons as when a turbine is off-line for part of the time. Firstly, a spike in the probability density of $\dot{d}$ again appears at $\dot{d} = 0$. Secondly, the reduced average amount of data with a non-zero numerical value causes greater uncertainty in the estimates.

The convergence is relatively insensitive to the method of sampling of the stress cycles. This is illustrated by the similar times required to achieve convergence across the whole operating range for the gearbox bearing damage and the structural damage. For the bearing damage, one complete torque cycle from 0 to $q_{lss}(t)$ is counted for every rotation of the low-speed shaft, while for the structural damage, cycle-counting is applied to $q_{lss}(t)$.

The importance of the second heuristic rule is demonstrated by Table 4.1. If the 10-minute results for the teeth damage are ignored, it can be seen that the convergence times have maxima in low and high wind-speeds with a minimum in the mid-range wind-speeds. Without the second heuristic rule, the estimated rate of teeth damage does not converge within acceptable time-scales in the 6 m/s and 9 m/s regions. The same problem occurs in 4 m/s and 5 m/s for the other measures of fatigue damage and for the energy capture.

The time to convergence also varies significantly as a function of power set point. In wind-speeds around the lower cut-in/out and upper cut-out/in, the convergence becomes slower as P_{sp} is reduced from P_{rated} (300 kW). Table 4.1 shows that the estimated rate of gearbox bearing damage requires 10 minutes to achieve convergence when $\overline{v} = 24$ m/s. Figure 4.4 illustrates that, if P_{sp} is then reduced to 10 kW, convergence is not achieved after 25 simulated days. (The threshold marked in Fig-

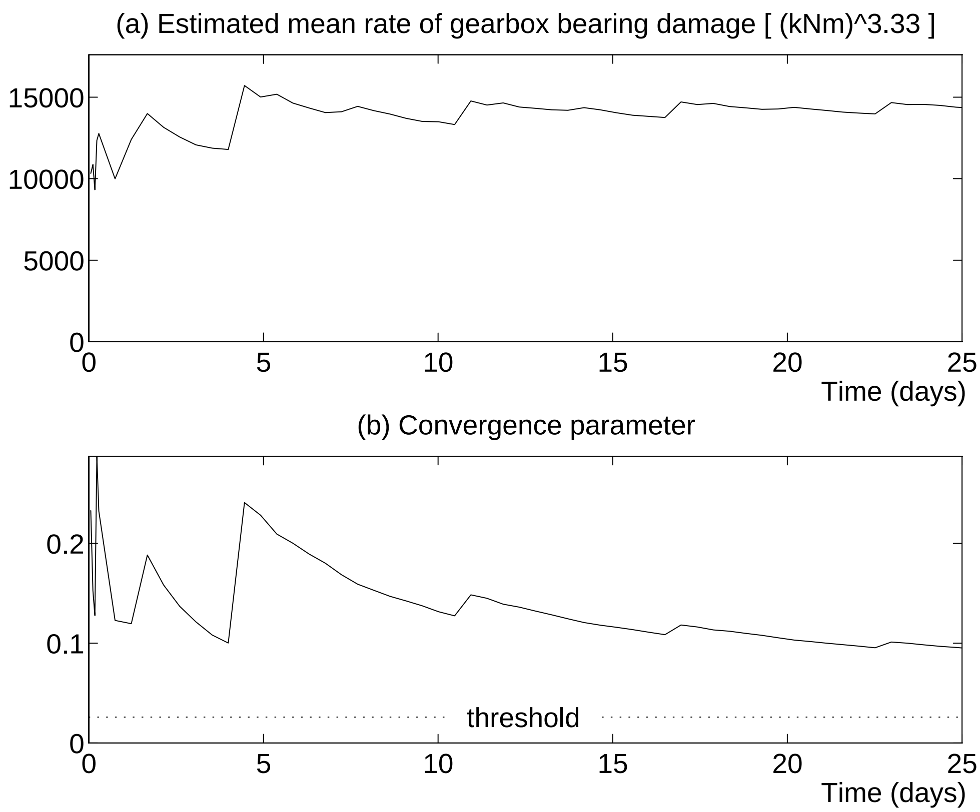

Figure 4.4: Convergence of the estimated mean rate of gearbox bearing fatigue damage with low power set point in high average wind-speed. (1 turbine; P_{sp} = 10 kW; $\overline{v}$ = 24 m/s; I = 15%.)

ure 4.4 is for a 95% confidence of the estimated mean being within +/- 5% of the true mean.) Consequently estimates for the various measures of fatigue damage were not obtained in part of the small region 23 m/s $< \overline{v} <$ 26 m/s; 0 kW $< P_{sp} <$ 140 kW. A periodic underlying trend with a period of approximately 6 days can be observed in Figure 4.4. This would indicate that the simulation package's random number generator cycles through its complete sequence several times during the run.

4.3.2 Validation of long-term gearbox bearing damage

In this section, the estimate of the annual cumulative gearbox bearing damage obtained using simulated data is compared with the equivalent figure obtained using measured data from the turbine manufacturer.

The manufacturer's data consists of estimated hourly power ($\propto$ torque) exceedance figures in 10 kW 'bins' for the turbine's 20-year lifetime. The power exceedance figures were produced using measured site data in 10-minute sets from a turbine at Carmarthen Bay in Wales. The measured data was extrapolated over 20 years using a wind-speed probability distribution with $\overline{v}_{year} = 8.5$ m/s (WEG 33-4115, 1991).

As the gearbox bearings experience complete torque cycles from zero to the low-speed shaft torque $q_{lss}(t)$ at every revolution of the shaft, and the number of stress cycles is proportional to time, from Equation 4.1 (page 104) the annual accumulated gearbox damage can be estimated from the site data as follows:

$$\hat{d}_{by1} = \frac{1}{20} \sum_{i=1}^{n_{q_{lss}}} T_i \, q_{lss_i}^{3\cdot 33}$$

where

$\hat{d}_{by1}$ = estimated annual gearbox bearing damage for site data

i = q_{lss} 'bin' number

$n_{q_{lss}}$ = number of 'bins' into which q_{lss} is divided

T_i = time spent at torque q_{lss_i}

The fraction $\frac{1}{20}$ scales down the damage estimate from over the period of 20 years to over the period of 1 year.

To estimate the annual bearing fatigue damage using simulated data ($\hat{d}_{by2}$), the summation on the left-hand side of Equation 2.13 (page 43) is used:

$$\hat{d}_{by2} = T_y \sum_{i=1}^{i_{hi}} \left\{ p\left(\overline{v}_i\right) . \overline{\hat{d}_b}\left[\overline{v}_i, P_{sp}(\overline{v}_i)\right] \right\}$$

For all i, $P_{sp}(\overline{v}_i)$ is set equal to P_{rated} and I is set to 15%. $\overline{v}_{year}$ is set to 8.5 m/s. To estimate $\overline{\hat{d}_b}$, the model of bearing fatigue damage in equation 4.2 (page 106) is used:

$$\overline{\hat{d}_b} = \frac{1}{T} \int_0^T q_{lss}^{3 \cdot 33}(t)\, dt$$

The resulting annual damage totals for the two methods are given in Table 4.2. The

Estimated annual gearbox bearing fatigue damage [(kNm)$^{3 \cdot 33}$.h]		Fractional error
From measured site data	From simulated data	
3.6×10^9	3.1×10^9	15%

Table 4.2: Comparison of estimates of annual gearbox bearing damage made with measured and simulated data.

discrepancy of 15% between the two long-term estimates is relatively low, considering the simplicity of the turbine model. One possible source of the discrepancy is that the measured data was taken from a 2-speed version of the turbine, whereas a model of a single-speed version is used in the simulation. The 2-speed design operates at a reduced rotational speed of 32 rpm (instead of 48 rpm) in low wind-speeds, to improve the tip-speed ratio and hence give higher energy capture and lower aerodynamic noise. The reduction in the rotational speed to $\frac{2}{3}$ of the rated speed increases the drive-train torques by 50%. Also, the measured data sets indicate that, for the assumed wind-speed probability distribution with $\overline{v} = 8.5$ m/s, the turbine runs at 32 rpm for approximately 30% of its operational time (WEG 33-4115, 1991). The combination of the significantly increased torques and the relatively high fraction of the time spent operating at 32 rpm therefore indicates that the gearbox fatigue damage will be higher for a 2-speed turbine than for a single-speed design. The increase in damage is heavily tempered, however, by the particularly low power levels that are experienced when operating at 32 rpm, as well as by the non-linear nature of the bearing damage $[\hat{\dot{d}}_b(t) \propto q_{lss}^{3.33}(t)]$.

4.3.3 Variation with wind-speed and turbulence intensity

The results in this section illustrate the variation of rates of fatigue damage and of cumulative fatigue damage with mean wind-speed and turbulence intensity. The power set point is constant throughout and equal to rated power (300 kW), as is standard industrial practice. The results are all normalised to the estimates obtained under the following conditions: $\overline{v} = v_{rated} = 11.5$ m/s; $I = 15\%$; $P_{sp} = P_{rated}$.

Mean rate of accumulation of fatigue damage

Results are presented in Figure 4.5 for the simulated mean rates of accumulation of fatigue damage. The solid line on each plot is for $I = 15\%$. For turbulence intensities of 10% and 20%, data has been obtained at the wind-speeds 10 m/s, 16 m/s, 22 m/s and 24 m/s, and this data is marked on the plots as individual points. In the case of the estimates for the gearbox teeth in plot (b), data for $I = 10\%$ and 20% was not recorded.

Figure 4.5(a) shows that the rate of fatigue damage of the gearbox bearings is relatively insensitive to changes in $\overline{v}$, except when $\overline{v}$ is low. The rate of damage is

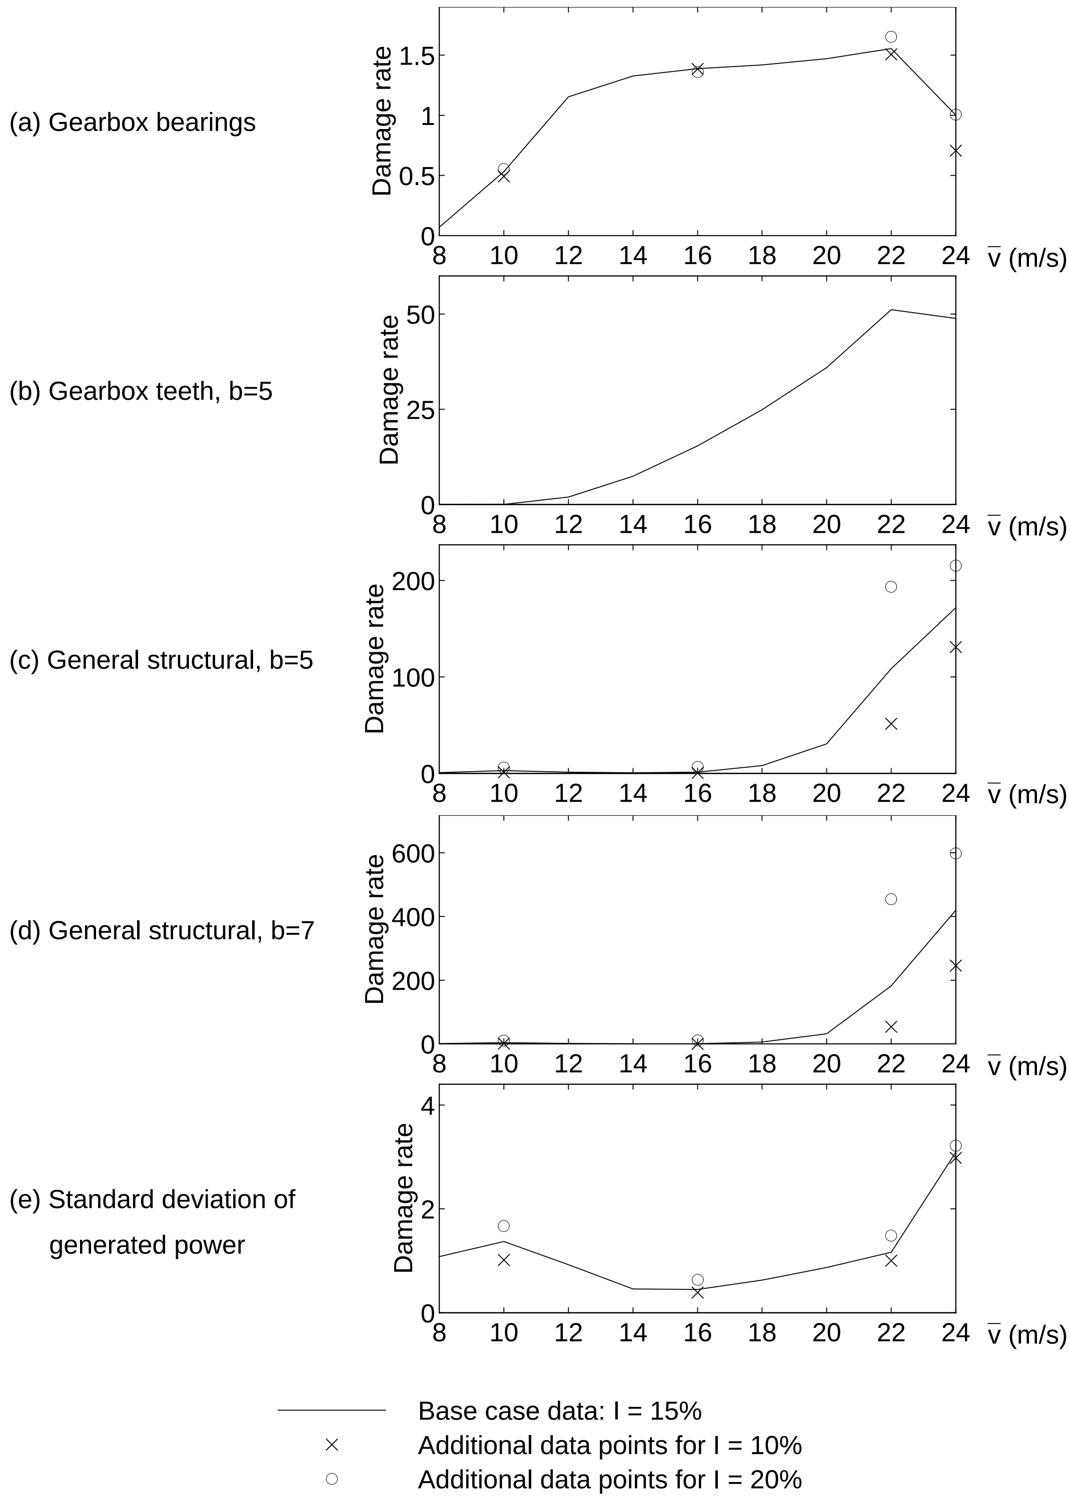

Figure 4.5: Normalised rate of accumulation of fatigue damage as a function of $\overline{v}$ *and* I *for 1 turbine*

insensitive to changes in I, even in high wind-speeds. Neither of these effects had been predicted in the intuitive basis for the investigation of the use of optimising control schemes, as it had been expected that all measures of fatigue damage would increase sharply with both $\overline{v}$ and I. One reason for this behaviour of the bearing damage is that the bearings experience a complete stress cycle at every revolution of the low-speed shaft. Another is that as the bearings do not have an endurance limit, all stress cycles count towards the total damage. As a result, non-negligible damage rates are experienced even in smooth, low-speed winds.

Figure 4.5(b) shows that, by contrast, the rate of fatigue damage of gearbox teeth with a slope index of 5 does increase sharply with $\overline{v}$. The maximum rate, at 22 m/s, is 51 times that at v_{rated}.

Figure 4.5(c) shows that the rate of general structural fatigue damage, with a slope index of 5, also increases sharply with $\overline{v}$. In this case the maximum damage rate for $I = 15\%$ is 170 times that experienced at v_{rated}. The variation as a function of turbulence intensity is also significant. When $\overline{v} = 22$ m/s, the ratio of the mean damage rate between $I = 20\%$ and $I = 10\%$ is 4 to 1, and this ratio has a maximum in 16 m/s of 14 to 1. The curve for $I = 15\%$ has a small local maximum at 10 m/s, with a normalised damage rate of 3.

Figure 4.5(d) shows that the rate of general structural fatigue damage, with a slope index of 7, has similar, but exaggerated, fatigue damage rates compared to those for the structural damage with a slope index of 5. The maximum damage rate for $I = 15\%$ is 420 times that experienced for $\overline{v} = 11.5$ m/s. When $\overline{v} = 22$ m/s, the ratio of the mean damage rate between $I = 20\%$ and $I = 10\%$ is 8 to 1, and the ratio has a maximum at 16 m/s of 140 to 1. The curve for $I = 15\%$ again has a small local maximum at 10 m/s, with a normalised damage rate of 4.

Figure 4.5(e) shows that the standard deviation of power is relatively insensitive to both $\overline{v}$ and I. The shape of the $I = 15\%$ curve is substantially different to that for the other four models of fatigue damage.

By comparing the graphs for the four simple models of fatigue damage in Figure 4.5(a)–(d), it can be seen that the rate of fatigue damage increases strongly with $\overline{v}$ in all cases, with the notable exception of the gearbox bearings. The structural damage in plots (c) and (d) exhibits the sharpest increases as a function of $\overline{v}$, and the damage rises more steeply when the slope index is larger. The standard deviation of power does not approximate to any of the damage models in plots (a) to (d).

Cumulative fatigue damage

Long-term cumulative fatigue damage is used when designing, and evaluating the performance of, windfarm supervisory controllers. Estimates of the cumulative damage are presented in this section. The damage is calculated by combining the results for the rates of fatigue damage with the annual Weibull distribution described in Section 3.4: $p(\overline{v}) \times \hat{\dot{d}}(\overline{v}, P_{sp})$. The parameters used in the Weibull distribution are $\overline{v}_{year} = 7.6$ m/s and $k_w = 1.9$. The results for the cumulative damage are presented in Figure 4.6. The solid line on each plot again represents $I = 15\%$, and individual points are plotted for the $I = 10\%$ and $I = 20\%$ data. As before, data for $I = 10\%$ and $I = 20\%$ is not available for the gearbox teeth damage.

If the estimates in Figure 4.6 are compared to those in Figure 4.5, it can be seen that the various measures of long-term cumulative damage all differ substantially as functions of $\overline{v}$ and I from the rates of damage.

Figure 4.6(a) shows that the cumulative gearbox bearing damage peaks at approximately rated wind-speed, and contributions in high wind-speeds are small. This

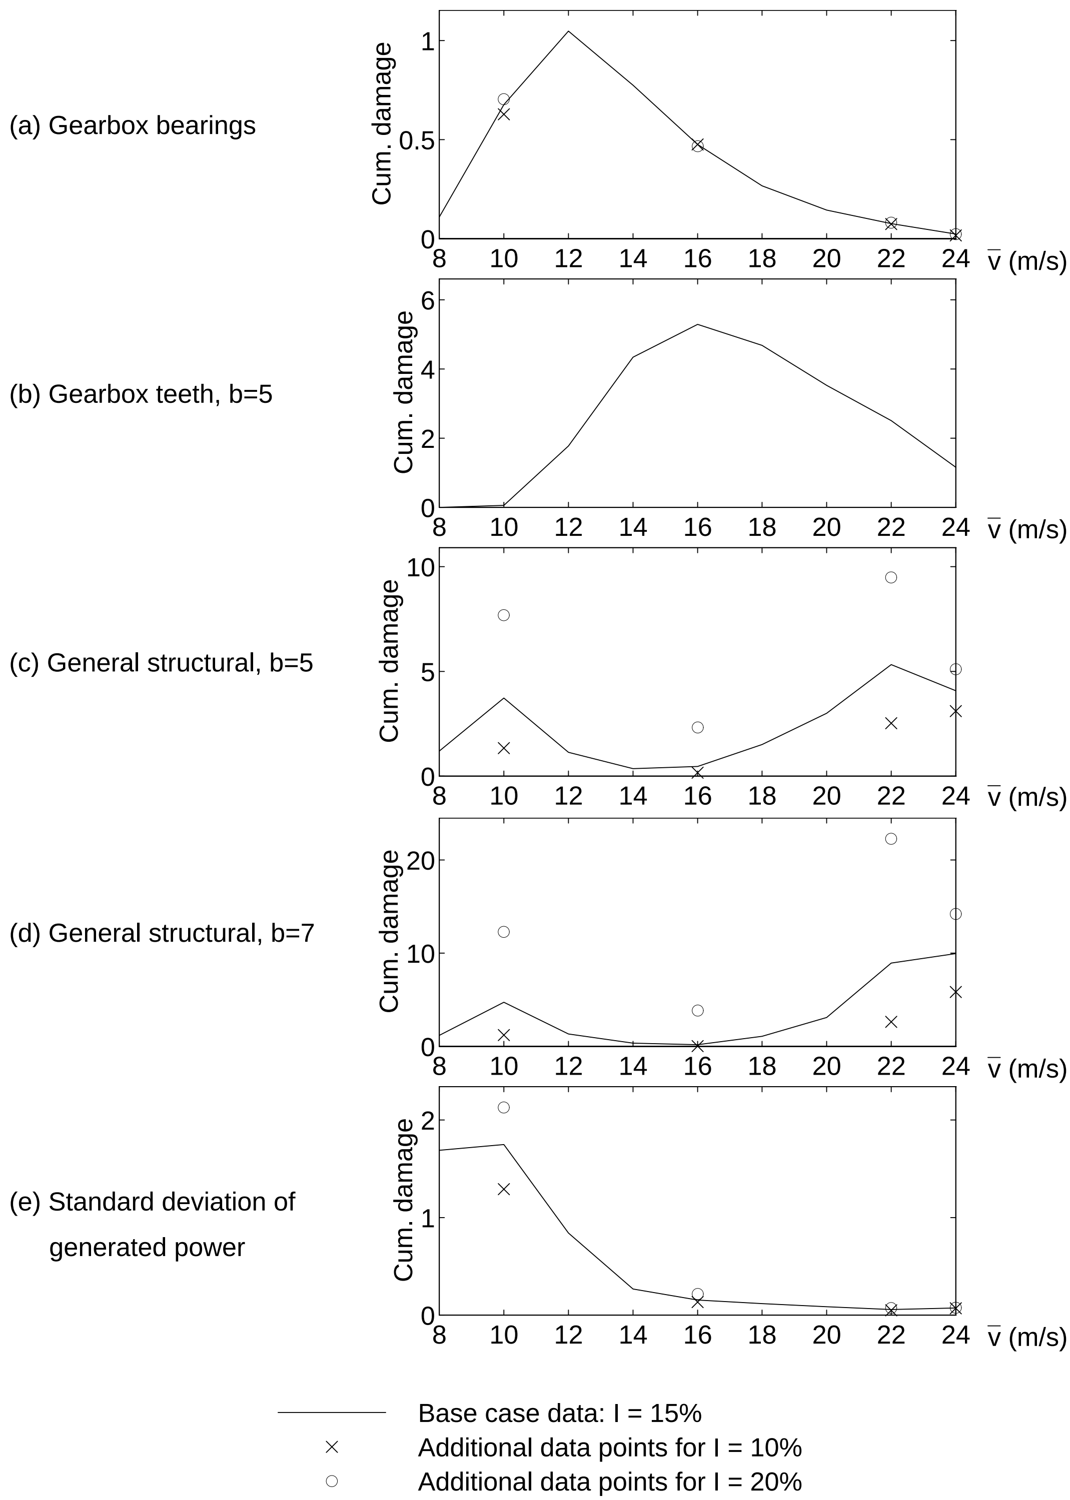

Figure 4.6: Normalised annual fatigue damage as a function of $\overline{v}$ and I for 1 turbine. ($\overline{v}_{year} = 7.6$ m/s; $k_w = 1.9$)

contrasts with the rate of damage, illustrated previously, which remains at a relatively constant level throughout the mid-range and high wind-speeds. There is no significant variation of the cumulative bearing damage with turbulence intensity.

Figure 4.6(b) shows that the curve of long-term gearbox teeth damage has a maximum in mid-range wind-speeds and minima in low and high wind-speeds. This again contrasts with the rate of damage, which increases steeply with wind-speed across the operating range.

Figure 4.6(c) shows that, for the general structural damage with $b = 5$, two regions of wind-speeds cause significant long-term damage. For $I = 15\%$, at 10 m/s there is a peak of damage of 4 times that at v_{rated}, and in 22 m/s there is a peak of 5 times that at v_{rated}. The cumulative damage function differs significantly from that for the rate of damage because of the high probability of occurrence of low wind-speeds and the low probability of occurrence of high wind-speeds. (See Figure 3.2 on page 63.) As described in the previous section, there is a small peak in the curve of the rate of damage at 10 m/s. This peak is heavily amplified by the probability distribution of wind-speed to give the relatively large maximum in the curve of cumulative damage. The reason for the peak in the curve of the rate of damage can be explained by reference to the transient responses for power ($\propto$ torque) at 9 m/s and 11.5 m/s, which are illustrated in Figures 3.10 (page 80) and 3.11. At 9 m/s, active control does not take place and the pitch angle of the blades remains constant at $\phi = 0°$ so that the turbine generates maximum power. Consequently a significant number of cycles of power with relatively large amplitudes occur. At 11.5 m/s, the number of large cycles is reduced due to extended periods spent operating at approximately rated power. As a result, the rate of damage and cumulative damage are both lower at 11.5 m/s than at 9 m/s. At any particular wind-speed, the ratio of damage between $I = 20\%$ and $I = 10\%$ is the same as for the rate of damage described previously, so the sensitivity to changes in turbulence intensity is again large.

Figure 4.6(d) shows that, for the structural damage with $b = 7$, the behaviour as a function of both $\overline{v}$ and I has the same basic shape as for $b = 5$ but with exaggerated features. For $I = 15\%$, the damage is 5 times that at v_{rated} in 10 m/s, and 10 times that at v_{rated} in 24 m/s. Again, for any particular wind-speed the ratio of damage between $I = 20\%$ and $I = 10\%$ is the same as for the rates of damage described previously, and so the sensitivity to changes in turbulence intensity is large.

Figure 4.6(e) shows that, in contrast to the standard deviation of power, the cumulative standard deviation of power over the long term falls to negligible levels in the medium and high wind-speed regions.

By comparing the graphs for the four simple models of fatigue damage in Figure 4.6(a)–(d), it can be seen that the cumulative fatigue damage varies very differently with $\overline{v}$ and I for the different models. The cumulative standard deviation of power presented in Figure 4.6(e) does not approximate to the cumulative damage for any of the models in plots (a)–(d). The general structural damage varies very widely with both $\overline{v}$ and I, which suggests that structural components will have the greatest sensitivity to changes in wind conditions due to operation within windfarms, that is, to widely varying $\overline{v}$ and I values due to micro-siting and wake effects.

4.3.4 Variation with power set point

The results in this section illustrate the variation of rates of fatigue damage and of cumulative fatigue damage with power set point, for wind-speeds across the whole operating range and for a constant turbulence intensity of 15%. The results are all normalised to the estimates obtained under the following conditions: $\overline{v} = v_{rated} = 11.5$ m/s; $I = 15\%$; $P_{sp} = P_{rated}$.

Mean rate of accumulation of fatigue damage

Results are presented in Figures 4.7–4.11 which illustrate the variation of the rates of fatigue damage. As described in Section 4.3.1, estimated damage rates have not been obtained in part of the small region 23 m/s $< \overline{v} <$ 26 m/s; 0 kW $< P_{sp} <$ 140 kW.

Figure 4.7 shows that the estimated rate of gearbox bearing damage ($\hat{\bar{\dot{d}}}_b$) falls steeply as a function of P_{sp} from the conventional operating curve $P_{sp} = P_{rated} = 300$ kW, across most of the operating range. Therefore at many wind-speeds modest fractional reductions in P_{sp} can provide useful reductions in $\hat{\bar{\dot{d}}}_b$.

Figure 4.8 shows that the rate of gearbox teeth damage ($\hat{\bar{\dot{d}}}_t$) falls very steeply as a function of P_{sp} from P_{rated}, for all wind-speeds where significant damage rates are experienced. Therefore large reductions in $\hat{\bar{\dot{d}}}_t$ can be obtained for modest reductions in P_{sp}. By comparison with the bearing damage, the steep rise in $\hat{\bar{\dot{d}}}_t$ with both set point and wind-speed for set points close to rated is due partly to the higher S–N curve slope index, and partly to the torque threshold which is introduced by the endurance limit. As it is assumed that torques below the endurance limit cause no fatigue damage, slight reductions in the torque spikes that are transmitted through the drive-train give large reductions in the mean rate of fatigue damage.

Figure 4.9 shows that the rate of accumulation of the general structural damage for $b = 3$ ($\hat{\bar{\dot{d}}}_{s3}$) falls steeply if P_{sp} is reduced from P_{rated}, for all wind-speeds where significant damage rates are experienced. Therefore useful reductions in $\hat{\bar{\dot{d}}}_{s3}$ can be obtained for modest reductions in P_{sp}.

Figure 4.10 shows that the rate of accumulation of structural damage for $b = 7$ ($\hat{\bar{\dot{d}}}_{s7}$) falls more steeply than for $b = 3$ if P_{sp} is reduced from P_{rated}, for all wind-speeds where significant damage rates are experienced. Therefore larger reductions in $\hat{\bar{\dot{d}}}_{s7}$ than in $\hat{\bar{\dot{d}}}_{s3}$ can be obtained for modest reductions in set point. $\hat{\bar{\dot{d}}}_{s7}$ also rises more

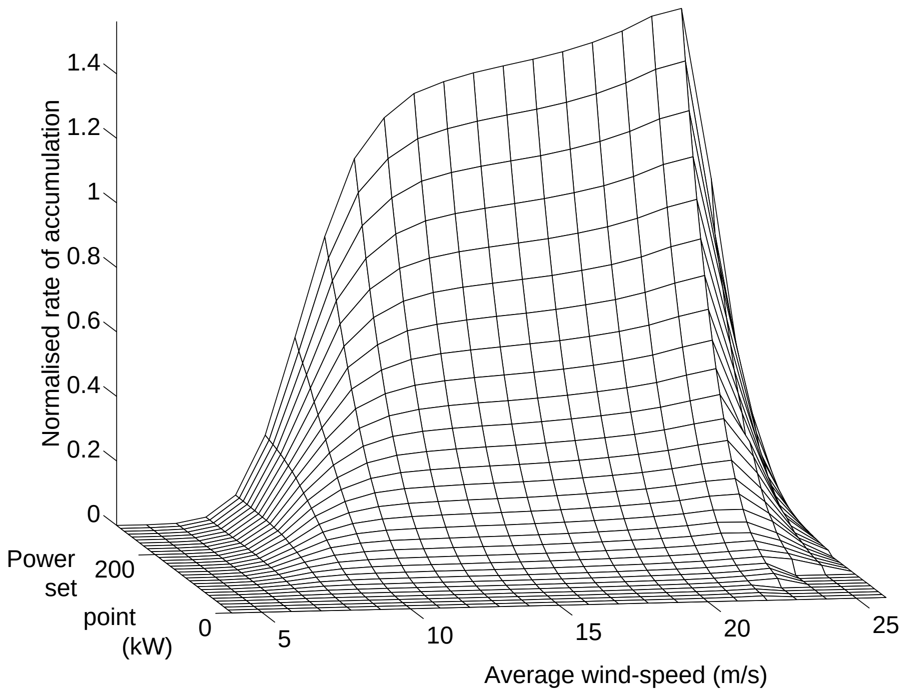

Figure 4.7: Normalised rate of accumulation of gearbox bearing fatigue damage for 1 turbine (I = 15%)

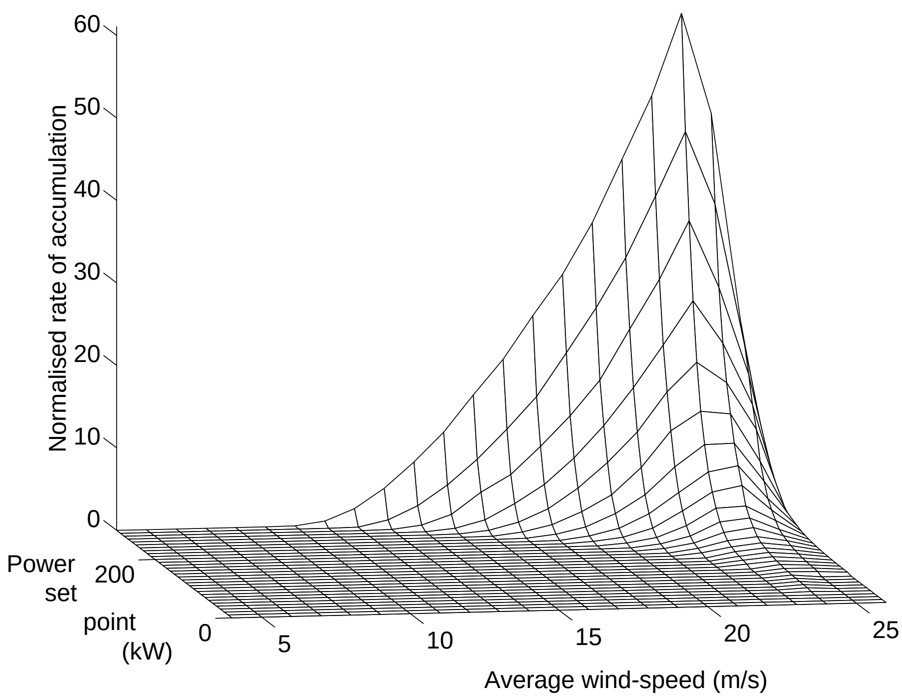

Figure 4.8: Normalised rate of accumulation of gearbox teeth fatigue damage with slope index of 5 for 1 turbine (I = 15%)

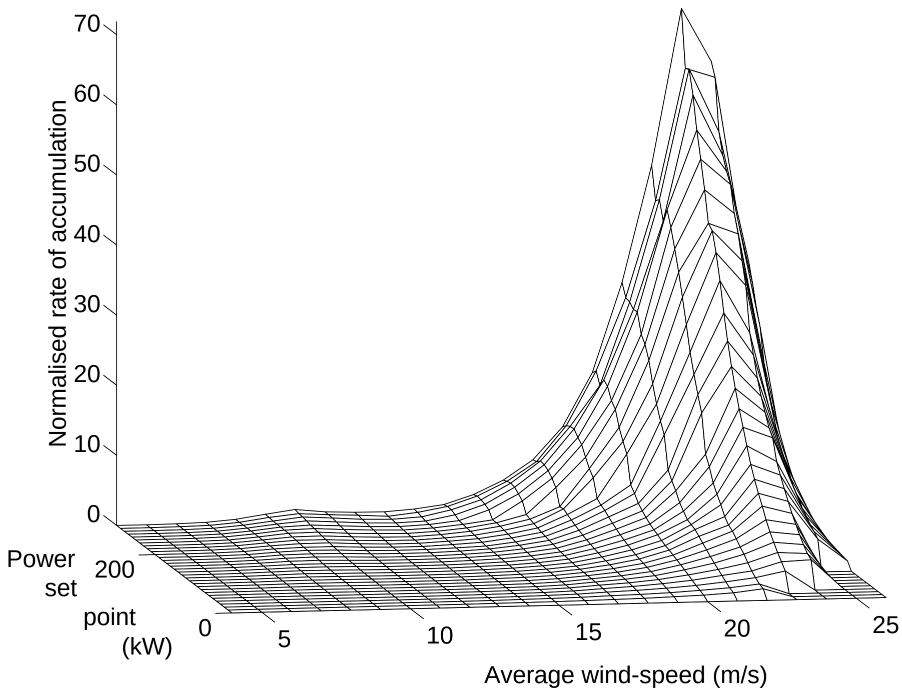

Figure 4.9: Normalised rate of accumulation of general structural fatigue damage with slope index of 3 for 1 turbine (I = 15%)

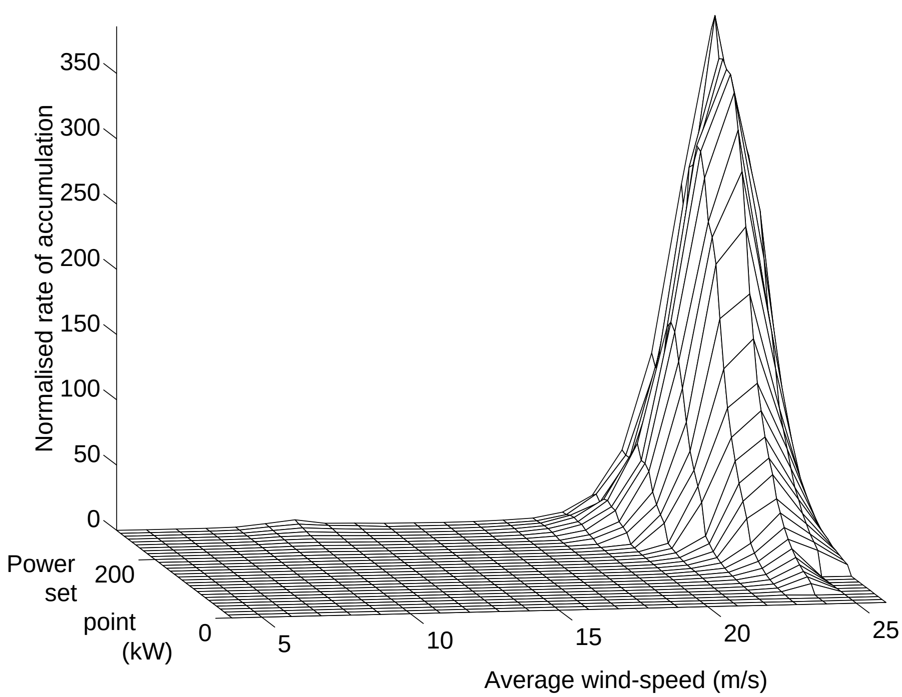

Figure 4.10: Normalised rate of accumulation of general structural fatigue damage with slope index of 7 for 1 turbine (I = 15%)

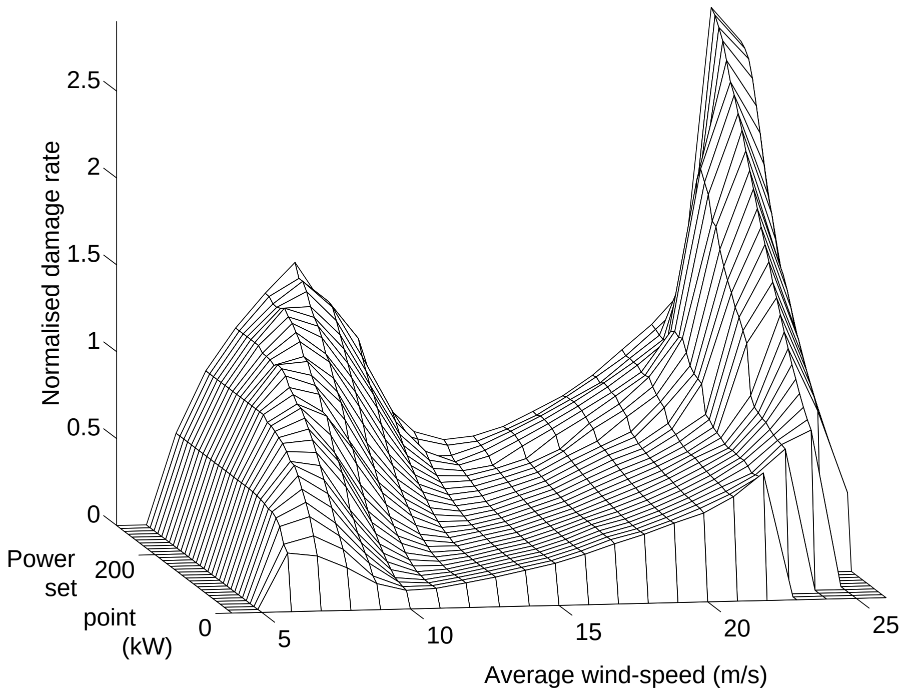

Figure 4.11: Normalised standard deviation of generated power for 1 turbine. (I = 15%)

steeply as a function of $\overline{v}$ than is the case for $\hat{\bar{\dot{d}}}_{s3}$, and the peak damage rate is substantially larger, by a factor of 5, than for $\hat{\bar{\dot{d}}}_{s3}$.

Figure 4.11 shows that the rate of accumulation of fatigue damage as measured by the standard deviation of power is a weak function of P_{sp} , if P_{sp} is reduced from P_{rated}, for all wind-speeds other than those around the upper cut-out/in. It is the only model of fatigue damage to have a significant discontinuous change in the rate of damage as P_{sp} is increased from 0 kW.

Cumulative fatigue damage

It is again informative to consider the variation of the long-term cumulative fatigue damage, $p(\overline{v}) \times \hat{\bar{\dot{d}}}(\overline{v}, P_{sp})$, as long-term damage is used in the design and performance evaluation of supervisory controllers.

To allow a complete comparison to be made of all the variables relevant to supervisory control for which results have been obtained, the annual energy capture surface is presented again in Figure 4.12. The results for the annual cumulative fatigue damage are presented in Figures 4.13–4.17. The Weibull wind-speed probability distribution that is used to produce both the energy capture and fatigue damage estimates has an annual mean wind-speed of 7.6 m/s and a shape factor of 1.9. (For details of the Weibull distribution see Section 3.4 on pages 61*ff.*)

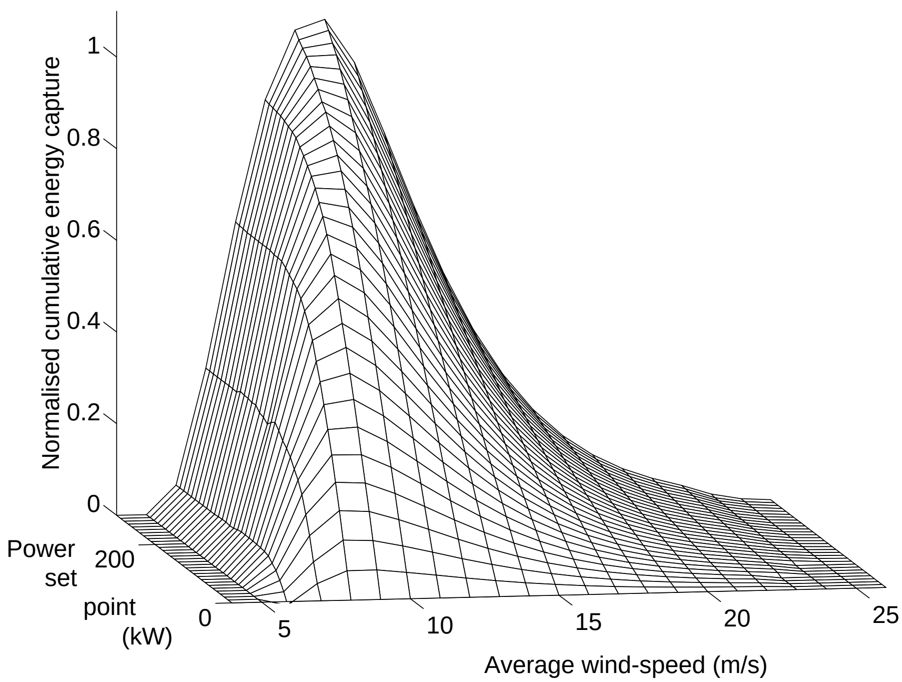

Figure 4.12: Normalised annual energy capture for 1 turbine ($\overline{v}_{year}$ = 7.6 m/s; k_w = 1.9; I = 15%)

Figure 4.13 shows that modest de-rating can give useful reductions in annual gearbox bearing damage. The bearing damage function has, however, a surprisingly similar shape to that of the annual energy capture, and consequently to achieve significant reductions in the bearing damage must necessitate significant reductions in the energy capture.

Figure 4.14 shows that the annual gearbox teeth damage is a very strong function of

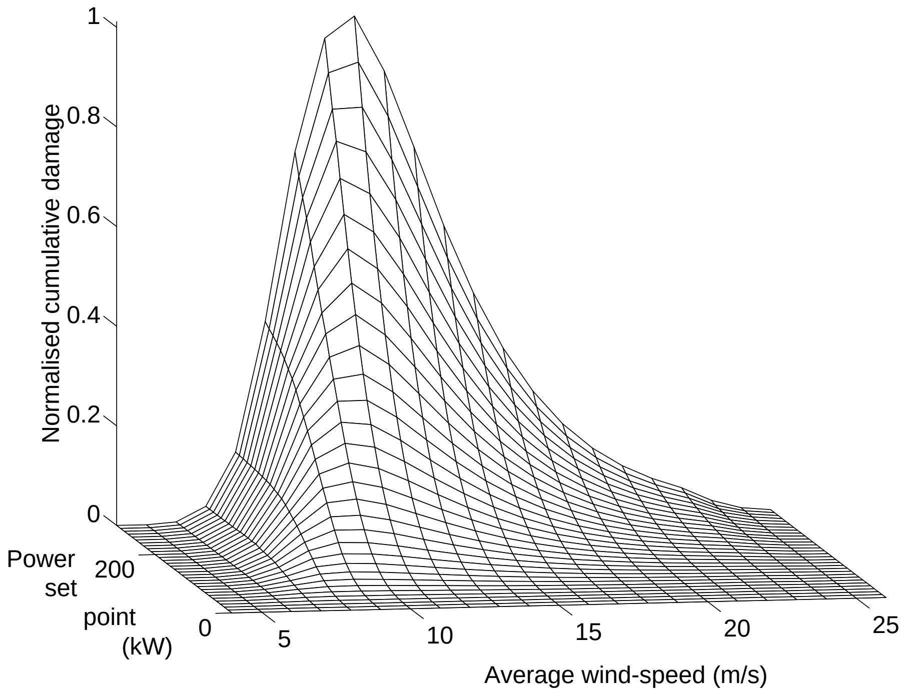

Figure 4.13: Normalised annual gearbox bearing fatigue damage for 1 turbine ($\overline{v}_{year}$= 7.6 m/s; k_w = 1.9; I = 15%)

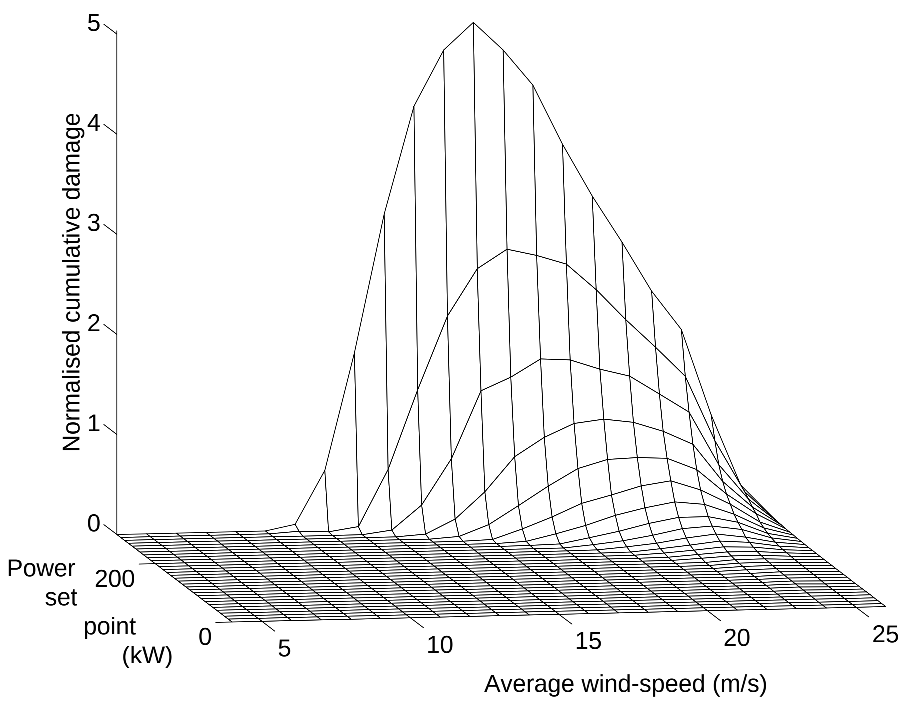

Figure 4.14: Normalised annual gearbox teeth fatigue damage with slope index of 5 for 1 turbine ($\overline{v}_{year}$= 7.6 m/s; k_w = 1.9; I = 15%)

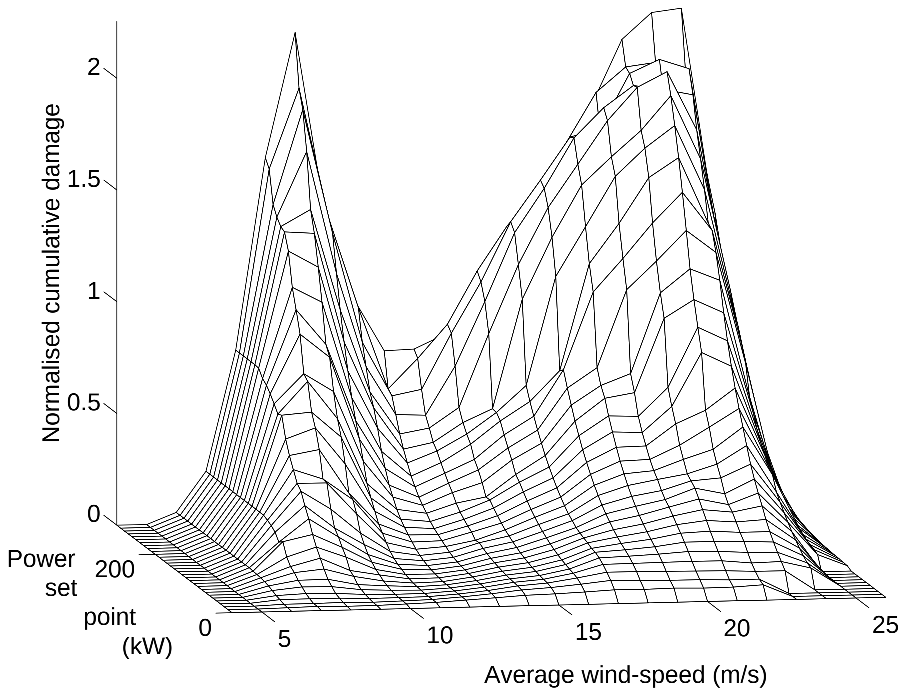

Figure 4.15: Normalised annual structural fatigue damage with slope index of 3 for 1 turbine ($\overline{v}_{year} = 7.6$ m/s; $k_w = 1.9$; $I = 15\%$)

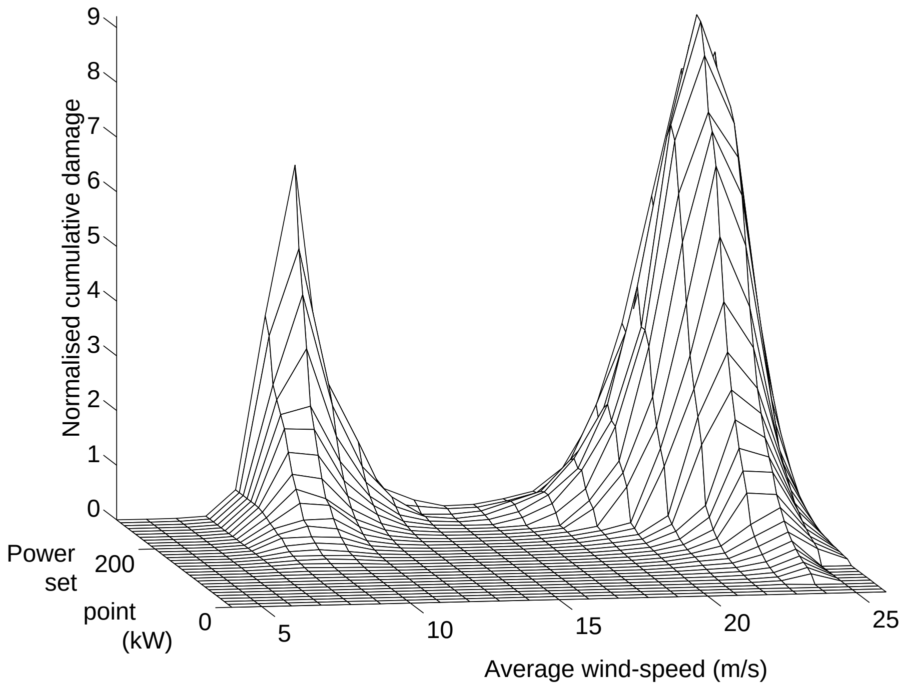

Figure 4.16: Normalised annual structural fatigue damage with slope index of 7 for 1 turbine ($\overline{v}_{year} = 7.6$ m/s; $k_w = 1.9$; $I = 15\%$)

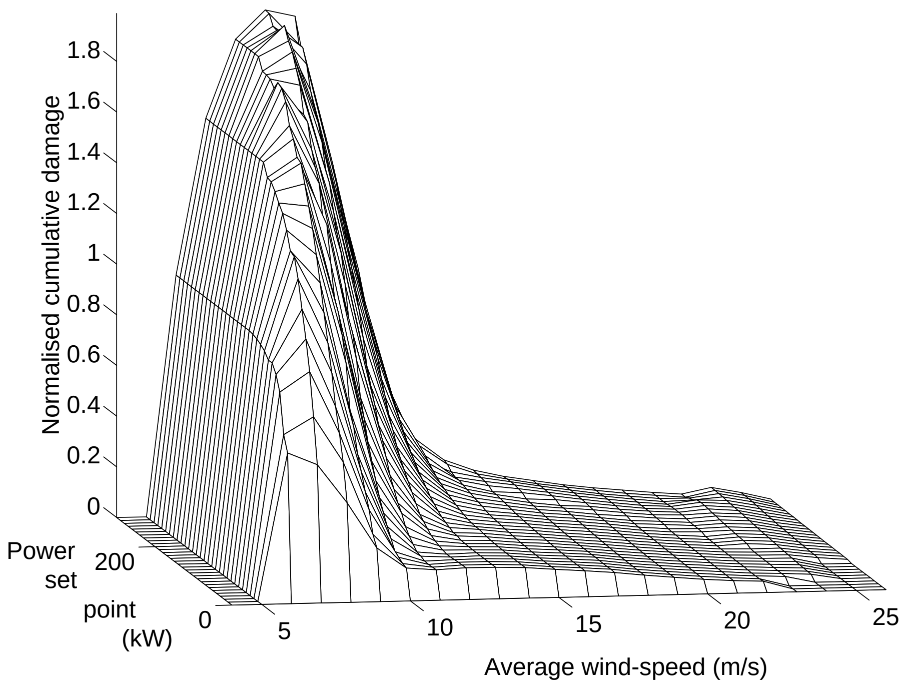

Figure 4.17: Normalised annual standard deviation of generated power for 1 turbine ($\overline{v}_{year}$ = 7.6 m/s; k_w = 1.9; I = 15%)

set point at all wind-speeds at which significant contributions to the total damage are made. By comparing Figure 4.14 with Figure 4.12, it can be seen that large reductions in gearbox teeth damage can be achieved for small reductions in energy capture, by making small reductions in set point in mid-range wind-speeds and large reductions in set point in high wind-speeds.

Figure 4.15 shows that the annual structural fatigue damage for $b = 3$ falls steeply if P_{sp} is reduced from P_{rated}, at all wind-speeds where significant contributions are made to the total damage, that is, in the region around 10 m/s and in high wind-speeds. By comparing Figure 4.15 with Figure 4.12, it can be seen that if P_{sp} is heavily reduced in high wind-speeds, where the contributions to the total energy capture are very low, substantial reductions in the structural damage can be achieved for small reductions in energy capture.

Figure 4.16 shows that the annual structural fatigue damage for $b = 7$ is a similar function to that for $b = 3$ but with higher, steeper peaks and lower, deeper valleys. If P_{sp} is heavily reduced in high wind-speeds, then for a given reduction in energy capture, larger reductions in the structural damage can be achieved than for $b = 3$.

A comparison of the four annual damage functions in Figures 4.13–4.16 shows that the shapes of the damage functions vary widely. The torque cycles from 0 kW to $q_{lss}(t)$ that are experienced by the gearbox bearings and teeth at every revolution of the low-speed shaft lead to damage levels in lower wind-speeds which are high relative to the general structural damage, the latter being evaluated using cycle-counting. The combination of the raw rates of accumulation of damage with an annual wind-speed probability distribution to obtain cumulative damage levels strongly increases the significance of the lower wind-speeds, and causes the bearing and teeth damage functions to contain maxima in low to mid-range wind-speeds.

Figure 4.17 shows that the annual standard deviation of power can only be reduced slightly if modest reductions in P_{sp} are made from the rated power*. Such reductions in P_{sp} can only give useful reductions in σ_p within a small range of wind-speeds, between 10 m/s and 12 m/s. The suitability of σ_p as a generic indicator of fatigue damage can be considered by comparing Figure 4.17 with the four damage functions in Figures 4.13–4.16. This comparison indicates that σ_p does not approximate to any of the four damage functions.

*Although the improvement of power quality is not an objective of the supervisory controllers considered here, the insensitivity of σ_p to P_{sp} suggests that there is little potential for improvement in the power quality of (non-interacting) constant speed turbines via set point adjustment. Further evidence is the insensitivity of σ_p to turbulence intensity, as illustrated by Figure 4.6(e), which means that there are no wind conditions which cause disproportionately high σ_p levels and in which significant reductions in the total σ_p could then be achieved for small reductions in energy capture.

4.4 Results for simulation of 3 turbines in a row

4.4.1 'Large' inter-turbine spacings

Results are presented in this section for the gearbox bearing damage of the turbines in the 3-turbine windfarm simulation that is described in Section 3.2. The two layouts that are investigated are illustrated in Figure 3.1 (page 58), and the turbines are numbered from "1" (furthest upwind) to "3" (furthest downwind). Estimated damage rates for the 'large' inter-turbine spacings (7.9D/7D) are presented in this section.

Figure 4.18 illustrates the total rate of bearing damage for the three turbines as a function of the power set points of turbines 1 and 2. The set point of turbine 3 is

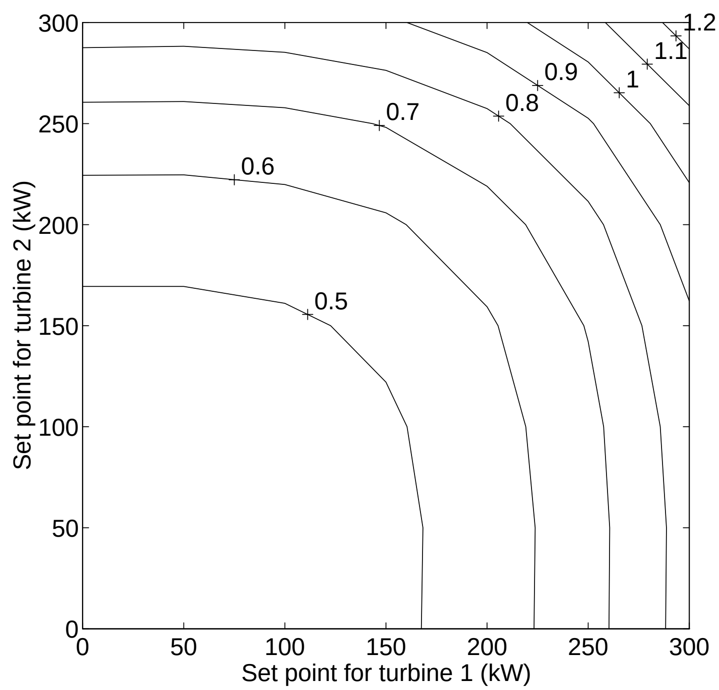

Figure 4.18: Total normalised rate of gearbox bearing damage for 3 turbines with 'large' inter-turbine spacings. ($\overline{v}$ = 14 m/s; I = 15%; $P_{sp}^{(3)} = P_{rated}$ = 300 kW)

held constant at P_{rated} throughout. The wind-speed is 14 m/s and the turbulence intensity is 15%. The numerical values for the total damage rates that are marked on the contours are normalised to three times the estimated rate for a single turbine under the following conditions: $\overline{v} = v_{rated} = 11.5$ m/s; $I = 15\%$; $P_{sp} = P_{rated}$.

The near symmetry of the plot about the line $P_{sp}^{(1)} = P_{sp}^{(2)}$ indicates that the effects upon the rate of bearing damage of operation in wakes with the 'large' inter-turbine spacings are negligible. The same result was also found for the gearbox teeth damage ($b = 5$) and for the energy capture (Spruce and Dexter, 1992). Further experiments showed that, at other wind-speeds between v_{rated} (11.5 m/s) and the upper cut-out (25 m/s), the effect of the wakes upon the gearbox damage and energy capture is similarly negligible.

4.4.2 'Small' inter-turbine spacings

Results are presented in this section for the annual fatigue damage of the turbines in the simulated windfarm with 'small' inter-turbine spacings (3D/3D). The estimates are normalised with respect to the damage, $p(\overline{v})\,\hat{\dot{d}}(\overline{v}, P_{sp})$, for a single turbine under the following conditions: $\overline{v} = v_{rated} = 11.5$ m/s; $I = 15\%$; $P_{sp} = P_{rated}$.

Gearbox

Figures 4.19 and 4.20 illustrate the variation of the individual turbines' cumulative bearing damage as functions of the free stream wind-speed and turbulence intensity. The turbines' power set points are all set to P_{rated} throughout. In contrast to the results for the 'large' inter-turbine spacings, Figures 4.19 and 4.20 show that with 'small' spacings the turbines experience significantly different levels of bearing damage as a result of the wake effects.

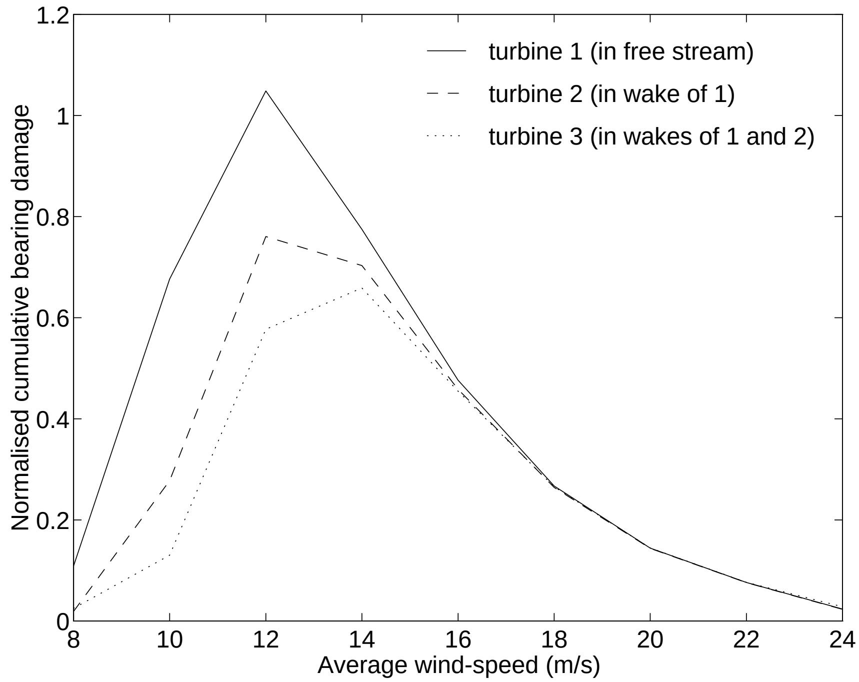

Figure 4.19: Normalised annual gearbox bearing fatigue damage for 3 turbines with 'small' inter-turbine spacings, as a function of $\overline{v}$.
($P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated} =$ *300 kW*; $I =$ *15%*)

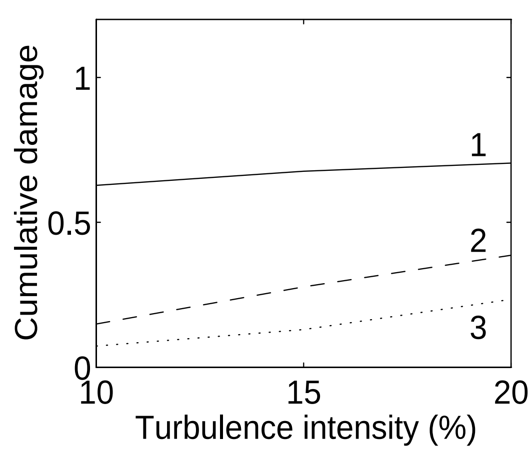

Figure 4.20: Normalised annual gearbox bearing fatigue damage for 3 turbines with 'small' inter-turbine spacings, as a function of I.
($P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated} =$ *300 kW*; $\overline{v} =$ *10 m/s*)

Figure 4.19 shows that at no point across the operating range does the bearing damage increase at the downstream turbines, and the damage decreases significantly downstream in low free stream wind-speeds. At the downstream turbines, $\overline{v}$ is lower than that of the free stream wind and I is higher than that of the free stream. As illustrated by Figure 4.6(a) (page 123), in lower wind-speeds the bearing damage decreases sharply as $\overline{v}$ decreases, yet the rate of damage is insensitive to changes in turbulence intensity. As a result, the damage decreases significantly downstream in lower wind-speeds. Figure 4.20 illustrates that higher levels of turbulence in the free stream wind cause only modest increases in the bearing damage of all three turbines.

Similar experiments investigated the variation of the gearbox teeth damage with wind-speed. The teeth damage is insensitive to the wake effects, with very small changes in damage levels at the downstream turbines for all wind-speeds across the operating range.

Structural components

Figures 4.21 and 4.22 illustrate the variation of the generalised structural fatigue damage for the individual turbines as functions of the free stream wind-speed and turbulence intensity, for the same 'small' inter-turbine spacings (3.0D/3.0D). The slope index is 7 and all of the turbines' power set points are again set to P_{rated}. The figures show that significantly higher damage levels are experienced at turbine 2, which operates within a single wake, and the damage levels increase still further at turbine 3, which operates within two wakes.

Figure 4.21 shows that the largest increases in damage are in wind-speeds around v_{rated} (11.5 m/s). This is because of the fall in mean wind-speed and rise in turbulence at the downstream turbines. As Figure 4.6(d) illustrates, reductions in $\overline{v}$ in

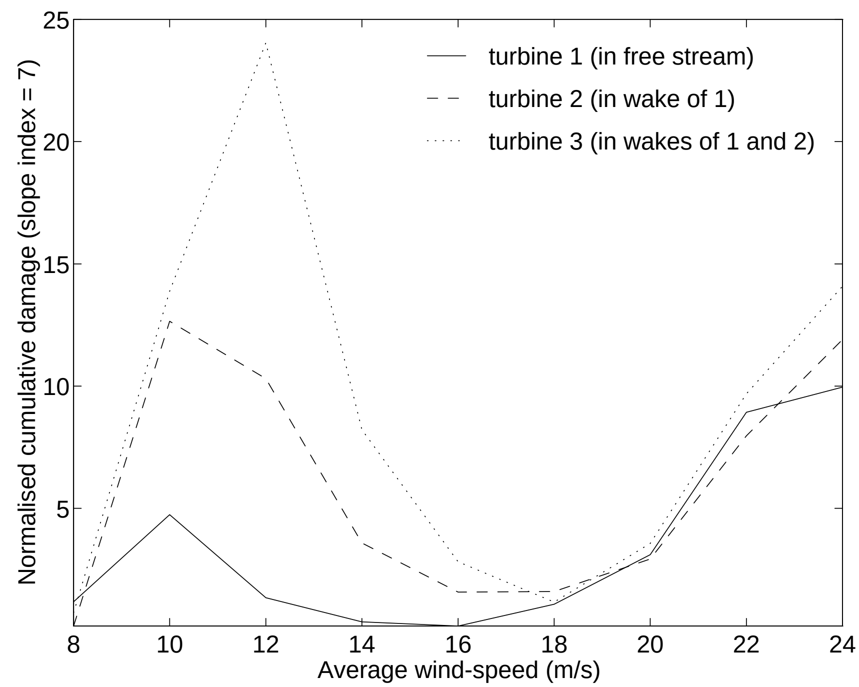

Figure 4.21: Normalised annual structural fatigue damage for 3 turbines with 'small' inter-turbine spacings, as a function of $\overline{v}$.
($b = 7$; $P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated} = 300$ kW; $I = 15\%$)

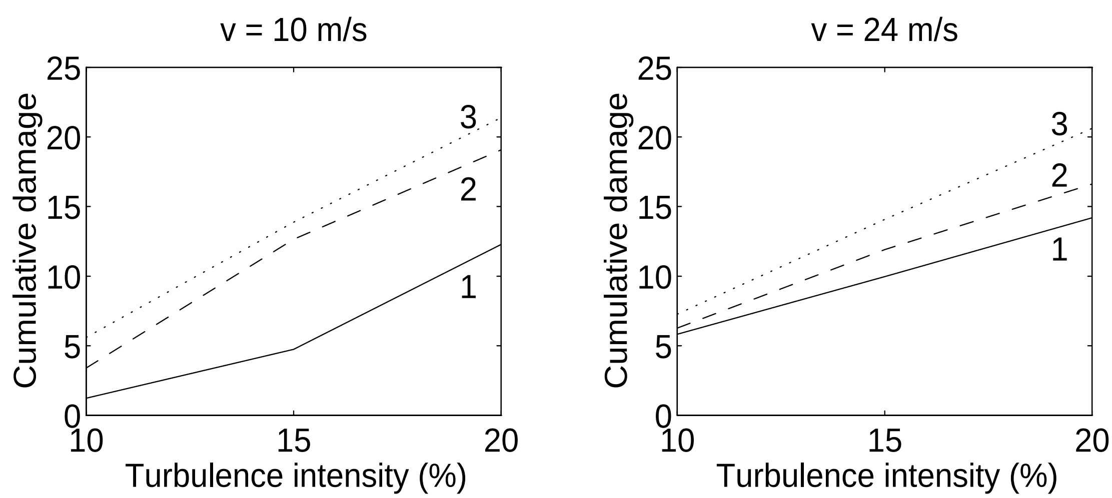

Figure 4.22: Normalised annual structural fatigue damage for 3 turbines with 'small' inter-turbine spacings, as a function of I.
($b = 7$; $P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated} = 300$ kW)

the range 10 m/s $< \overline{v} <$ 16 m/s and increases in I both cause substantial increases in the general structural damage.

Figure 4.22 shows that there are large increases in the general structural fatigue damage at all the turbines as the free stream turbulence intensity increases, for both $\overline{v} = 10$ m/s and $\overline{v} = 24$ m/s. An increase in the free stream turbulence intensity from 10% to 20% results in a fourfold to tenfold increase in the turbines' damage levels when $\overline{v} = 10$ m/s, and a twofold to threefold increase when $\overline{v} = 24$ m/s. The strong relationship between turbulence intensity and fatigue damage helps to explain why there are higher levels of damage at the downstream turbines. Similar results were obtained for various free stream turbulence intensities at wind-speeds of 16 m/s and 22 m/s, and in both cases there are again large increases in the structural damage as a function of free stream turbulence intensity.

The variation in the general structural damage for a slope index of 5 was also investigated. The shapes of the curves of fatigue damage as functions of both wind-speed and turbulence intensity are all similar to those for a slope index of 7.

4.5 Conclusions

All the models of fatigue damage that have been used in the design and performance evaluation of supervisory controllers have been described in this chapter. The estimated annual gearbox bearing damage predicted by the simulation has been validated against an estimate from measured data. Suitably close agreement is achieved between the two results and it can be concluded that, for the purposes of testing supervisory controllers, the single turbine model provides sufficiently accurate estimates of the gearbox bearing damage.

The estimated rates of damage for an individual turbine have been presented, as functions of $\overline{v}$, I and P_{sp}, and the results for the various models of fatigue damage differ widely from one another. Similarly, the results for the annual damage of an individual turbine differ widely from one another as functions of $\overline{v}$, I and P_{sp}. The shapes of the cumulative damage functions are all significantly different to those of the rates of damage.

Part of the basis for the investigation of the use of optimising control schemes was the expectation that all measures of fatigue damage would increase sharply with both $\overline{v}$ and I. However, it has been found that the cumulative damage of the most expensive component in the turbine, the gearbox, contradicts this supposition. The cumulative gearbox bearing damage has small contributions in high wind-speeds and is insensitive to changes in turbulence intensity, and the teeth damage has modest contributions in high wind-speeds. The generalised structural damage does rise sharply with $\overline{v}$ in high wind-speeds, though, and also increases sharply with I. The principal cause of the differences between the gearbox damage and the generalised structural damage is the fundamentally different mechanisms by which the components' fatigue loads are generated. The gearbox teeth and bearings sample the low-speed shaft torque at every revolution of the shaft, thereby experiencing complete torque cycles between zero and the low-speed shaft torque at every revolution. The generalised structural 'component', by contrast, experiences torque cycles which are simply equal to the low-speed shaft torque itself.

Comparison of the annual damage functions with the annual energy capture function shows that the components for which the supervisory control of individual turbines offers the greatest potential for reductions in fatigue damage include: the gearbox teeth; those structural components whose fatigue loads are approximately proportional to the low-speed shaft torque. The shape of the gearbox bearing damage function is so similar to that for energy capture that attaining significant reductions in the bearing damage of an isolated turbine necessitates significant reductions in energy capture.

Results for the standard deviation of power have been compared with those for the other models of fatigue damage. It has been found that σ_p is substantially different as a function of $\overline{v}$, I and P_{sp} from all of the other models, and therefore it can be concluded that σ_p does not provide a suitable generic indication of fatigue damage.

Annual fatigue damage results have been presented for the 3-turbine windfarm simulation, both for 'large' inter-turbine spacings (7.9D/7D) and for 'small' spacings (3D/3D). The data obtained for the 'large' spacings is for the gearbox damage in the region $\overline{v} > v_{rated}$. The results indicate that the wake and wake superposition effects cause negligible changes in gearbox damage. The results for the 'small' spacings, by contrast, indicate that some significant interaction effects occur. There are large rises in the generalised structural damage of the downstream turbines in low to mid-range wind-speeds, and it is found that, for the gearbox, the wake effects either cause significant reductions in the damage at the downstream turbines (bearings) or cause no change in the damage downstream (teeth). The results for the gearbox damage again contradict an initial basis for the investigation of optimising supervisory controllers, that all measures of fatigue damage would increase sharply at turbines operating in the wakes of other turbines.

Chapter 5

Windfarm controller design

5.1 Introduction

This chapter describes the implementation of, and results arising from, the methodology for the development of windfarm supervisory controllers which is presented in Chapter 2. Sections 5.2, 5.3, 5.4 and 5.5 in this chapter run directly parallel to Sections 2.3, 2.4, 2.5 and 2.6 respectively.

The 'ideal' power curves for an isolated turbine, for both two and one tariff seasons per annum, are evaluated in Section 5.2. The weighting parameters of the controller's cost function are evaluated in Section 5.3, again for an isolated turbine and for both two and one tariff seasons per annum. The cost function for the isolated turbine is applied to the 3 interacting turbines in the simulated small windfarm in Section 5.4, for the case of one tariff season per annum only. The resulting operation of a simple on-line optimising control algorithm, using the single tariff cost function developed in Section 5.4, is presented in Section 5.5.

5.2 'Ideal' power curves

The first stage in the development of a windfarm control algorithm is the evaluation of the 'ideal' power curves for an isolated turbine, as described in Section 2.3. To evaluate the ideal power curves, searches are carried out to solve for the functions $P_{sp}(i_s, \overline{v})$ and $P_{sp}(\overline{v})$ in the 'ideal' performance functions (Equations 2.10 to 2.13 on pages 41 and 43).

A genetic algorithm is used to carry out the searches, as this class of algorithm has the following attributes:

- capable of operation with a large number of independent variables
- able to operate with a large search space
- capable of high off-line performance
- robust with respect to possible local minima in the search surface

5.2.1 Genetic search

A genetic search algorithm aims to maximise the 'fitness', or performance, of a function. An overview of the genetic algorithm that is used is given in Appendix E. In this application, the fitness is equal to the estimated annual financial income $\hat{m}_{year}$, as defined in Equations 2.10 and 2.12. The constraint on fatigue damage is limited to the gearbox and, as the bearing damage is more critical than the teeth damage for the particular design of gearbox that is modelled (Hicks, 1992), the bearing damage alone is used in the constraint. A simple brick wall penalty function is used to implement the constraint, whereby the fitness of any power curve for which the damage exceeds its limit is set to zero.

The optimisation problem is converted into a form suitable for solution by a genetic algorithm by using a multiparameter, binary, discretised coding of the power set points. 5-bit binary numbers are used to represent each set point, and 31 of the 32 possible values are allocated the following power set points:

$$P_{sp} = 0\ \text{kW},\ 10\ \text{kW},\ 20\ \text{kW}, \ldots, 300\ \text{kW}$$

It is not necessary to distribute the set points uniformly throughout the range, and the set point 295 kW is allocated to the 32nd point.

The range of wind-speeds is discretised to give 18 or 36 independent variables for the case of 1 or 2 seasons per annum respectively:

$$\overline{v} = 8\ \text{m/s},\ 9\ \text{m/s},\ 10\ \text{m/s}, \ldots, 25\ \text{m/s}$$

The use of these wind-speeds and set points gives a search space of 32^{36} points in the 2-tariff case, which is equal to approximately 10^{54} points.

In the genetic search algorithm, each power curve, or set of two power curves, is represented by a single string. To form one string, a set of 18 or 36 5-bit numbers are concatenated together, so each string consists of either 95 or 190 bits. Several separate randomly selected populations are formed, each containing 100 strings. The genetic algorithm then allows these populations to evolve for up to approximately 100 000 generations, until the populations have converged to the region of maximum fitness. (Details of the use of multiple 'parallel' populations is given in Appendix E.) The strategy with the highest fitness in the populations is the ideal power curve.

The genetic algorithm that is implemented is Goldberg's 'Simple Genetic Algorithm', and several extensions to the basic algorithm are also investigated (Goldberg, 1989). A brief description of the algorithm is given in Appendix E, together with details of the extensions and their effect upon the performance of the search.

Constraint on feasible power set points in high wind-speeds

A constraint is applied to the feasible region for the power set points in wind-speeds around the upper cut-out. This constraint prevents the ideal power curves from defining set points in high wind-speeds that cause the turbine to generate low average power levels. A turbine often cuts in for short periods of time in wind-speeds around the upper cut-out. If the mean power level that is generated while the turbine is on-line is low, then the total energy that is generated during the period spent on-line will not be worth the fatigue damage that the turbine undergoes during the synchronisation and de-synchronisation to and from the grid.

Table 5.1 gives estimates of the average fraction of the time that the turbine spends off-line for mean wind-speeds around the upper cut-out. The table also presents the mean rate of energy capture at these wind-speeds, for a set point of 300 kW. To quantify the constraint upon the power set point, a target curve for the average rate of energy capture is specified. The feasible region for the power set point is then constrained such that the estimated average rates of energy capture on the target curve are always either achieved or exceeded. The target curve that is specified is plotted in Figure 5.1. The curve increases linearly from 22 m/s to 26 m/s, following the pattern of the fraction of time spent off-line that is listed in Table 5.1.

The resulting constraint upon the feasible region of the power set points is presented in Figure 5.2. The constraint rises to a set point of 100 kW in 23 m/s and to 270 kW in 24 m/s. In 25 m/s, no power set point is capable of achieving the required average power level of 225 kW, and so the turbine must remain off-line. The dotted line on

Average wind-speed $\overline{v}$ (m/s)	Estimated average fraction of time that turbine is off-line (%)	Mean rate of energy capture $\hat{\bar{e}}$ for 300 kW power set point (kW)
23	3	290
24	39	180
25	76	71
26	97	9.9

Table 5.1: Mean rate of energy capture and fraction of time off-line in high wind-speeds

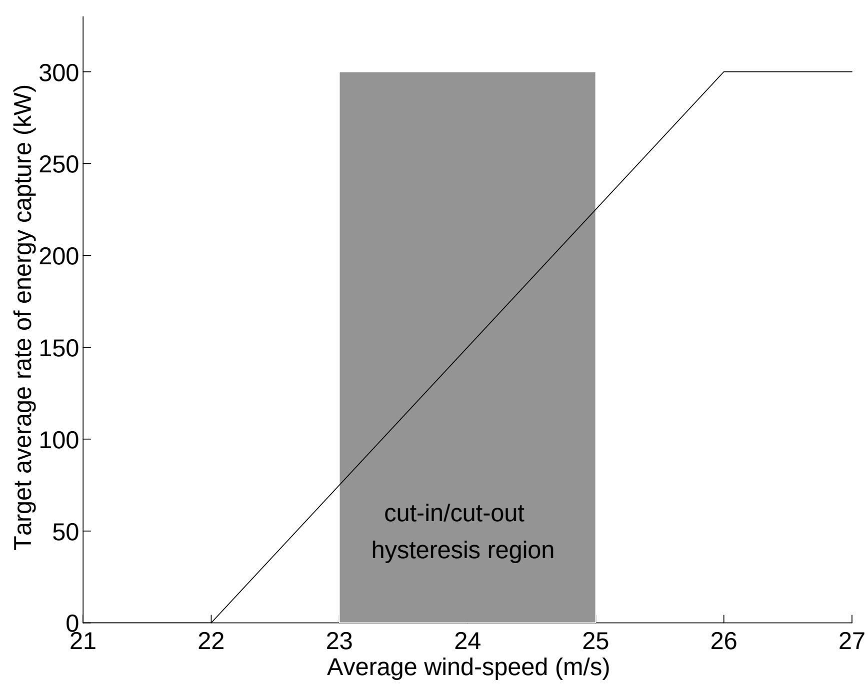

Figure 5.1: Target average rate of energy capture in high wind-speeds

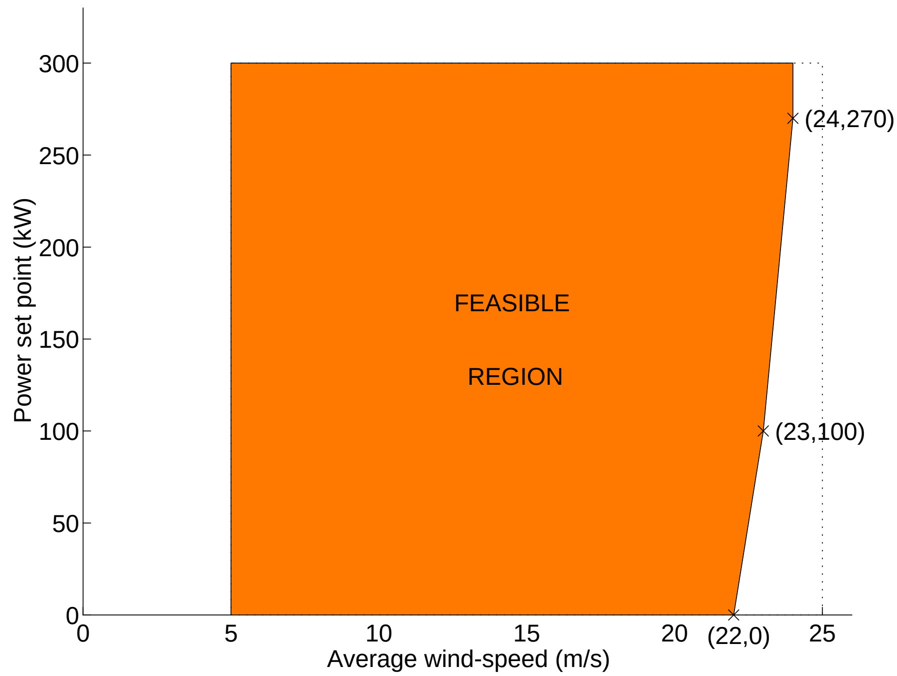

Figure 5.2: Feasible region for power set points

the figure is the original feasible region before the constraint is applied.

The set point constraint is again implemented using a simple brick wall penalty function, whereby the fitness of any power curve that contains set points outside the feasible region is set to zero. Tests were carried out which showed that the use of the simple penalty function has a negligible detrimental effect upon the performance of the genetic algorithm. For one tariff season per annum and for a constraint upon the lifetime fatigue damage of 80% times the maximum possible damage, the performance of the ideal power curves that result from searches with and without the high wind-speed constraint is as follows:

Without high wind-speed constraint: fitness (1988 £) = 24 589 p.a.

With high wind-speed constraint: fitness (1988 £) = 24 583 p.a.

Fractional difference = 0.02 %

The reason for the negligible difference between the performance with and without the constraint is that the contribution to the total cumulative energy capture and bearing damage in high wind-speeds is small.

Convergence of search

Figure 5.3 illustrates the convergence of the search algorithm for the more testing case of two tariffs per annum, with a gearbox fatigue damage safety factor A_{gear} of 0.8. The upper graph shows that the fitness of the fittest individual in the population

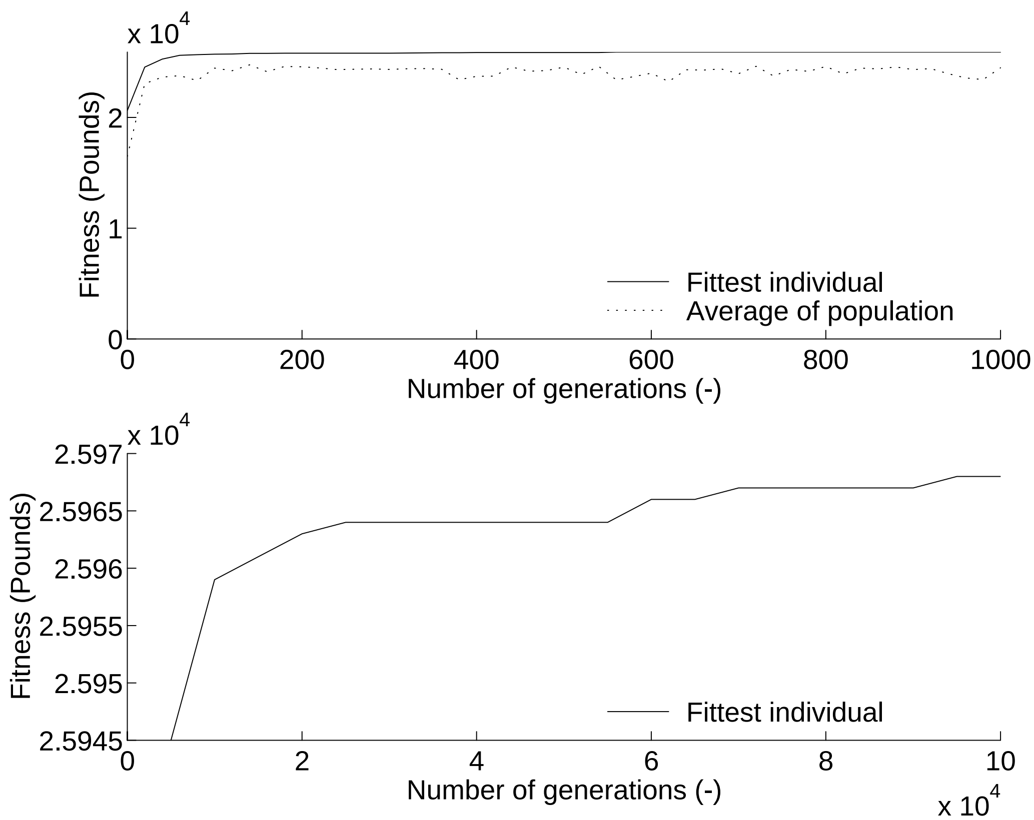

Figure 5.3: Convergence of genetic search for 'ideal' power curves

and the mean fitness of the population both achieve approximate convergence after 100 generations. The convergence of the mean fitness indicates that the 100 points in the search space that are generated by the 100 individuals in each population cluster in around the optimum, as required. The flat nature of the mean fitness for more than approximately 100 generations is due to the continual slow mutation of individuals in the population and the mating of these individuals with others. In this sense, the mutation process holds the mean fitness down. However, without occasional random mutation 'in-breeding' will take place. The population will evolve to a certain point and then not be able to evolve further, because no new gene structures will be introduced.

The lower graph in Figure 5.3 illustrates that the algorithm continues to find improved power curves for 100 000 generations.

5.2.2 Two tariff seasons per annum

Figure 5.4 presents results of the searches for the ideal power curves*, for two tariff seasons per annum and for fatigue damage reliability factors A_{gear} of 0.7, 0.8 and

*The ideal power curves usually contained one or two individual points which lay off the line of the curves, because the genetic algorithm does not explicitly search for 'smooth' power curves. To increase the fitness, these points are altered by hand to give smooth curves. The resulting changes in the fitness and fatigue damage are negligible, however, and so the inclusion of a mechanism in the search, either manual or automatic, to give smooth power curves is not necessary.

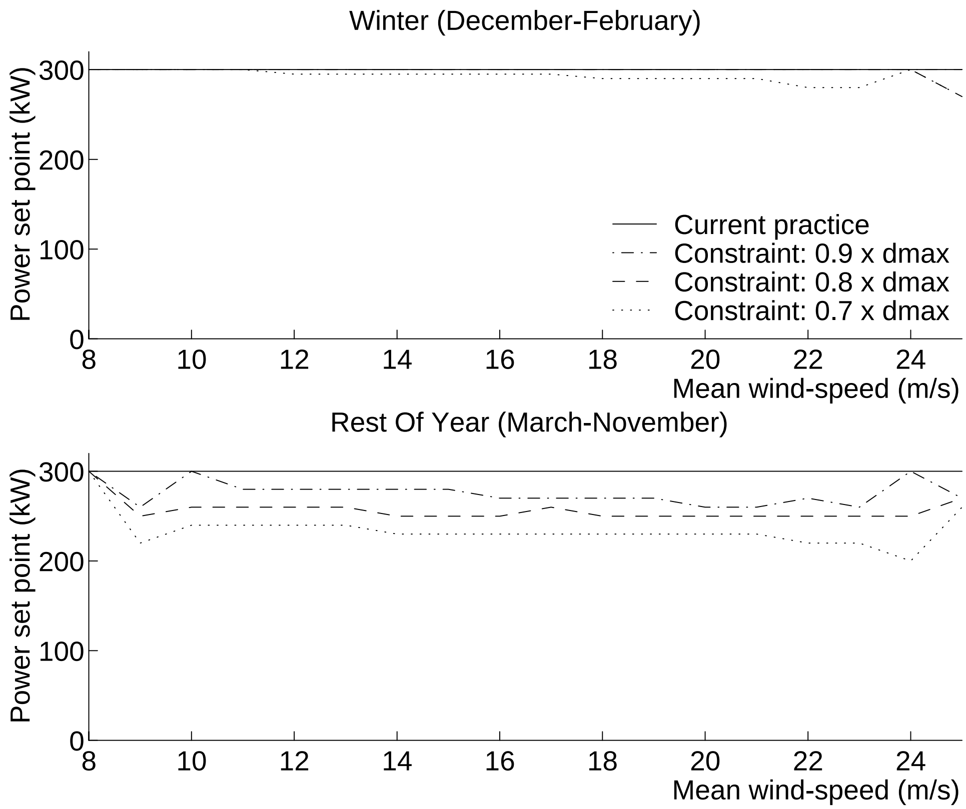

Figure 5.4: 'Ideal' power curves for 1 turbine with 2 tariffs p.a. (I = 15%)

0.9. The important features of these curves are:

- the absence of de-rating in the Winter season in any other than very high wind-speeds†, for all cases except $A_{gear} = 0.7$
- the flatness of all the ideal power curves
- the uneven nature of the curves in the Rest Of Year season for $\overline{v} \leq 9$ m/s and $\overline{v} \geq 24$ m/s

The negligible de-rating in the Winter months is due to the effect of the higher tariffs in Winter compared to in the Rest Of Year season.

The flatness of the ideal power curves is due to the use of the gearbox bearing damage in the damage constraint. This can be illustrated by considering the shapes of the surfaces of the two variables in the optimisation problem, gearbox bearing damage and energy capture, as functions of power set point and mean wind-speed. The surfaces for the complete year are repeated in Figures 5.5 and 5.6. The two surfaces have very similar shapes, thereby giving no obvious region in the operating range where de-rating has a particularly strong effect on bearing damage.

The uneven nature of the ideal power curves in the Rest Of Year season for $\overline{v} \leq 9$ m/s and $\overline{v} \geq 24$ m/s has two causes. Firstly, the performance of the search algorithm is insensitive to changes in power set point in the region around P_{rated} at 8 m/s and, to a lesser extent, 9 m/s. The curves of fatigue damage and energy capture as functions of power set point are relatively flat in this region, as illustrated by Figures 5.5 and 5.6. Secondly, the use of a fixed-point coding in the genetic algorithm causes the power set point to be raised above the level indicated by the underlying trend at either $\overline{v} = 24$ m/s or at $\overline{v} = 25$ m/s, to bring the annual fatigue damage closer

†All three 'ideal' curves fall to 270 kW at 25 m/s in the Winter season.

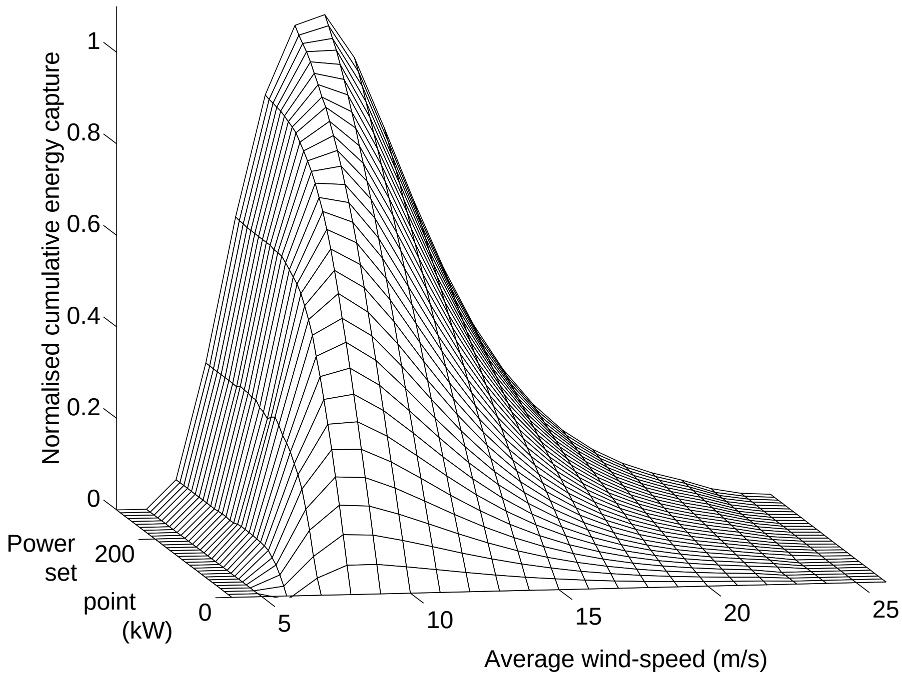

Figure 5.5: Normalised annual energy capture for 1 turbine ($\overline{v}_{year} = 7.6$ m/s; $k_w = 1.9$; $I = 15\%$)

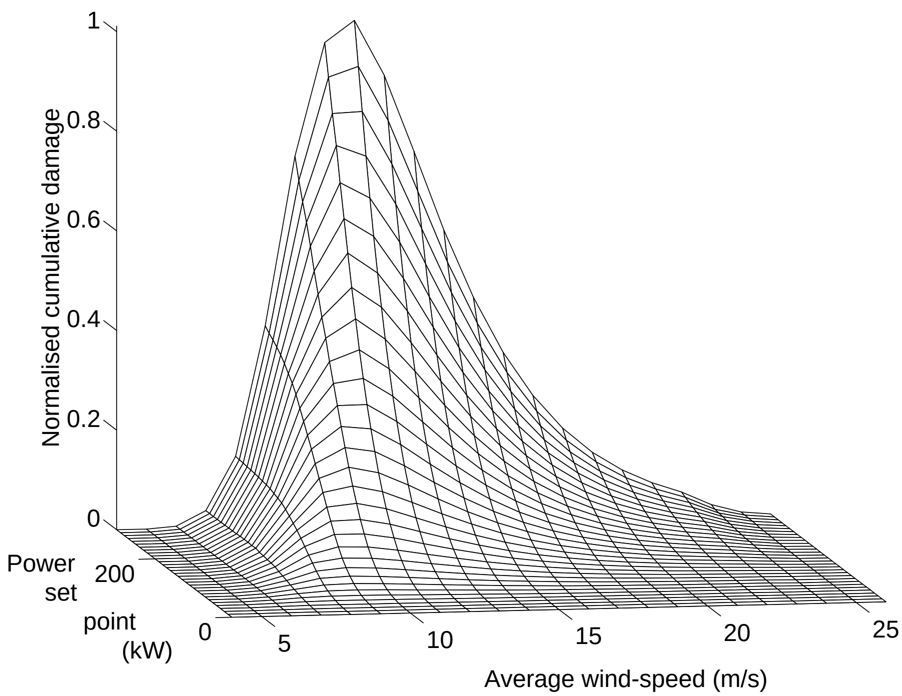

Figure 5.6: Normalised annual gearbox bearing fatigue damage for 1 turbine ($\overline{v}_{year} = 7.6$ m/s; $k_w = 1.9$; $I = 15\%$)

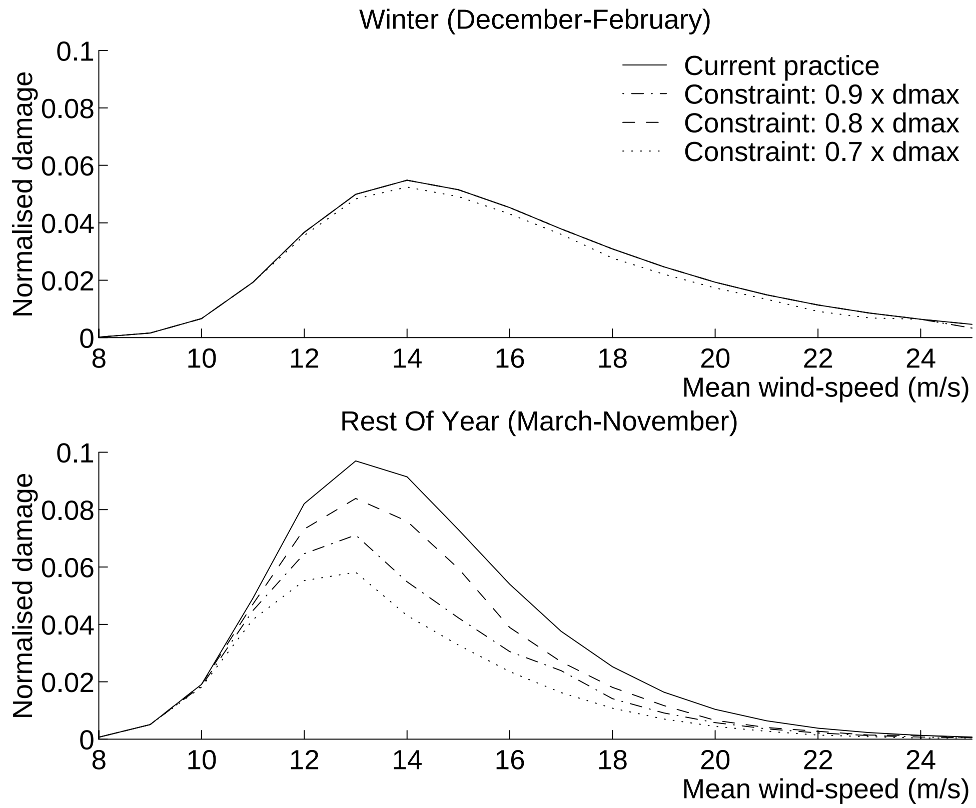

Figure 5.7: Distributions of gearbox bearing fatigue damage for 'ideal' power curves and for current practice, with 2 tariffs p.a. (1 turbine; I = 15%)

to its constraining value. This does not occur in other wind-speeds because of the relatively very small levels of cumulative damage that are found at $\overline{v} = 24$ m/s and $\overline{v} = 25$ m/s. If continuous functions had been used in the optimisation, the fluctuations in set point at high wind-speeds would not have occurred.

The distributions of the gearbox bearing damage and the financial income are presented in Figures 5.7 and 5.8 respectively, for the case of $P_{sp} = P_{rated}$ and for the ideal power curves. These figures indicate the importance of the mid-range wind-speeds and the relative unimportance of the low and high wind-speeds. The figures

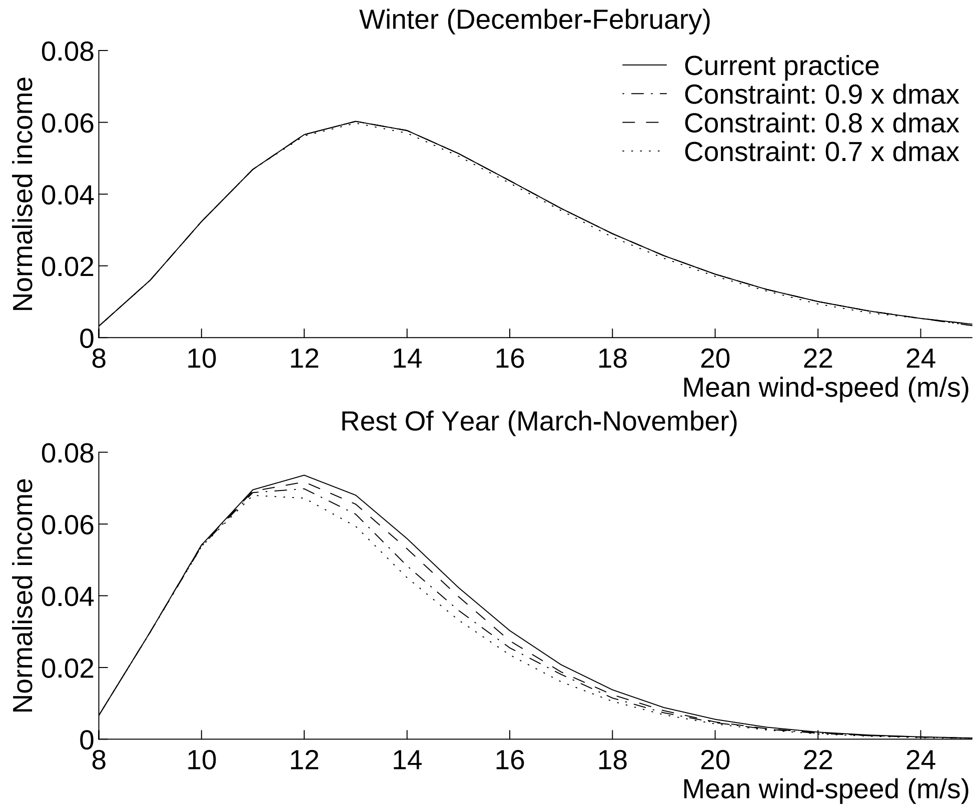

Figure 5.8: Distributions of income for 'ideal' power curves and for current practice, with 2 tariffs p.a. (1 turbine; I = 15%)

also provide some evidence of the convergence of the searches for the ideal power curves, as the distributions of damage and income in the Rest Of Year season form a uniform pattern. The distributions of damage and income all increase and decrease monotonically on both sides of their respective maxima, and at all wind-speeds the damage and income decreases monotonically as A_{gear} decreases.

5.2.3 One tariff season per annum

Figure 5.9 presents the ideal power curves for one tariff season per annum. The

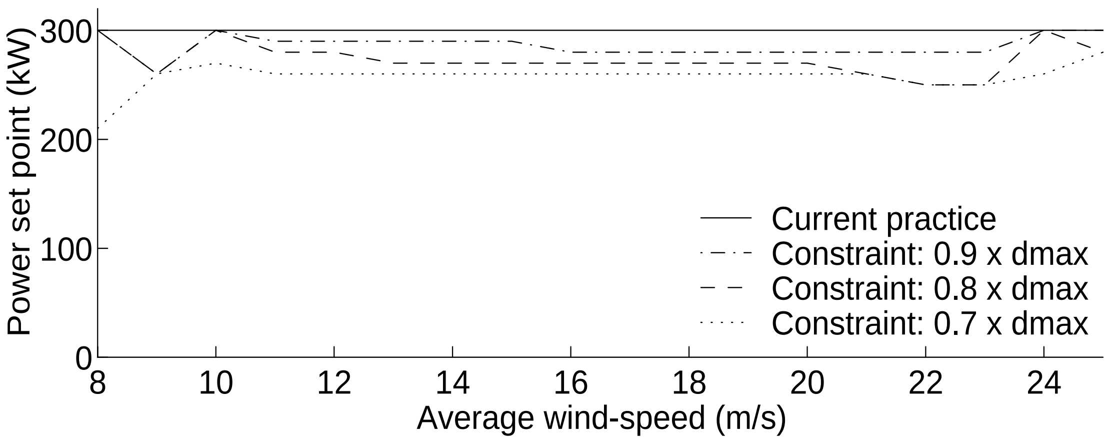

Figure 5.9: 'Ideal' power curves for 1 tariff p.a. (1 turbine; I = 15%)

significant features are again the general flatness of the curves across the majority of the operating range and the uneven nature of the curves for $\overline{v} \leq 9$ m/s and $\overline{v} \geq 24$ m/s. The same mechanisms cause these effects as in the case of two tariff seasons per annum. The distributions of gearbox bearing fatigue damage and energy capture are presented in Figure 5.10. The curves have similar shapes to those for two tariff seasons per annum, again providing evidence for the reliability of the search results.

5.2.4 Comparison of performance

A comparison of the performance of the 2-tariff and 1-tariff ideal power curves is presented in Table 5.2. The annual income increases when multiple tariffs are taken into account in the operating strategy, for all three constraints upon fatigue damage. Increases in income are achieved even though the underlying optimisation problem

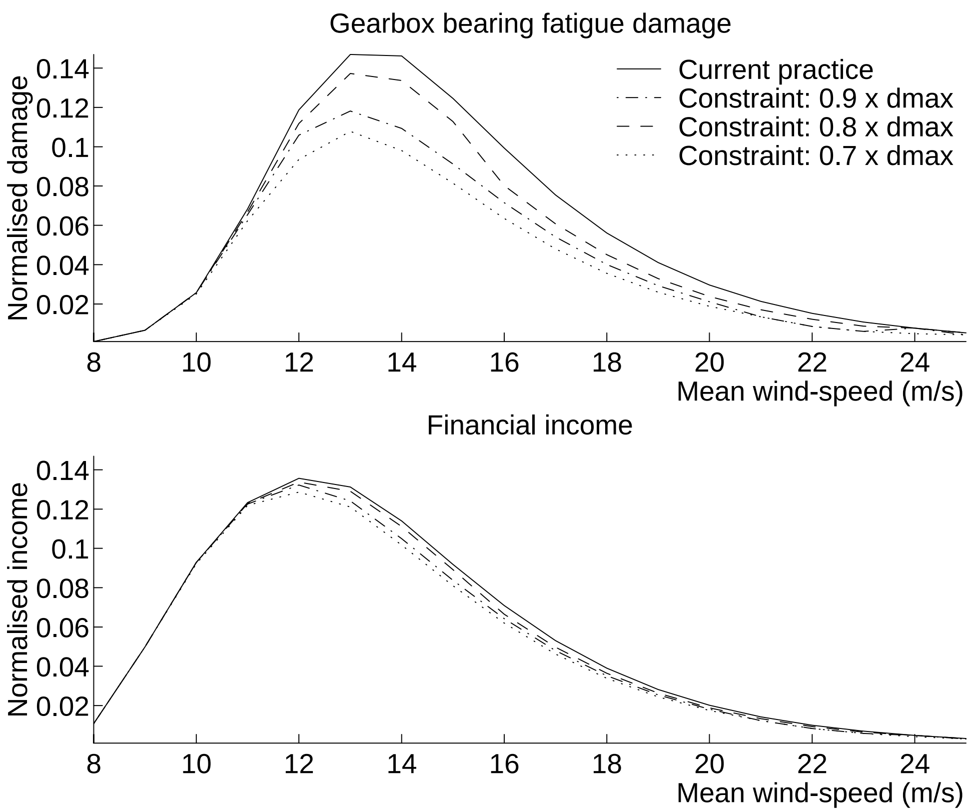

Figure 5.10: Distributions of gearbox bearing fatigue damage and income for 'ideal' power curves and for current practice, with 1 tariff p.a. (1 turbine; $I = 15\%$)

1 tariff			
Fraction of maximum possible income	Fatigue damage constraint		
	$70\% \times \hat{d}_{y_{max}}$	$80\% \times \hat{d}_{y_{max}}$	$90\% \times \hat{d}_{y_{max}}$
$\hat{m}_y/\hat{m}_{y_{max}}$ (%)	92.0	94.9	97.6
2 tariffs Seasons: December–February; March–November Tariff ratio: 1.7:1			
Fraction of maximum possible income	Fatigue damage constraint		
	$70\% \times \hat{d}_{y_{max}}$	$80\% \times \hat{d}_{y_{max}}$	$90\% \times \hat{d}_{y_{max}}$
$\hat{m}_y/\hat{m}_{y_{max}}$ (%)	92.6	95.6	98.0

$\hat{d}_{y_{max}}$ = maximum possible annual damage, for which $P_{sp} = P_{rated}$ throughout

$\hat{m}_y$ = annual financial income for ideal power curve

$\hat{m}_{y_{max}}$ = maximum possible annual income, for which $P_{sp} = P_{rated}$ throughout

Table 5.2: Comparison of performance of 'ideal' power curves for single and multiple tariff seasons per annum (1 turbine; $I = 15\%$)

presents apparently little room for improvement: firstly, the gearbox damage and energy capture functions have similar shapes; secondly, the extremes of the tariff structure are averaged out when reducing the raw tariffs to two per annum. The results in Table 5.2 therefore indicate that control strategies which include tariff variations in their design are likely to possess considerable potential for improving performance.

5.3 Simple cost function for non-interacting turbines

The cost functions that are used for non-interacting turbines are given in Equations 2.14 and 2.15 (page 45), and the cost functions' weighting parameters are evaluated using the method described in Section 2.4.2. The same discrete values of wind-speed and power set point are used as in the searches for the ideal power curves. The searches for the weighting parameters utilise optimisation functions in a proprietary software package (MATLAB, 1992). For one tariff season per annum, with only one independent variable, a golden section search with parabolic interpolation is used (function '`fmin`'). For two tariff seasons per annum, with two independent variables, a simplex search is used (function '`fmins`').

The two search functions are for unconstrained minimisation. The function to be minimised is $-\hat{m}_y$, and a penalty function is applied to prevent the annual fatigue damage from exceeding its constraint. The penalty function reduces the income in inverse proportion to the amount by which the fatigue damage violates the constraint:

$$\text{if } \hat{d}_y(\lambda_w) \leq d_{y_{con}} \quad \text{then} \quad m'_y(\lambda_w) = \hat{m}_y(\lambda_w)$$
$$\text{else} \quad m'_y(\lambda_w) = \frac{\hat{m}_y(\lambda_w)}{1 + a_{pen}\frac{(\hat{d}_y(\lambda_w) - d_{ycon})}{d_{ycon}}}$$

where

$$d_{y_{con}} = \frac{1}{20} A_{gear} \hat{d}_{gear_{max}} = \text{fatigue damage constraint}$$

$m'_y(\lambda_w)$ = penalised annual income

a_{pen} = penalty function constant

Tests were carried out which gave the values of a_{pen} for which the damage is equal to its constraining value. For an A_{gear} of 0.8, the following values of a_{pen} result:

$$1 \text{ tariff:} \quad a_{pen} = 0.35 \qquad 2 \text{ tariffs:} \quad a_{pen} = 0.5$$

The weighting parameters are the independent variables in this optimisation and therefore it is not possible to apply the 300 kW constraint to the power set points. It is necessary to identify whether any set points exceed 300 kW for a particular value of λ_w and consequently $\hat{\dot{d}}$ and $\hat{\dot{e}}$ are evaluated for P_{sp} = 310 kW. A second penalty function is then used to penalise the income of any power curve which contains any power set points that exceed 300 kW. The penalty function reduces the income by a constant factor for every set point that exceeds 300 kW:

$$\begin{aligned} &\text{for each } P_{sp} \leq P_{rated}\ , \quad m'_y(\lambda_w) = \hat{m}_y(\lambda_w) \\ &\text{for each } P_{sp} > P_{rated}\ , \quad m'_y(\lambda_w) = \frac{\hat{m}_y(\lambda_w)}{b_{pen}} \end{aligned}$$

b_{pen} is the penalty function constant. Suitable values for b_{pen} lie in the range 1–1.005.

Initial search results gave power curves for which the annual bearing damage is equal to its constraint, as required. However, at least one set point in each power curve exceeds 300 kW in either low or high wind-speeds. Significant increases in λ_w are necessary to ensure that no set points exceed 300 kW, but then greater weight is attached to the damage term in the cost function, the power curves move to lower levels, and the annual damage decreases. The power curves are lowered to such an extent that the income is reduced to a level significantly below that of the ideal power curves. As a result, two regions are excluded from the searches for the weighting parameters: $\overline{v} \leq v_{rated}$ and $\overline{v} \geq 24$ m/s.

In the region $\overline{v} \leq v_{rated}$, it is assumed that no on-line optimising control takes place and the set points are fixed at rated power.

In the region $\overline{v} \geq 24$ m/s, the set points are assumed to be equal to those for $\overline{v} = 23$ m/s. In practice, a supervisory controller is unlikely to attempt on-line hill-climbing control in this region, as the turbines cut out due to high wind-speeds for part of the time and the estimation of $\hat{\overline{d}}$ and $\hat{\overline{e}}$ is difficult. For $\overline{v} \geq 24$ m/s, a sensible strategy for an on-line controller is to maintain the set point at the 23 m/s level. This set point will have already been optimised for the prevailing level of turbulence and will provide some protection against excessive fatigue damage for $\overline{v} \geq 24$ m/s.

The full implementation of the optimisation, for the case of one tariff season per annum, is as follows:

$$m_{y_{opt}}(\lambda_w) = max\left(m'_y(\lambda_w)\right)$$

such that:

$$\begin{array}{lll}
\text{(i)} & \text{if } \hat{d}_y(\lambda_w) \leq d_{y_{con}} & \text{then} \quad m'_y(\lambda_w) = \hat{m}_y(\lambda_w) \\
 & & \text{else} \quad m'_y(\lambda_w) = \dfrac{\hat{m}_y(\lambda_w)}{1 + a_{pen}\frac{(\hat{d}_y(\lambda_w) - d_{y_{con}})}{d_{y_{con}}}} \\
\text{(ii)} & \text{for each } P_{sp} \leq P_{rated}\ , \quad m'_y(\lambda_w) = \hat{m}_y(\lambda_w) & \\
 & \text{for each } P_{sp} > P_{rated}\ , \quad m'_y(\lambda_w) = \dfrac{\hat{m}_y(\lambda_w)}{b_{pen}} & \\
\text{(iii)} & \text{if } \overline{v} \leq v_{rated} & \text{then} \quad P_{sp} = P_{rated} \\
 & \text{else if } v_{rated} < \overline{v} < 24 \text{ m/s} & \text{then} \quad P_{sp} = P_{sp}(\lambda_w) \\
 & \text{else} & \qquad\quad\ P_{sp} = P_{sp}(\lambda_w, 23 \text{ m/s})
\end{array}$$

$m_{y_{opt}}(\lambda_w)$ = optimal (maximum) annual income for weighting parameter λ_w.

For two tariffs per annum, λ_w in the above equations becomes $\lambda_w(i_s)$, a function of tariff season i_s.

5.3.1 Two tariff seasons per annum

Figure 5.11 presents the annual financial income, and power curves, that result from the use of cost functions with weighting parameters optimised by the above procedure, for two seasons per annum and for an A_{gear} of 80%. The ideal power curves and corresponding annual income are also plotted for comparison. The simple cost function successfully provides almost the same level of performance as the ideal

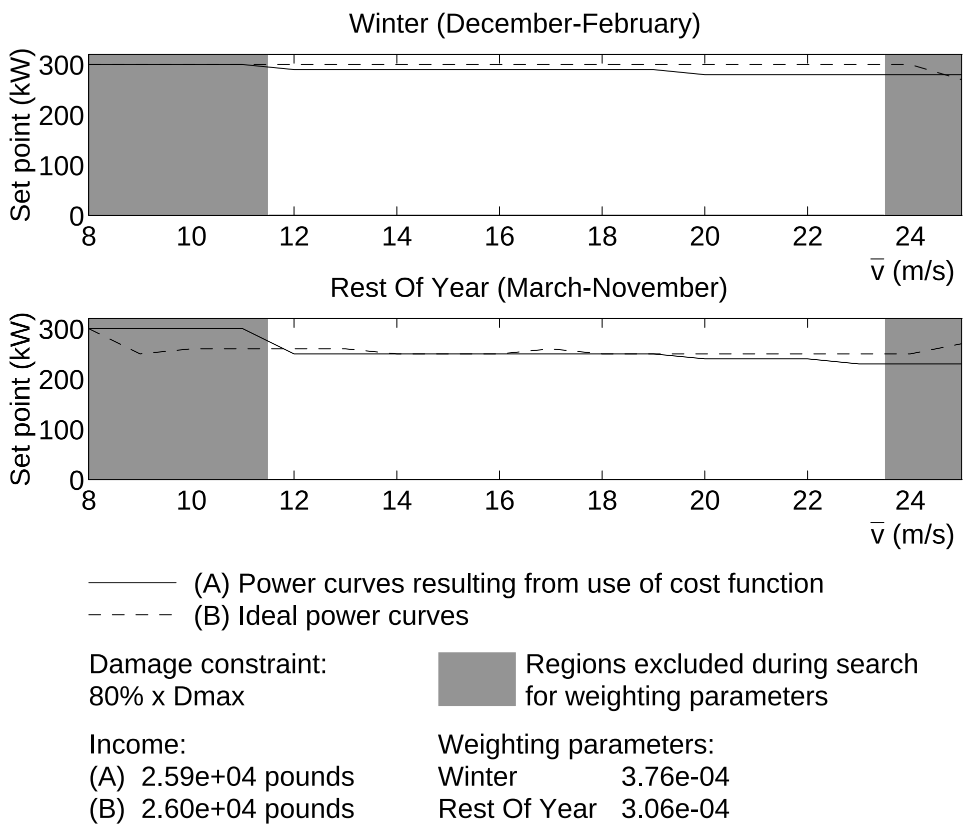

Figure 5.11: Comparison of performance of simple cost function with 'ideal' operation, for 2 tariffs p.a. and 2 weighting parameters (1 turbine; I = 15%)

power curves, with an annual income that is only 0.4% lower than the ideal. In addition, the two pairs of power curves are very similar. For $A_{gear} = 70\%$ and $A_{gear} = 90\%$, the income also differs negligibly from that for the ideal power curves.

The golden section and simplex algorithms are less robust than a genetic algorithm with respect to possible convergence to local minima, as genetic algorithms are inherently 'noisy' and conduct a relatively broad search across the search space. However, the close agreement between the performance of the simple cost function and the performance of the ideal power curves for two tariffs per annum shows that the simplex technique provides satisfactory convergence when searching for the cost function's weighting parameters.

The same weighting parameter ($\lambda_{w_{year}}$) can be used in both seasons in the 2-tariff case. The cost function is equal to Equation 2.14 with $\lambda_w(1) = \lambda_w(2) = \lambda_{w_{year}}$:

$$J = \int_{t-T_c}^{t} \left[\lambda_{w_{year}} \, \dot{d}(t) - \dot{m}(t) \right] dt$$

The resulting annual income and power curves for $A_{gear} = 80\%$ are as shown in Figure 5.12. The performance is again good, as the annual income differs negligibly (by 0.3%) from that for the 'ideal' case. A similar result is obtained for $A_{gear} = 70\%$, however for $A_{gear} = 90\%$ the annual income is 3% lower than the ideal.

5.3.2 One tariff season per annum

Figure 5.13 presents the annual financial income, and power curve, for one tariff season per annum and for an A_{gear} of 80%. As in the two-tariff case, the use of the cost function successfully results in a similar level of performance as is obtained when using the ideal power curve. Also, for $A_{gear} = 70\%$ and $A_{gear} = 90\%$ the income

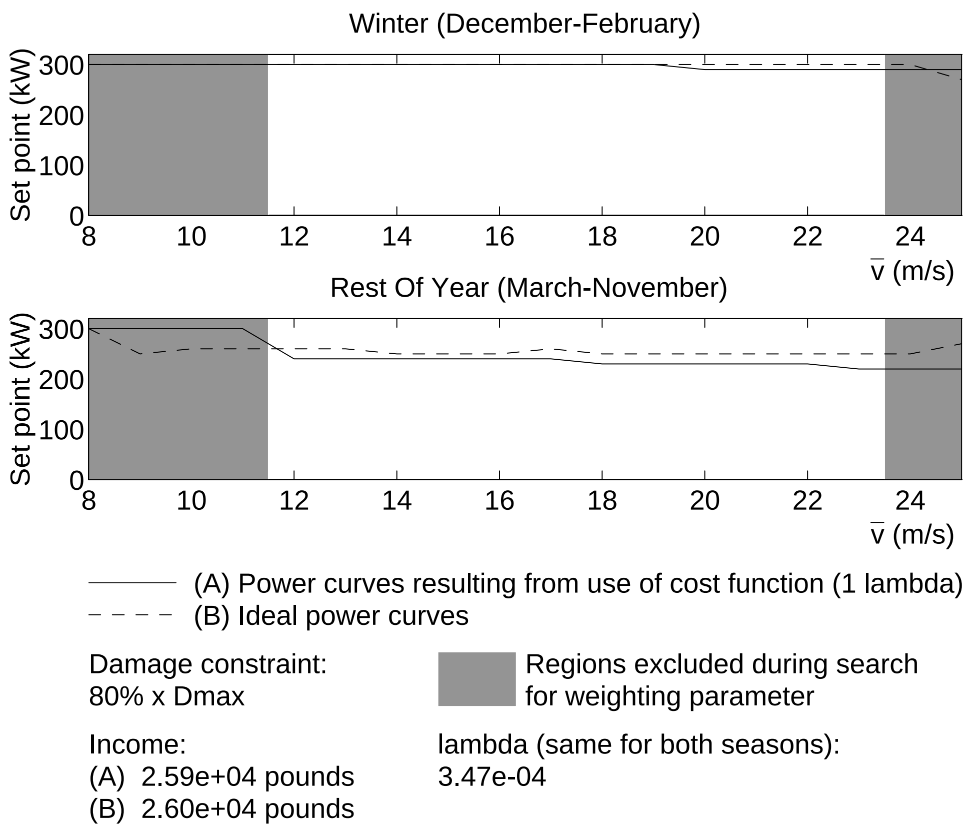

Figure 5.12: Comparison of performance of simple cost function with 'ideal' operation, for 2 tariffs p.a. and 1 weighting parameter (1 turbine; $I = 15\%$)

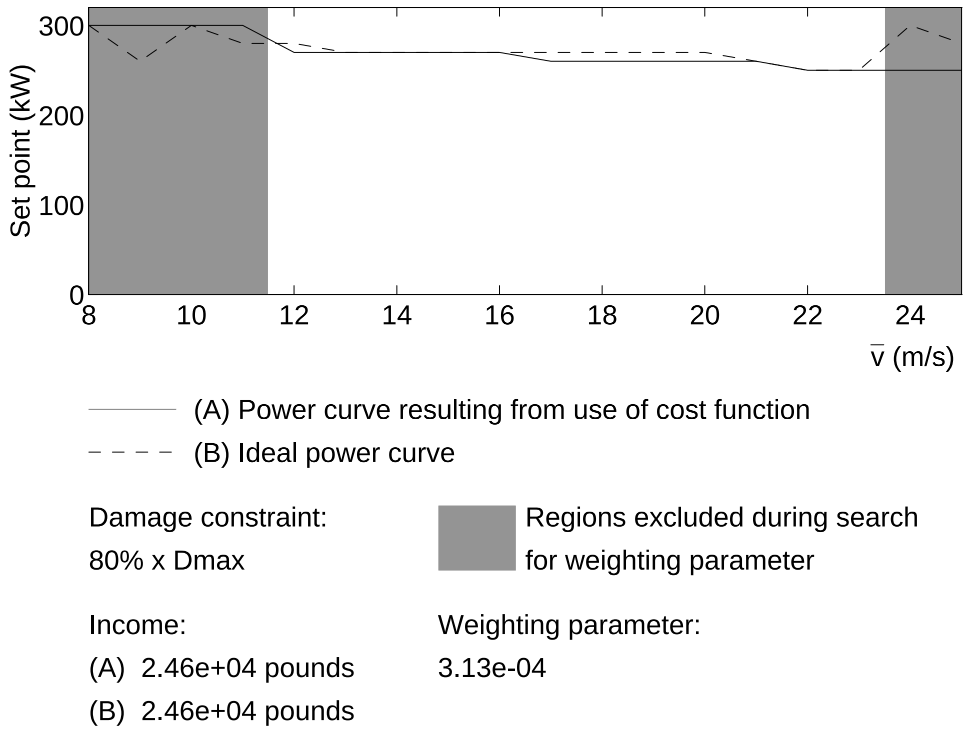

Figure 5.13: Comparison of performance of simple cost function with 'ideal' operation, for 1 tariff p.a. (1 turbine; I = 15%)

levels are the same as for the ideal power curves. The close agreement between the performance of the simple cost function and the performance of the ideal power curves shows that the golden section technique provides satisfactory convergence when searching for the cost function's weighting parameter in the case of one tariff season per annum.

The shape of the cost function that gives the power curve of Figure 5.13 is illustrated in Figures 5.14 and 5.15. The general view in Figure 5.14 shows that, at any wind-speed, the cost only has one minimum as a function of P_{sp}. The cost varies widely

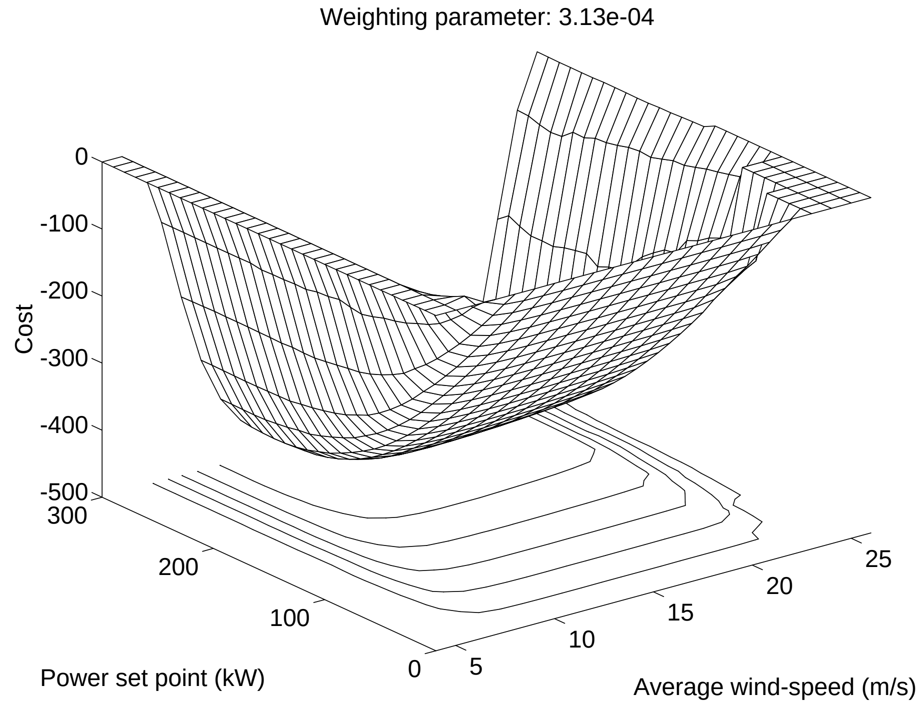

Figure 5.14: General view of simple cost function for 1 tariff per annum and $\lambda_w = 3.13 \times 10^4$ *(1 turbine; I = 15%)*

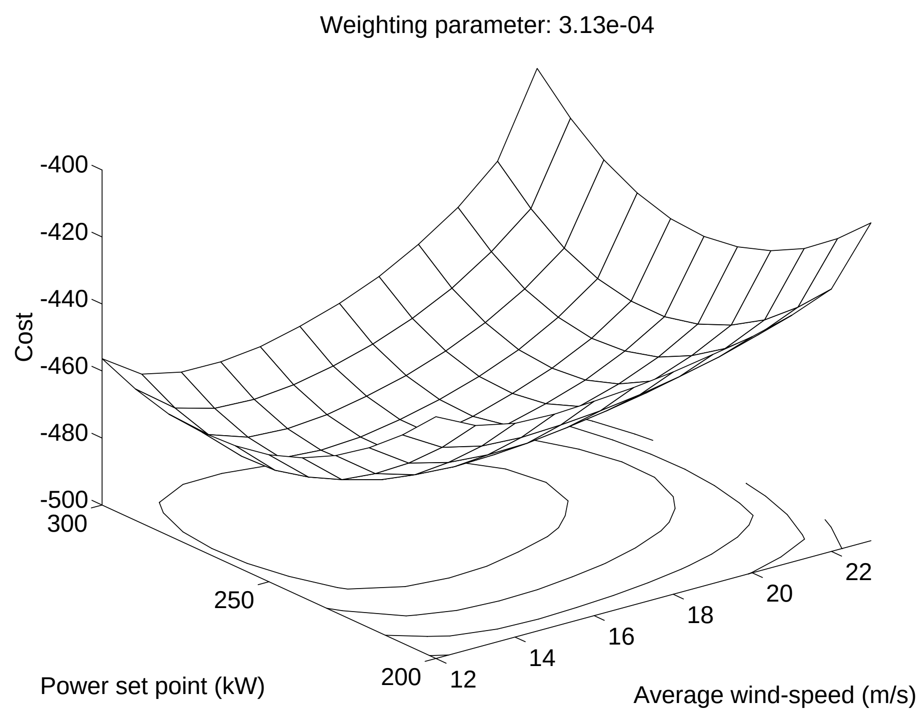

Figure 5.15: Detail of simple cost function for 1 tariff p.a. and $\lambda_w = 3.13 \times 10^4$ *(1 turbine; I = 15%)*

as a function of both $\overline{v}$ and P_{sp} but is relatively flat in the general area where an on-line controller would be used, that is, in the region 200 kW $\leq P_{sp} \leq$ 300 kW. This is illustrated by the more detailed view in Figure 5.15, in which the line of minimum cost for all values of $\overline{v}$ can be seen, at approximately P_{sp}= 270 kW. The flatness of the cost function in the region 200 kW $\leq P_{sp} \leq$ 300 kW mainly results from the similarity of the shapes of the gearbox bearing damage and energy capture functions.

5.3.3 On-site tuning

It has been shown that control based upon the proposed simple cost function can give similar performance to the 'ideal', for both one and two tariff seasons per annum. The weighting parameters have been evaluated using simulation, for one particular design of turbine and for one particular cost function. In practice, it is likely that on-site tuning of the weighting parameter would be employed directly.

For one tariff season per annum, a possible approach to direct on-site tuning would be as follows. λ_w would be adjusted only when the mean wind-speed is close to the rated wind-speed and the turbulence intensity has a moderate value. The target value of λ_w would be that which, under these wind conditions, causes the controller to select a set point close to P_{rated}. A large value for λ_w would be used initially, as this would give a large weight to the term for fatigue damage in the cost function and would therefore cause significant de-rating to take place. λ_w would then be progressively reduced, thereby causing P_{sp} to increase, until P_{sp} is close to P_{rated}. In moderately turbulent conditions, the resulting value of λ_w should give no de-rating when $\overline{v} < v_{rated}$ and would give slight de-rating when $\overline{v} > v_{rated}$.

For two tariff seasons per annum, it is likely that the same value of λ_w would be used in both seasons. Exactly the same procedure could be used as for the case of one tariff per annum, except that in this case the tuning would take place in the Winter season.

5.4 Simple cost function for interacting turbines

The cost function for a single non-interacting turbine is extended to multiple interacting turbines, as described in Section 2.5. The cost functions for interacting turbines are given in Equations 2.16 and 2.17 (page 48). The simulated windfarm of three turbines in a row is again considered, with 'small' (3D/3D) inter-turbine spacings. The same discrete values of wind-speed and power set point are used as in the searches both for the ideal power curves and for the weighting parameters in the single turbine's cost function.

The case of a single tariff season per annum is investigated, with $\lambda_w = 2.49 \times 10^{-4}$. This value of λ_w gives a small reduction of 5% in the annual bearing damage of a non-interacting turbine and Figure 5.16 gives the resulting power curve and annual income. The annual income is reduced by only 1.3%. It is likely that a conservative value for λ_w such as this would be used in practice, so that in moderate conditions (moderate turbulence intensity, no wake effects) there is a minimal loss in revenue. The positive effects of the supervisory controller are then experienced under unusual or difficult conditions. The regions $\overline{v} \leq v_{rated}$ and $\overline{v} \geq 24$ m/s are again excluded from the search for the weighting parameter in the single turbine's cost function. However, when the cost function is used with the interacting turbines, it is assumed that on-line hill-climbing does take place when $\overline{v} \leq v_{rated}$ and $\overline{v} \geq 24$ m/s.

Off-line searches are conducted to find the power curves for the interacting turbines that result from the use of the cost function (see Section 2.5). To carry out a

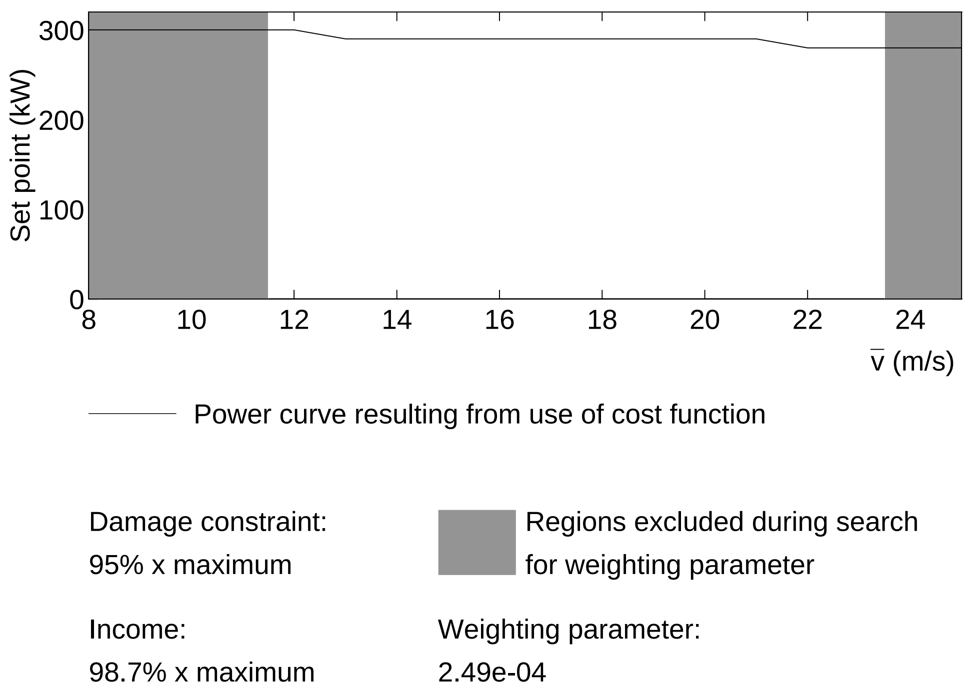

Figure 5.16: Power curve resulting from use of cost function for 1 tariff p.a. with $\lambda_w = 2.49 \times 10^{-4}$. *(1 turbine; $I = 15\%$)*

search for the power set points at a particular mean wind-speed, the simulation and optimisation software is run simultaneously. Parameter sweeps are carried out in the simulation for each set of three power set points that are generated by the optimisation algorithm. The same random number seed is used for each parameter sweep, and the data transfer between the simulation and optimisation software is carried out automatically via intermediate files, with hand-shaking to give reliable transfers. The search for the power set points utilises an optimisation function in the same software package as is used for the search for the single turbine's weighting parameter (MATLAB, 1992)[Optimisation Toolbox]. A constrained minimisation algorithm is used which is based upon a sequential quadratic programming method, whereby a quadratic programming subproblem is solved at each iteration (function '`constr`'). The algorithm is highly efficient in this application, generally requiring approximately 40 steps to achieve convergence.

The power set point curves for the three interacting turbines that result from the searches, with $I = 15\%$ and 8 m/s $\leq \overline{v} \leq 24$ m/s, are presented in Figure 5.17. The curves may be considered as consisting of two distinct parts:

- For $\overline{v} > 14$ m/s, the set point curves are similar to those for 3 turbines operating in isolation, as can be seen by comparing the curves with that for a single turbine in Figure 5.16. Although I increases at the downstream turbines, the insensitivity of both the rate of bearing damage and rate of energy capture to I causes little difference in the 3 turbines' power curves.
- For $\overline{v} < 14$ m/s, the set point curves are substantially different to those for a single turbine. The set point of turbine 1 is heavily reduced, and the set point of turbine 2 is also reduced. The set point of turbine 3 is maintained at its maximum throughout.

The effect of the three power curves on the rates of energy capture and fatigue damage is illustrated in Figure 5.18. This figure presents the changes in the total rate of accumulation of the energy capture and fatigue damage across the three turbines as a fraction of the total rates of accumulation for $P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated}$. Figure 5.18 suggests that large improvements in the performance of the 3 turbines can be achieved.

Small changes in the total rate of energy capture (curve 'A') of between -3.7% and +0.9% occur in the lower wind-speeds, where the majority of a non-interacting turbine's annual energy capture occurs. The reductions in the rate of energy capture

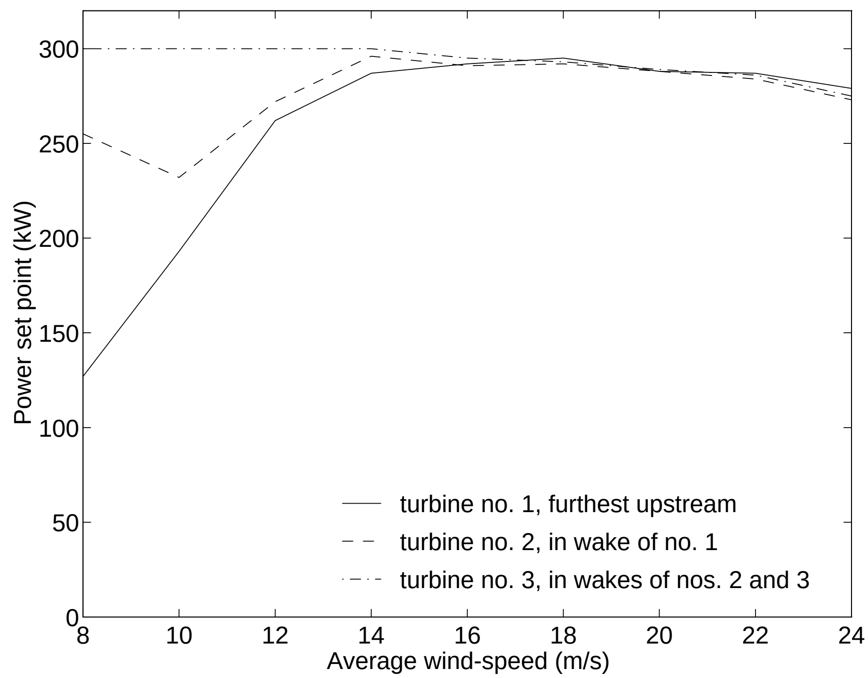

Figure 5.17: Power curves for 3 interacting turbines resulting from use of cost function for 1 tariff p.a. ($\lambda_w = 2.49 \times 10^{-4}$; $I = 15\%$)

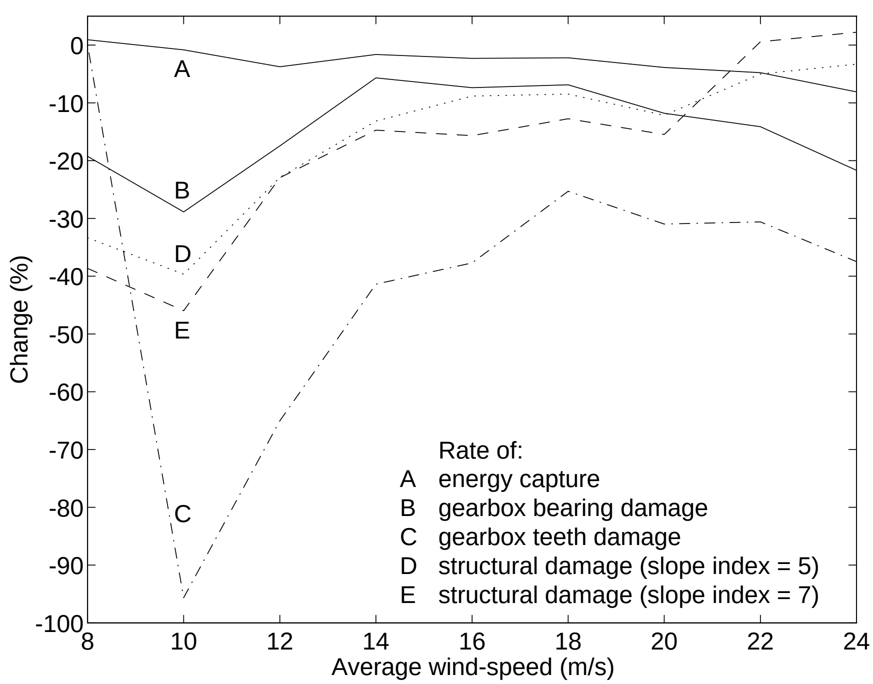

Figure 5.18: Changes in total energy capture and fatigue damage for 3 interacting turbines after optimisation of power set points. ($\lambda_w = 2.49 \times 10^{-4}$; $I = 15\%$)

in the less frequent higher wind-speeds are small. All are between 2% and 4%, except in the rare 24 m/s winds for which the reduction in energy capture increases to 8%.

The gearbox bearing damage (curve 'B') falls by between 6% and 29%, with the largest reductions in the lower wind-speeds where the majority of an isolated turbine's bearing damage is concentrated.

Large reductions in the total rate of gearbox teeth damage (curve 'C') occur, with a maximum reduction of 96% when $\overline{v} = 10$ m/s. (There is no change in 8 m/s because the damage for the $P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated}$ case is already zero.) The reductions in teeth damage are consistently much greater than those for the bearing damage, even though only the bearing damage is included in the cost function. The principal reason for the large reductions in teeth damage is the heavily non-linear relationship between teeth damage and power set point.

There are significant reductions in the general measures of structural fatigue damage (curves 'D' and 'E') in lower wind-speeds. The maximum reductions are 40% for the slope index of 5 and 46% for the slope index of 7. There are large contributions to the annual structural damage of a non-interacting turbine in two principal regions: around 10 m/s and around 23 m/s. The former contributions are lessened significantly by operation with the optimised set points whilst the latter are slightly increased.

Results are presented in Table 5.3 which illustrate the mechanism by which increases in energy capture and decreases in bearing damage are brought about in low wind-speeds. For the purposes of illustration, the data presented in Table 5.3 is for a lower turbulence intensity of 10%, as the effects are exaggerated under these conditions. The wind-speed is 10 m/s. If the rates of energy capture for the original set points (A) are compared with those for the optimised set points (B), it can be seen that the rate of energy capture drops at turbine 1 but increases both at turbine

Turbine number		1	2	3	Total (3 s.f.)
Power set point (kW)	original (A)	300	300	300	
	optimised (B)	180	190	260	
Mean rate of energy capture (kW)	A	200	100	60	366
	B	170	120	80	370
Mean rate of gearbox bearing fatigue damage ($(\mathrm{kNm})^{3.33} \times 10^5$)	A	3.6	0.9	0.4	4.91
	B	1.7	1.0	0.7	3.33

Table 5.3: Improvement in performance of 3 interacting turbines with optimised set points, for $\overline{v} = 10$ m/s and $I = 10\%$. ($\lambda_w = 2.49 \times 10^{-4}$)

2 and, particularly strongly, at turbine 3. The combined effect is to give a slight increase of 1% in the total rate of energy capture. The rate of accumulation of bearing damage falls substantially at turbine 1 and increases at turbines 2 and 3. However, the increases at turbines 2 and 3 are from relatively low levels, and the net effect is to give a substantial reduction of 32% in the total bearing damage.

The process by which the optimised power set points improve the turbines' performance can be explained by considering the mean wind-speeds which are incident upon the downstream turbines before and after optimisation. The de-rating of turbine 1 causes it to extract less power from the wind, so $\overline{v}^{(2)}$ and $\overline{v}^{(3)}$ increase. The de-rating of turbine 2 similarly assists turbine 3. The highly non-linear nature of both the rate of bearing damage and rate of energy capture with wind-speed means that it is beneficial to de-rate turbine 1, as then the operating point of turbine 1 moves away from the region of relatively high damage and $\overline{v}^{(2)}$ and $\overline{v}^{(3)}$ are boosted. The increases in $\overline{v}^{(2)}$ and $\overline{v}^{(3)}$ give rise to large increases in the rate of energy capture at turbines 2 and 3 relative to the increases in fatigue damage, due to the low wind-speeds at turbines 2 and 3.

The results in Table 5.3 also show that, with the optimised power set points, higher performance is achieved at lower turbulence intensities. This trend is verified by further results, for $I = 20\%$ (and $\overline{v} = 10$ m/s), which indicate smaller reductions in bearing damage and smaller increases in energy capture.

Table 5.4 illustrates the long-term performance of the proposed operating strategy. The annual totals for the optimised strategy are presented as fractions of the totals for the current operating practice ($P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated}$). The proposed operating strategy results in a slight drop in the annual energy capture and gives significant reductions of between 13% and 36% in the different measures of annual fatigue damage.

	Annual total following optimisation as a fraction of total for $P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated}$ (%)
Energy capture	97.9
Gearbox bearing fatigue damage	87
General structural fatigue damage ($b = 7$)	83
General structural fatigue damage ($b = 5$)	82
Gearbox teeth fatigue damage	64

Table 5.4: Annual energy capture and fatigue damage for 3 interacting turbines with optimised set points. ($\overline{v}_{year} = 7.6$ m/s; $k_w = 1.9$; $I = 15\%$)

The results in Table 5.4 for three interacting turbines can be compared with those for the 'ideal' operation of a non-interacting turbine in Table 5.2. For the interacting turbines, the bearing damage falls to 87% $\times \hat{d}_{y_{max}}$ for a fractional income $\hat{m}_y/\hat{m}_{y_{max}}$ of 97.9%. ($\hat{d}_{y_{max}}$ is the maximum possible annual damage, for which $P_{sp} = P_{rated}$ throughout; $\hat{m}_y$ is the annual financial income; $\hat{m}_{y_{max}}$ is the maximum possible annual income.) In Table 5.2, results for $\hat{m}_y/\hat{m}_{y_{max}}$ are presented for a non-interacting turbine with constraints of 80% $\times \hat{d}_{y_{max}}$ and 90% $\times \hat{d}_{y_{max}}$. If the value of 87% $\times \hat{d}_{y_{max}}$ is used to interpolate for the fraction $\hat{m}_y/\hat{m}_{y_{max}}$ for a non-interacting turbine, $\hat{m}_y/\hat{m}_{y_{max}} = 96.8\%$ results. This figure is lower than the 97.9% for the three interacting turbines, so higher performance is achieved for the interacting turbines than for non-interacting ones.

5.5 On-line control of interacting turbines

The simulated windfarm of three turbines in a row with 'small' (3D/3D) inter-turbine spacings is again considered. A very simple hill-climbing, or extremum control, scheme is implemented. The hill-climbing algorithm consists of a simple univariate search (Dixon, 1972), with logic inserted to firstly prevent the set points from exceeding their constraints and secondly stop the same set point from being specified more than once. Details of the algorithm are given in Appendix F. The set point of turbine 1 is stepped for 4 hours, followed by the set point of turbine 2 and then the set point of turbine 3. Therefore the longest time period over which the low frequency fluctuations are assumed to be negligible is 4 hours. Tests were carried out to determine a suitable time-step for the controller and a time-step of 1 hour resulted.

The operation of the simulated windfarm under extremum control is illustrated in Figures 5.19 and 5.20, for $\overline{v} = 10$ m/s and $I = 15\%$. The variation of the power set points is presented in Figure 5.19(a). The numbers "1", "2" and "3" on the

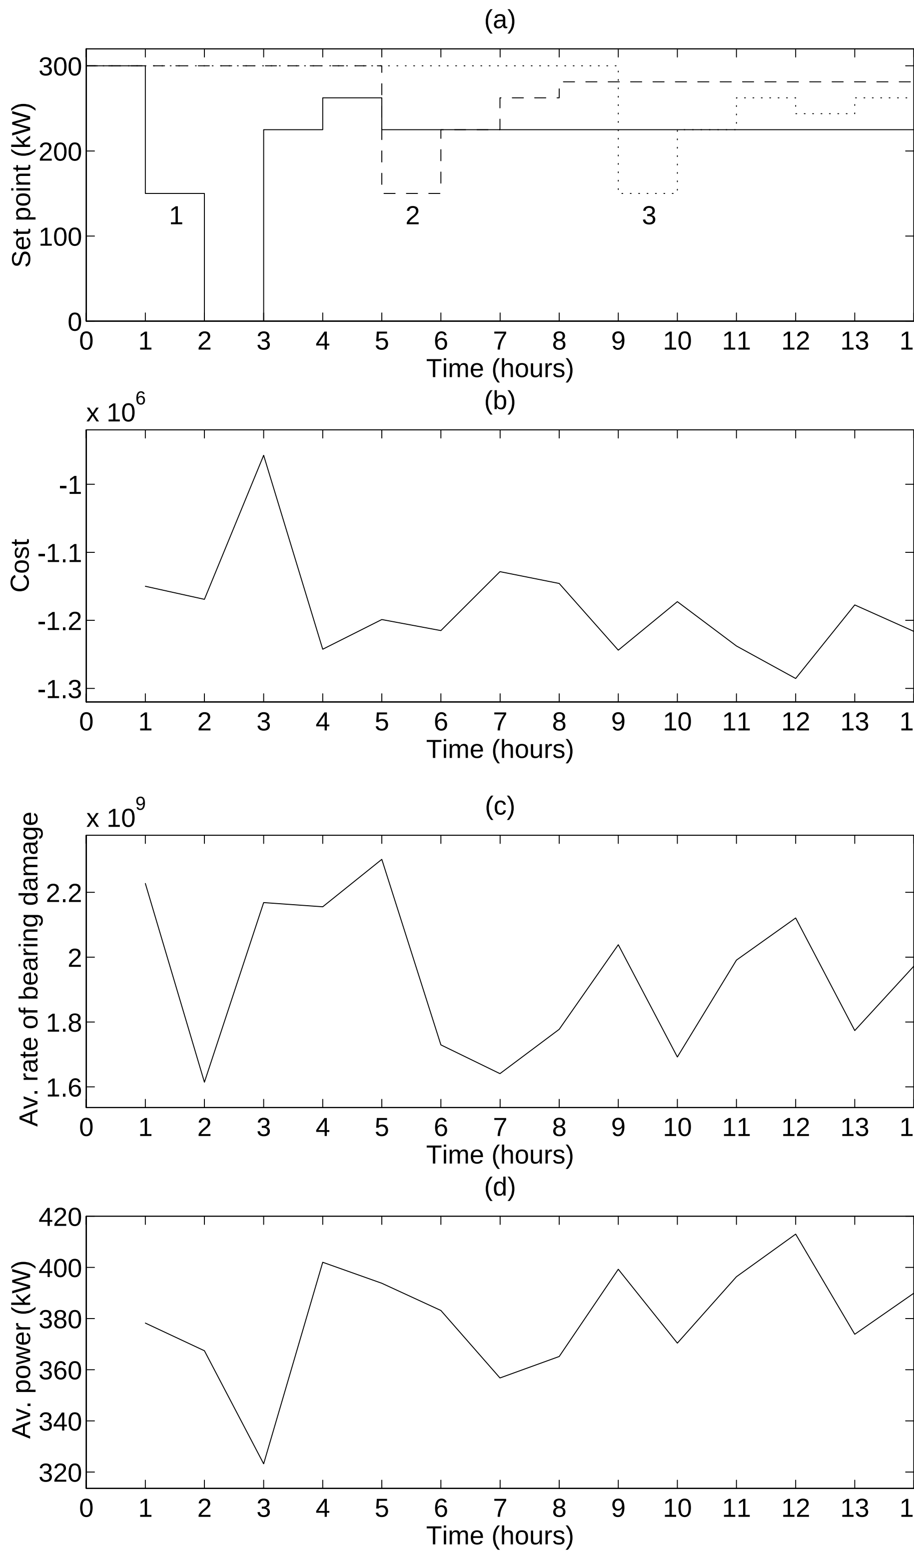

Figure 5.19: On-line extremum control of 3 interacting wind turbines. ($\overline{v} = 10$ m/s; $I = 15\%$; $\lambda_w = 2.49 \times 10^{-4}$)

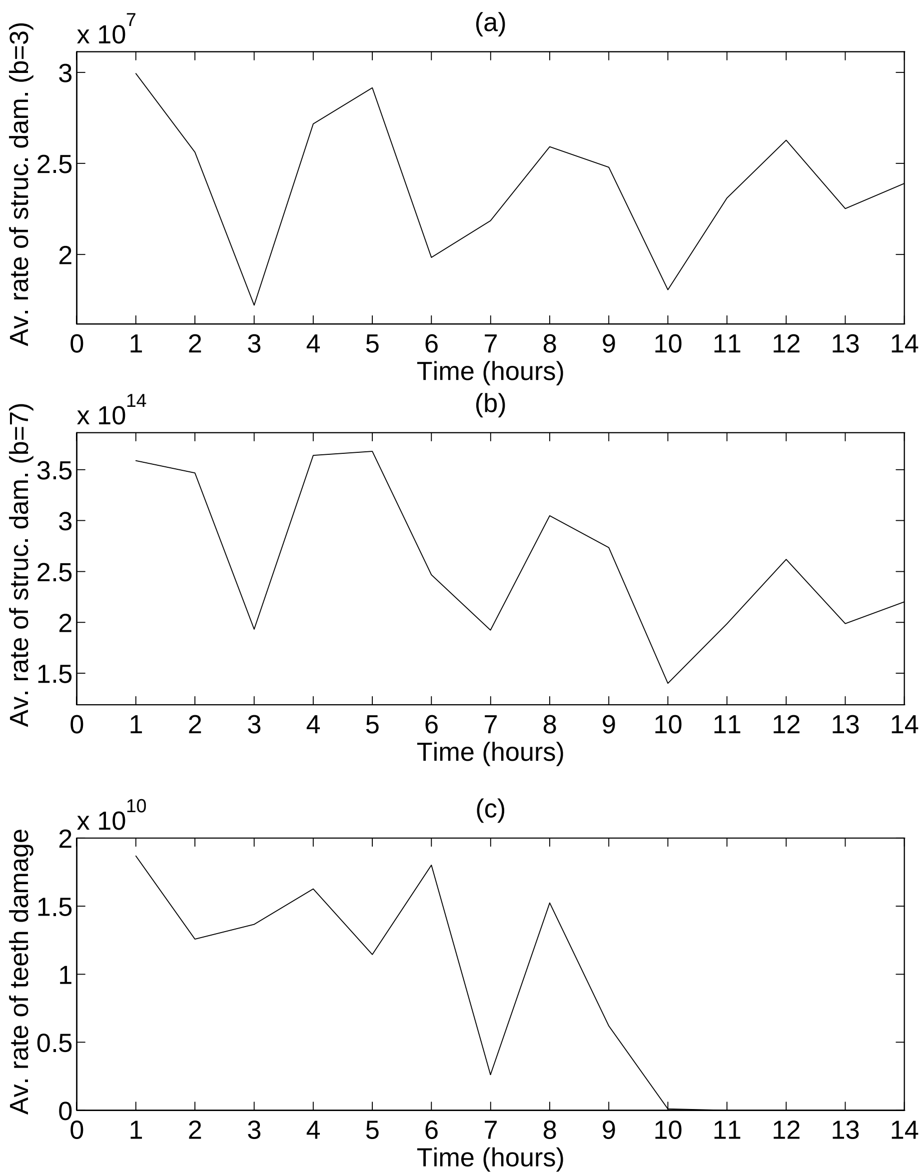

Figure 5.20: Effect upon fatigue damage of on-line extremum control of 3 interacting wind turbines. ($\overline{v}$ = 10 m/s; I = 15%; $\lambda_w = 2.49 \times 10^{-4}$)

graph are the numbers of the turbines in the row. The starting point for the control is $P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated}$, and 4 steps of 1 hour each are carried out for each set point, with a final 13th step to adjust $P_{sp}^{(3)}$ to its most successful value so far. The simple control algorithm is stable and convergent, and the power set points at the end of the 14 hours of operation are presented in Table 5.5. The set points

Optimisation of power set points carried out:	Power set point (kW)		
	Turbine 1	Turbine 2	Turbine 3
(i) on-line	230	280	260
(ii) off-line	190	230	300

Table 5.5: Comparison of power set points for 3 interacting turbines resulting from on-line and off-line optimisation. ($\overline{v}$ = 10 m/s; I = 15%; $\lambda_w = 2.49 \times 10^{-4}$)

that result from the off-line search of Section 5.4 are also presented in the table for comparison. It can be seen that the results of the on-line search are all reasonably close to the target values given by the off-line search.

The variation in the total mean rate of accumulation of the cost during the period of simulated operation is illustrated in Figure 5.19(b). The cost follows a satisfactory downward trend, as does the rate of bearing damage in Figure 5.19(c), while the rate of energy capture in Figure 5.19(d) remains approximately constant. The changes in all of these variables between $t = 1$ hour and $t = 14$ hours are in the same directions as indicated by the results of the off-line searches. (See the results for $\overline{v} = 10$ m/s in Figure 5.18.) The total mean rates of general structural fatigue damage for material slope indices of 3 and 7 are presented in Figures 5.20(a) and 5.20(b) respectively. Both of these measures of fatigue damage exhibit significant downward trends, again as is the case for the off-line searches.

The total mean rate of gearbox teeth damage, with a slope index of 5, is illustrated in Figure 5.20(c). The off-line search indicates that the teeth damage falls by 96% if the optimal set points are employed. The on-line control achieves a reduction of similar magnitude, as the teeth damage falls by 100% by the end of the run.

5.6 Conclusions

A genetic search algorithm can be used to determine the 'ideal' power curves for isolated (non-interacting) turbines. Weighting parameters in the proposed simple cost function for isolated turbines can be evaluated which, when used in a controller, give power curves whose performance differs negligibly from the 'ideal'. The searches for the weighting parameters are computationally much cheaper than the searches for the 'ideal' power curves. When using a weighting parameter that causes moderate de-rating, the cost function has only one minimum as a function of power set point at each wind-speed across the operating range. Therefore an on-line controller would not encounter problems associated with convergence to local minima.

When the simple cost function is applied to a group of three interacting turbines, the power curves that would result if the cost function was to be used in an on-line controller can be found via an off-line search. For a site with a 7.6 m/s annual mean wind-speed, the power curves give estimated annual totals for the four measures of fatigue damage which are between 13% and 36% lower than those given by the standard practice of $P_{sp}^{(1)} = P_{sp}^{(2)} = P_{sp}^{(3)} = P_{rated}$. The annual total for the energy capture is 2.1% lower. This represents a beneficial trade-off, particularly because the wind only blows along rows of turbines infrequently in well-designed windfarms, and at these times the fatigue damage of structural components is disproportionately high due to the wake effects.

The annual gearbox teeth damage of the interacting turbines, with a slope index of 5, is substantially reduced by operation with the power curves which result from the use of the cost function. As the slope index is the exponent to which the low-speed shaft torque is raised in the model of teeth damage, it can be concluded that the damage of bath-nitrided teeth, which have a slope index of 16 and are used in the design of turbine that is modelled, would also be substantially reduced.

It has been shown that the simulation can be used for testing on-line optimising control algorithms. Results for the simulated operation of a small windfarm under extremum control have illustrated the stability and convergence of a simple control algorithm under one set of wind conditions.

Chapter 6

Conclusions and further work

6.1 Conclusions

A windfarm supervisory controller has been designed whose objectives are the maximisation of financial income from the generated electricity and the minimisation of the turbines' fatigue damage. The strategic trade-off between financial income and fatigue damage has been formalised, both a lifetime performance function and a cost function suitable for use in an on-line controller have been proposed, and the effects of the two functions on windfarm operation have been quantified.

A windfarm simulation that has been developed for the purposes of testing supervisory controllers has been described, and it has been shown that the simulation is a suitable test-bed for this application. In general, the results show that the use of supervisory control can increase windfarm performance, by significantly reducing fatigue damage for small to negligible decreases in financial income.

It has been demonstrated that it is advantageous for a supervisory controller to exploit the variable power set point that pitch-controlled wind turbines possess, and to operate the turbines as set point followers instead of simple power regulators. This capability of pitch-controlled turbines is an important, although to date unnoticed, advantage of this class of turbine over stall-regulated designs.

It has been shown that significant performance improvements can be achieved if the operation of the turbines is altered to take account of variable electricity tariffs. Therefore when trading off financial income with fatigue damage when operating (and, ideally, designing) wind turbines, financial income should not be assumed to be proportional to energy capture, as is generally the case currently.

A windfarm control strategy that splits the turbines into 'interacting' and 'non-interacting' categories has been found to give good results. 'Interacting' turbines are those which, at any time, are either standing within the wake of another turbine or are creating a wake within which another turbine is standing.

Ideal operation of individual turbines

'Ideal' power set point curves have been evaluated from the lifetime performance function. The 'ideal' curves indicate that, if multiple tariff seasons per annum are included in the analysis, significantly higher performance can be achieved than for the case of one tariff season per annum.

Cumulative gearbox bearing fatigue damage and energy capture have been found to be very similar functions of both average wind-speed and power set point. Partly as a result of this, the 'ideal' power curves show that for the case of 1 tariff season per annum, attaining significant reductions in the bearing damage necessitates significant, and probably unacceptable, reductions in financial income.

Further results indicate that the components for which supervisory control of individual turbines offers significant potential for reductions in fatigue damage, for small to negligible reductions in financial income, include the gearbox teeth and those structural components whose fatigue loads are proportional to the drive-train torque. In general, it is likely that the components which can benefit most from supervisory control include: components with an endurance limit; structural components; components with high material slope indices.

Suitable cost functions for use in on-line supervisory controllers can be developed independently from the 'ideal' power curves. The 'ideal' results, however, give the highest financial income that can be achieved for one or more constraints on one or more measures of fatigue damage.

Operation of non-interacting turbines in a windfarm

Simple cost functions can be developed for non-interacting turbines in windfarms which give performance that is close to the ideal for specific constraints on annual gearbox bearing fatigue damage. Further results indicate that it may be possible to use the same value of weighting parameter in multiple tariff seasons and still attain performance which is close to the ideal.

Operation of interacting turbines in a windfarm

Cost functions developed for individual, non-interacting turbines have been found to give significant performance improvements when applied to groups of interacting turbines. A cost function for an isolated turbine has been applied to a group of three interacting turbines, with a weighting parameter which causes negligible de-rating for the isolated turbine in typical near-neutral meteorological conditions, and off-line searches have been carried out to find the optimal set point curves. These optimal set point curves are very different to that for the single turbine, with significant de-rating of the two upstream turbines, and substantial reductions in all measures of total annual fatigue damage are achieved for a small fall in total annual financial income. The largest reductions in fatigue damage are experienced in low average wind-speeds. When the turbulence intensity is also low, a slight increase in energy capture is found in addition to significant reductions in all measures of fatigue damage.

The performance of the supervisory controller has been investigated for a line of three turbines, and the minimum inter-turbine spacings that have been investigated are three rotor diameters. It is likely that the demonstrated improvements in performance would be significantly greater both for a longer line of turbines and for smaller inter-turbine spacings, because of the increase in the magnitude of wake effects with the number of interacting wakes and with smaller inter-turbine spacings.

On-line supervisory control

For use in on-line supervisory control and in controller design, all time-averages of fatigue damage and energy capture must continue for long enough such that the estimation errors caused by fluctuations in wind-speed due to atmospheric turbulence are small. This means, in practice, carrying out estimation for a minimum of approximately 10 minutes. In addition, an on-line controller should not attempt to continue optimisation when the average wind-speed is close to the upper cut-out wind-speed, because the estimates of fatigue damage and energy capture are too unreliable. The power set points can be left at those values that are found when the wind-speed is just below the upper cut-out.

It is beneficial to incorporate a constraint on power set point at wind-speeds close to the upper cut-out when implementing an on-line supervisory controller. This prevents high wind-speed start/stops from taking place at times when only low average power levels can be generated while on-line because of set point de-rating.

The effect upon the annual financial income of a simple high wind-speed constraint has been found to be negligible.

It has been shown that a very simple extremum control algorithm operating on three turbines approaches the potential performance indicated by the off-line searches within a small number of steps in the set points. For practical application, a simple method has been proposed that would tune a cost function's weighting parameter directly on site. And it has been found that the dividing line between the classification of turbines into the 'interacting' and 'non-interacting' categories lies, in terms of inter-turbine spacings, between 3 rotor diameters and 7 rotor diameters.

Estimation of fatigue damage

It can be concluded that the standard deviation of power does not provide a suitable generic indicator of fatigue loading. Indeed, the variation of the standard deviation of power with average wind-speed, turbulence intensity and power set point is significantly different to that of all the other measures of fatigue damage. In addition, a large degree of variation has been found between the various damage functions, which indicates that no single function can adequately represent the fatigue damage of the principal components in a turbine.

Results suggest that relatively large reductions in types of damage other than that for which a term is included in the cost function can be obtained as a side effect of the supervisory control, and that these side effects can be greater than the effect on the primary variable itself.

One intuitive basis for the development of supervisory controllers is that wind turbines experience high levels of damage and low contributions to the total energy capture in high wind-speeds, and so large reductions in damage can be achieved for small reductions in energy capture via heavy de-rating in high wind-speeds. However, this assumption is contradicted by the results that have been obtained for the cumulative damage of the most expensive component in the turbine that is modelled, the gearbox. The gearbox damage functions do not even contain maxima in high wind-speeds, with the peak contributions to the annual bearing and teeth damage both occurring closer to the rated wind-speed than to the upper cut-out, at 12 m/s and 16 m/s respectively. The cumulative damage functions for a generalised structural 'component', by contrast, fully support the intuitive basis. Therefore from the perspective of supervisory controller development, the fatigue damage mechanisms of components in a wind turbine can be split into at least two groups: components in the gearbox and structural components.

Windfarm simulation

The results have demonstrated that a simple windfarm simulation is able to characterise the features of a windfarm that are necessary for the purposes of testing supervisory controllers. The behaviour of all the principal components within the simulation has either been validated against site data or has been shown to operate in a manner that is representative of windfarm operation.

Simple wake and wake interaction models can be developed that are adequate for this application. The lack of coherence of high frequency fluctuations in wind-speed in the longitudinal direction can be exploited to allow the wake effects to be split up into two components: a low frequency part and a high frequency part.

6.2 Further work

Further work is required in the following areas:

Cost functions and fatigue damage models

- Transfer functions from the low-speed shaft torque to the blade root bending moment and tower overturning moment should be incorporated; models of the fatigue damage of the blades and tower and a model of the generator damage should be added.

- A methodology should be developed for evaluating the weighting parameters in cost functions which contain multiple terms for fatigue damage. The relative magnitudes of the weights should probably be calculated on the basis of the respective components' replacement costs.

- A comparison should be made of the performance that can be attained from the use of cost functions which contain terms for the damage of any combination of the following components: gearbox teeth, gearbox bearings, blades, tower, generator. A cost function which includes a single term for fatigue damage is likely to be most promising in practice, as the number of weighting parameters is then minimised. A successful cost function should, when used in a controller, reduce the fatigue damage of as many components as possible by the maximum possible. In addition, the estimate of fatigue damage should converge swiftly at all wind-speeds. Results suggest that a damage model which has the greatest potential for achieving both of these goals is likely to be for a structural component whose fatigue loads are approximately proportional to the drive-train torque and whose slope index is relatively low.

- The magnitude of the improvements in performance that can be achieved if more than two tariff seasons per annum are used in the analysis should be investigated, particularly if diurnal variations in tariff are incorporated. It is likely that substantial further improvements in performance would be attained, because in this study the two averaged tariffs vary by a factor of 1.7 while in practice the tariffs vary (across four bands) by a factor of 4.6.

- The effects on turbine performance and design of allowing operation, under certain smooth wind conditions, with power set points that are greater than the rated power should be examined. Either:

  - the use of current designs should be assumed and the additional energy capture that can be achieved should be investigated, or

  - a broader study should be carried out into the optimal turbine design and operational strategy that results from the ability to operate at above rated power.

  In the former case, areas that should be investigated include the possible effects on generator heating as well as any effects on the fatigue damage of other components. The restriction of the use of the above-rated power set point to periods when the standard deviation of wind-speed is low might allow the strengthening of many components to be avoided, as there would be no effect on the worst case operational conditions used in some design codes.

- Experimental validation of the fatigue damage models should be carried out.

On-line supervisory control

- The on-line control performance that can be achieved by the use of self-tuning or neural network extremum controllers should be investigated (Wellstead and Scotson, 1990) (Nascimento et al., 1992). More sophisticated schemes should be capable of delivering significantly better performance than the very basic hill-climbing controller that has been used in this study.

Windfarm simulation

- An investigation should be carried out to determine whether the spectra of the wind-speed at the downstream turbines are distorted significantly by the division of the wake model into low and high frequency signals. Only the means and standard deviations of the wind-speed at the downstream turbines are validated in this thesis. If significant errors do exist, one solution could be to replace the first order filters that are currently used for the low- and high-pass filtering of the two signals with higher order filters.

- The effect upon the estimates of fatigue damage and energy capture of removing the transport delay from the wake model should be investigated. It has been shown that all estimates of fatigue damage and energy capture must continue for a minimum period of the order of 10 minutes, and the time-delay of the wind-flow from turbine to turbine is negligible by comparison. Hence it is possible that if the data queue and reconstruction filters in the wake model are removed, the effect upon the convergence of the estimates would be negligible.

Appendix A

Nomenclature and symbols

General nomenclature

$a^{(i)}$	value of variable a for turbine i in windfarm
$\dot{a}$	rate of change of $a(t)$
$\overline{a}$	mean value of fluctuations in $a(t)$ that are caused by the turbulent wind, for which $f > \sim \frac{1}{10 \text{ mins}}$
$\hat{a}$	estimated value of a
a_{gear}	value of a for gearbox
a_{hss}	value of a for high-speed shaft
a_{life}	value of a for lifetime of turbine
a_{lss}	value of a for low-speed shaft
a_{roy}	value of a for Rest Of Year season (March to November)
a_{win}	value of a for Winter season (December to February)
a_y	value of a for year

Principal symbols

a_{icvd}	initial centre-line velocity deficit behind rotor	–
A_{gear}	fatigue damage safety factor for gearbox ($0 \leq A \leq 1$)	–
A_{blades}	fatigue damage safety factor for blades ($0 \leq A \leq 1$)	–
A_{tower}	fatigue damage safety factor for tower ($0 \leq A \leq 1$)	–
b	slope index of S–N curve	–
C_p	power coefficient	–
C_t	thrust coefficient	–
d	fatigue damage	various
d_b	fatigue damage of gearbox bearings	$(\mathrm{kNm})^{3.33}\,\mathrm{s}$
d_{max}	maximum possible annual fatigue damage, for which $P_{sp} = P_{rated}$ throughout	various
d_s	fatigue damage of generalised structural 'component'	various
d_t	fatigue damage of gearbox teeth	various
D	rotor diameter	m
e	energy capture	kWs
f	frequency	Hz
i	wind-speed interval number (Chapter 2 only)	–
i_{hi}	wind-speed interval number for upper cut-out wind-speed	–
i_p	period number	–
i_s	season number	–
i_y	year number	–
I	turbulence intensity	–
I_{inc}	turbulence intensity of wind incident upon a turbine	–
I_{target}	target turbulence intensity	–
J	cost of operation (to be minimised)	–
k_w	Weibull shape factor	–

K_m	tariff for electricity generated	p/kWh
m	money (financial income from generated electricity)	£
m_{max}	maximum possible annual income, for which $P_{sp} = P_{rated}$ throughout	£
n_d	number of days in season	–
n_p	number of periods in day	–
n_s	number of seasons in year	–
n_y	number of years in design lifetime of turbine	–
$p(\overline{v})$	probability of wind-speed $\overline{v}$	–
$p_{\{i_s,i_p\}}(\overline{v})$	probability of wind-speed $\overline{v}$ for period i_p in season i_s	–
P	power	kW
P_{aero}	aerodynamic power extracted by rotor	kW
P_{rated}	rated power	kW
P_{sp}	power set point	kW
q	torque	kNm
q_{lss_∞}	endurance limit torque for gearbox teeth on low-speed shaft	kNm
s	Laplace transform operator, $s \equiv d/dt$	s^{-1}
t	time	s
T_c	time-step of supervisory controller	s
T_{pi}	time-step of turbine's pitch controller	s
T_p	length of period	h
v	wind-speed	m/s
$\overline{v}_i$	mean of the turbulent wind-speed for interval i	m/s
v_{inc}	wind-speed incident upon a turbine	m/s
v_{lf}	low frequency component of wind-speed	m/s
v_{lo}	lower cut-in/cut-out wind-speed	m/s
v_{hf}	high frequency component of wind-speed	m/s
v_{hi}	upper cut-out/cut-in wind-speed	m/s

v_{rated}	rated wind-speed	m/s
$\overline{v}_{year}$	annual mean wind-speed	m/s
λ	tip-speed ratio	–
λ_w	cost function's weighting parameter	–
σ_p	standard deviation of power	kW
σ_v	standard deviation of turbulent wind-speed, for which $f > \sim \frac{1}{10 \text{ mins}}$	m/s
σ_{vlf}	standard deviation of low frequency component of wind-speed in model of the turbulent wind	m/s
σ_{vhf}	standard deviation of high frequency component of wind-speed in model of the turbulent wind	m/s
$\sigma_{v\infty}$	standard deviation of free stream wind-speed	m/s
ϕ	pitch angle of blades Convention: $\phi = 0°$ maximum aerodynamic torque generated $\phi = 90°$ minimum aerodynamic torque generated	°
ω	rotational speed	rad/s

Appendix B

Wind turbine model

Many models have been produced and reported on for the simulation of single wind turbines for control purposes. See (de la Salle et al., 1990) and (Mattsson, 1984) for reviews. In the case of simulation for testing windfarm controllers, however, the requirements of the models of the individual turbines differ somewhat. The models can be relatively simple, as it is only necessary to model those characteristics of the machine operation to which a windfarm supervisory controller will be sensitive.

For the testing of windfarm supervisory controllers, the model of the individual wind turbines must contain elements to represent at least the following:

- rotor aerodynamics
- drive train
- pitch controller

A list of the principal parameters of the turbine that is modelled in this thesis is given in Table 3.1 (page 59). Additional parameters are presented both in this appendix and in Appendix C (pages 204*ff*).

A block diagram of the model of the turbine is presented in Figure B.1. The model contains one input, one disturbance signal and four outputs. The input is the power set point, the disturbance signal is the wind-speed incident upon the turbine and the outputs are the generated power, the low-speed shaft torque, the thrust coefficient of the rotor and the initial centre-line velocity deficit.

Although the turbine that is modelled incorporates an induction generator, and hence the rotational speed can vary slightly during operation, it is assumed that the rotational speeds of both the low and high speed shafts are constant and equal to their nominal values.

B.1 Rotor aerodynamics

As Figure B.1 illustrates, the first element in the model of the rotor aerodynamics is a spatial filter (Madsen and Frandsen, 1984a) (Leithead et al., 1989). This filter represents the swept disc averaging effect of the rotation of the rotor in the air-stream and can be modelled with the following transfer function:

$$G_{sf}(s) = \frac{1 + s\frac{\beta}{\sqrt{2}}}{\left(1 + s\frac{\beta\sqrt{a}}{\sqrt{2}}\right)\left(1 + s\frac{\beta}{\sqrt{a}}\right)}$$

where the nominal parameters given in the literature are used:

$$a = 0{\cdot}5 \qquad \beta = \frac{1{\cdot}3\ R}{\overline{v}}$$

R is the radius of the turbine's rotor.

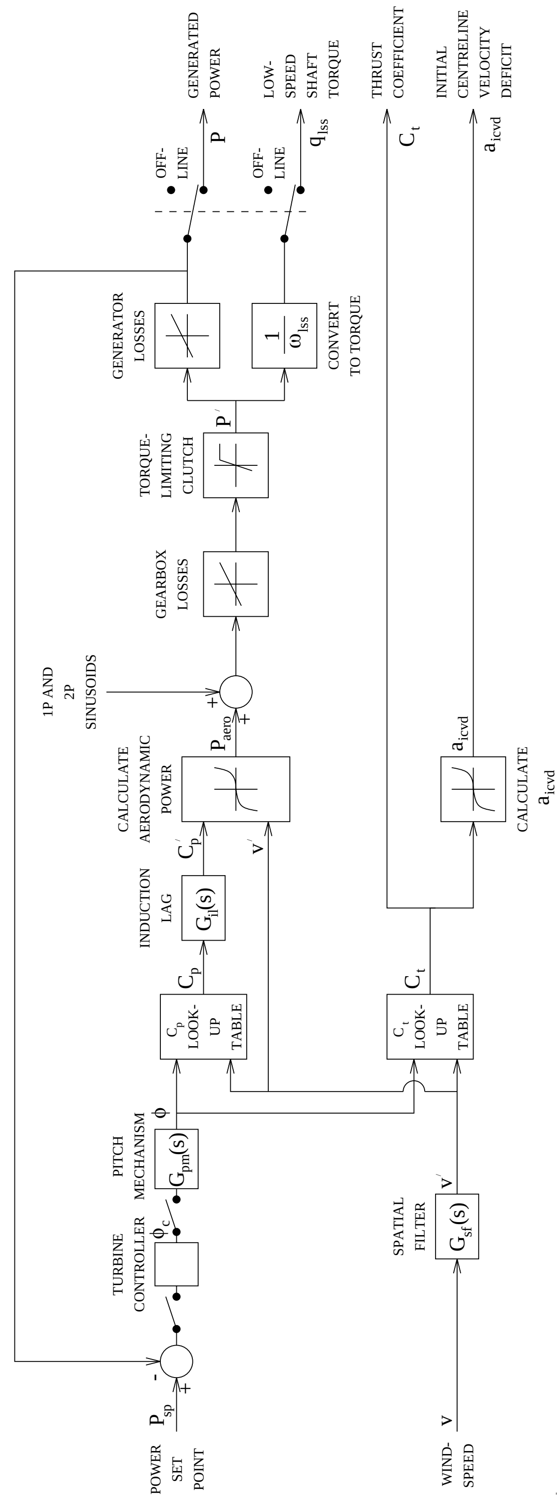

Figure B.1: Block diagram of wind turbine model

Aerodynamic power P_{aero} and thrust T_{aero} can be represented as follows:

$$P_{aero} = \frac{1}{2} C_p \rho A_r v^3 \qquad T_{aero} = \frac{1}{2} C_t \rho A_r v^2$$

where C_p is the power coefficient, C_t is the thrust coefficient, ρ is the density of air and A_r is the cross-sectional area of the rotor. C_p and C_t are highly non-linear functions of both wind-speed and blade pitch angle (ϕ), and are modelled by the use of interpolation tables:

$$C_p = f(v, \phi) \qquad C_t = f(v, \phi)$$

The data used in the tables was supplied by the turbine manufacturer and had been generated using a proprietary software package (BLADES, 1988). See Appendix C for numerical values.

The aerodynamic induction lag is modelled as the dominant forward path dynamics in the turbine model (Stig Øye, 1986) (Leithead et al., 1989). This induction lag is caused by the wake immediately downstream of the rotor having to re-align itself following any change in the incident wind-speed or in the pitch angle of the blades. Dynamically, the result can be considered as an instantaneous over- or under-shoot in aerodynamic power followed by a first order lag:

$$G_{il}(s) = \frac{1 + a_{il} s}{1 + \tau_{il} s}$$

The values used for the constants are $a_{il} = 1.37$ and $\tau_{il} = 9$ seconds (Leithead et al., 1989).

The pitch mechanism has the dominant dynamics in the feedback path and can be modelled with a second order transfer function (Spruce, 1988a):

$$G_{pm}(s) = \frac{1}{1 + a_{pm}s + b_{pm}s^2}$$

where $a_{pm} = 0.12$ and $b_{pm} = 5.9 \times 10^{-3}$. The function $G_{pm}(s)$ is a fit to results obtained from a non-linear simulation of the pitch mechanism. In addition to $G_{pm}(s)$, the physical end stop is modelled which prevents the blades from moving to a pitch angle* of less than $-2°$.

The power harmonics that the 2-bladed turbine experiences at frequencies of once per rotor revolution (1P) and twice per revolution (2P) are assumed to be sinusoids of constant amplitude and phase. These sinusoids are added to the aerodynamic power P_{aero} as indicated by Figure B.1, and their amplitudes are evaluated as follows. At any wind-speed, the amplitudes of the 1P and 2P sinusoids must be zero when $P_{sp}(v) = 0$ and have maxima when $P_{sp} \geq P_{max}(v)$, where $P_{max}(v)$ is given by:

$$v < v_{rated}: \quad P_{max}(v) = \overline{P}(v) \quad \text{when} \quad \phi = 0°$$

$$v > v_{rated}: \quad P_{max}(v) = P_{rated} = 300 \text{ kW}$$

In addition, aerodynamic power is proportional to the wind-speed cubed. Consequently the amplitudes of the 1P and 2P sinusoids (A_{1P} and A_{2P}) are approximated by the following functions:

$$A_{1P} = K_{1P}\ r\ v^3 \qquad\qquad A_{2P} = K_{2P}\ r\ v^3$$

where K_{1P} and K_{2P} are constants, and r is a fraction representing the effect of

*The convention for pitch angle ϕ is as follows:

$\phi = 0°$	blades at 'fine'	maximum aerodynamic torque generated
$\phi = 90°$	blades at 'feather'	minimum aerodynamic torque generated

variations in power set point upon the amplitudes of the 1P and 2P sinusoids:

$$K_{1P} = 2.2 \times 10^{-3}\ \mathrm{kW/(m/s)^3} \qquad K_{2P} = 4.5 \times 10^{-3}\ \mathrm{kW/(m/s)^3}$$

$$\begin{array}{lll} \text{if } v < v_{rated}: & & \\ & \text{if } P_{sp} < P_{max}(v) & r = \dfrac{P_{sp}}{P_{max}(v)} \\ & \text{else} & r = 1 \\ \text{if } v \geq v_{rated}: & & r = \dfrac{P_{sp}}{P_{rated}} \end{array}$$

K_{1P} and K_{2P} were calculated using six power spectra of measured data which were supplied by the turbine manufacturer. The spectra are for wind-speeds of between 7.0 m/s and 20.7 m/s, and all are for $P_{sp} = P_{rated}$. The amplitudes of the 1P and 2P components are firstly estimated from the spectra, and then the above functions for A_{1p} and A_{2P} are fitted to the measured data points. Satisfactory agreement between the fitted functions and all the data points is attained using this approach. A_{1P} and A_{2P} in the above equations are both divided by the fractional drive-train losses ($K_{loss_{gear}} \times K_{loss_{gen}}$) before being used in the model, as the 1P and 2P sinusoids are added to the model's low-speed shaft power. $K_{loss_{gear}}$ is the gearbox losses and $K_{loss_{gen}}$ is the generator losses. Both are defined in Section B.2 below.

Two of the outputs from the model, power P and low-speed shaft torque q_{lss}, are required for the assessment of energy capture, fatigue damage and other variables. The other two outputs, the thrust coefficient C_t and initial centre-line velocity deficit a_{icvd}, are required for the purposes of calculating the wind conditions in the turbine's wake. a_{icvd} is evaluated by using a simple one-dimensional fluid flow model of the rotor, the Rankine–Froude actuator disk. In this model the rotor is replaced by an 'actuator disc' through which the static pressure decreases discontinuously (Eggleston and Stoddard, 1987)[pp. 20–25] (Katić et al., 1986), and a_{icvd} is related to C_t

as follows:

$$a_{icvd} = \frac{1 - \sqrt{1 - C_t}}{2}$$

B.2 Drive train

The drive-train power losses of the gearbox and the generator, indicated in Figure B.1, are modelled using simple linear relationships of the following form:

$$P' = K_{loss}P + C_{loss}$$

Numerical values for the loss constants K_{loss} and C_{loss} for the gearbox and generator are given in Appendix C.

The model of the drive-train includes a term for a slipping clutch in the high-speed shaft which limits the instantaneous torque and, equivalently, the power:

$$P' = \begin{cases} K_{clutch}\, P_{rated} & \text{if } P > K_{clutch}\, P_{rated} \\ P & \text{otherwise} \end{cases}$$

The power limit is 500 kW in the case of the MS–3 turbine, so $K_{clutch} = 1.7$.

B.3 Wind turbine controller

The performance of any supervisory windfarm controller is strongly dependent upon the performance and capabilities of the plant level controllers, because it is the plant level controllers that carry out the dynamic control to the set points which are specified by the supervisory controller. As a result the model of the turbines' controllers is relatively detailed. All data and algorithms described in this section were supplied by the turbine manufacturer.

A discrete-time model is used to represent the pitch controller, which employs proportional–integral (PI) control with gain scheduling. The PI algorithm is modelled as follows:

$$\phi'_t = \phi'_{t-1} + b_0\, e_t + b_1\, e_{t-1}$$

where

$$b_0 = K(\phi)\,.\left(1 + \frac{T_{pi}}{2T}\right)$$

$$b_1 = K(\phi)\,.\left(-1 + \frac{T_{pi}}{2T}\right)$$

ϕ'	= pitch angle demand
e	= power error = $P_{sp} - P$
$K(\phi)$	= proportional gain
T_{pi}	= time-step of pitch controller
T	= integral gain

The proportional gain is scheduled as a function of pitch angle, to compensate for the heavily non-linear nature of the torque-pitch gain:

$$K(\phi) = \frac{K(0^\circ)}{1 + 0{\cdot}33\,\phi}$$

The numerical values for the gains and time-step are as follows:

$$K(0^\circ) = 0.147\ ^\circ/\text{kW} \qquad T = 0.0952\ \text{s} \qquad T_{pi} = 60\ \text{ms}$$

The blades spend a high fraction of their operational time in the region close to 0° and therefore accurate re-production of the low pitch angle end-stop is necessary. Limits are applied to the pitch angle demand (ϕ') that is generated by the turbine's controller, such that $0^\circ \leq \phi' \leq 90^\circ$. As described in Section B.1, the lower physical end-stop for the blades of -2° is also modelled. This allows the pitch angle to go transiently negative when it is moving from a positive angle towards 0°.

The start/stop algorithm that is used in the turbine's controller is itself a supervisory control action, and therefore has important implications for the design of a windfarm supervisory controller. As a result the turbine's start/stop algorithms are also carefully re-produced. An average wind-speed ($\overline{v}'$) is calculated by filtering the incident wind-speed using a first order filter with a time-constant of 33 seconds, and then the turbine cuts in or out as $\overline{v}'$ crosses certain thresholds. The thresholds are different for cutting in and cutting out, thereby providing hysteresis that prevents the turbine from starting and stopping too frequently. The turbine also incorporates a cut-out that is triggered by low rotational speed of the rotor, and this can be modelled by taking a running average of power ($\overline{P}'$) using a 20-second rectangular window and applying a low power threshold (Smith, 1990). The complete cut-in and cut-out logic for the wind-speed and power is as follows:

Turbine cuts in: IF $\overline{v}'$ exceeds 6 m/s OR IF $\overline{v}'$ drops below 23 m/s
Turbine cuts out: IF $\overline{v}'$ drops below 5 m/s OR IF $\overline{v}'$ exceeds 25 m/s
Turbine cuts out: IF $\overline{P}'$ drops below 0 kW

Appendix C

Parameters used in the wind turbine model

The principal characteristics and parameters of the turbine are given in Table 3.1 (page 59) and additional parameters are given throughout Appendix B. All data in this appendix was supplied by the turbine manufacturer.

Gearbox and generator losses

$$K_{loss_{gear}} = 0.98 \qquad C_{loss_{gear}} = -3.9 \text{ kW}$$
$$K_{loss_{gen}} = 0.97 \qquad C_{loss_{gen}} = -8.1 \text{ kW}$$

Power and thrust coefficients

The power coefficients are presented in Table C.1 and the thrust coefficients are presented in Table C.2.

		Tip-speed ratio								
		3	4	5	6	7	8	10	13	16
		Wind-speed (m/s)								
		27.6	20.7	16.6	13.8	11.9	10.4	8.29	6.38	5.18
	-2	0.057	0.194	0.333	0.418	0.444	0.442	0.403	0.279	-0.221
	0	0.076	0.215	0.339	0.402	0.435	0.450	0.433	0.262	-0.422
	2	0.107	0.233	0.330	0.383	0.411	0.426	0.431	0.090	-0.890
ϕ	5	0.129	0.236	0.294	0.325	0.338	0.335	0.143	-0.652	-2.120
(°)	10	0.143	0.198	0.204	0.160	0.041	-0.149	-0.777	-2.475	-5.318
	15	0.127	0.123	0.038	-0.130	-0.391	-0.758	-1.867	-4.608	-8.757
	20	0.086	0.009	-0.167	-0.452	-0.821	-1.275	-2.556	-5.680	-10.650
	25	0.026	-0.119	-0.335	-0.622	-1.012	-1.525	-3.004	-6.694	-12.611

Table C.1: Power coefficient as a function of pitch angle and wind-speed

		Wind-speed (m/s)								
		27.6	20.7	16.6	13.8	11.9	10.4	8.29	6.38	5.18
	-2	0.157	0.291	0.470	0.636	0.743	0.829	0.984	1.198	1.420
	0	0.158	0.297	0.460	0.577	0.664	0.736	0.841	0.978	1.121
	2	0.171	0.304	0.432	0.521	0.581	0.630	0.694	0.778	0.849
ϕ	5	0.177	0.293	0.367	0.413	0.439	0.443	0.452	0.439	0.383
(°)	10	0.174	0.231	0.240	0.224	0.196	0.156	0.041	-0.220	-0.583
	15	0.145	0.145	0.109	0.051	-0.027	-0.126	-0.384	-0.863	-1.322
	20	0.098	0.061	-0.013	-0.116	-0.215	-0.298	-0.465	-0.740	-1.067
	25	0.050	-0.021	-0.092	-0.148	-0.202	-0.255	-0.366	-0.619	-1.021

Table C.2: Thrust coefficient as a function of pitch angle and wind-speed

Appendix D

Determination of the convergence of estimated means

An estimate is considered to have converged once the value of a convergence parameter falls within a confidence interval. Two heuristic rules are also applied.

In this appendix, details are given of the use of confidence intervals and then of the complete procedure that is used to determine when an estimate of a mean of a variable has converged.

D.1 Confidence intervals

For a random variable x, the estimate of the mean ($\overline{x}$) is also a random variable, and the estimated variance of $\overline{x}$, $\hat{\sigma}_{\overline{x}}$, can be written as (Jacobs, 1993)[pp. 266–269]:

$$\hat{\sigma}_{\overline{x}} = \frac{\sigma}{\sqrt{n}} \tag{D.1}$$

where σ = standard deviation of distribution; n = number of samples.

If it is assumed that the samples are normally distributed, the confidence interval (Kreyzig, 1983)[pp. 947–955] can be written as:

$$\mathrm{CONF}\left\{\overline{x} - c\hat{\sigma}_{\overline{x}} \leq \mu \leq \overline{x} + c\hat{\sigma}_{\overline{x}}\right\}$$

where

$$\begin{aligned} \mu &= \text{mean} \\ c &= f(\gamma),\ \text{where } \gamma = \text{required confidence level} \end{aligned}$$

Let the tolerance of the confidence interval be applicable to $\overline{x}$ lying within $\pm\delta.\mu$ of μ, where δ is a 'small' number:

$$c\hat{\sigma}_{\overline{x}} = \delta.\mu$$

For the required degree of accuracy of the estimates, $c\hat{\sigma}_{\overline{x}} \ll \overline{x}$ after convergence has been achieved. Assuming that the estimator is biased and that a large number of samples have been taken, $\mu \approx \overline{x}$. Therefore the confidence interval is for $\overline{x}$ lying within approximately $\pm\delta.\overline{x}$ of μ:

$$c\hat{\sigma}_{\overline{x}} \approx \delta.\overline{x}$$

Hence for satisfaction of the required confidence level,

$$c\hat{\sigma}_{\overline{x}} \leq \delta.\overline{x}$$

and the convergence criterion can be written as:

$$a_{conv} \leq \frac{\delta}{c} \tag{D.2}$$

where

$$a_{conv} = \frac{\hat{\sigma}_{\overline{x}}}{\overline{x}} \tag{D.3}$$

$$= \text{convergence parameter}$$

Continuous time-averaging is used for estimating the means and standard deviations of the wind-speed and rates of energy capture, gearbox bearing damage and gearbox teeth damage:

$$\overline{x} = \frac{1}{T}\int_0^T x(t)\,dt$$

$$\hat{\sigma}_{\overline{x}} = \frac{1}{T}\sqrt{\int_0^T x^2(t)\,dt - \frac{1}{T}\left(\int_0^T x(t)\,dt\right)^2}$$

where the equation for $\hat{\sigma}_{\overline{x}}$ is derived from Equation D.1. The simulation package's integration algorithm is used for the continuous-time estimation.

For the estimation of generalised structural damage, data is logged continuously during a simulation run. However, the cycle-counting and calculation of the means and standard deviations of the rates of structural damage are carried out at discrete intervals (of 50s):

$$\overline{x} = \frac{1}{n}\sum_{i=1}^{n} x_i$$

$$\hat{\sigma}_{\overline{x}} = \frac{1}{n}\sqrt{\sum_{i=1}^{n} x_i^2 - \frac{1}{n}\left(\sum_{i=1}^{n} x_i\right)^2}$$

D.2 Procedure

The procedure for determining if an estimate of the mean of a signal has converged is as follows (Kreyzig, 1983)[p. 950]:

Calculation of convergence threshold

1. Choose a confidence level γ. ($\gamma = 95\%$ is used for the rate of gearbox bearing damage and energy capture.)
2. Choose δ, where $\pm\delta.\overline{x}$ is the tolerance within which μ will lie on convergence. ($\delta = 5\%$ is used for the rate of gearbox bearing damage and energy capture.)
3. Evaluate $F(c) = \frac{1}{2}(1+\gamma)$.
4. Determine c from $F(c)$ using the t-distribution of Student.
5. Calculate convergence threshold b_{conv}, where $b_{conv} = \delta/c$.

Test for convergence

1. Monitor a_{conv}, where $a_{conv} = \hat{\sigma}_{\overline{x}}/\overline{x}$.
2. Convergence is achieved when the confidence interval is satisfied, that is, when $a_{conv} \leq b_{conv}$, subject to the following heuristic rules:
   (a) Every simulation must run for a minimum of 10 minutes.
   (b) If an estimated mean is less than 10^{-3} times the maximum possible value of the estimate across the operating range, and if more than 10 minutes have elapsed, the run is terminated.

The maximum possible value of the estimated mean rate of fatigue damage or energy capture that is required for the application of the second heuristic rule is evaluated as follows. The basic shapes of the surfaces of $\hat{\dot{d}}$ (all measures) and $\hat{\dot{e}}$ as functions of $\overline{v}$ and P_{sp} are found by conducting simulations for all wind-speeds across the operating range and for all set points in the range $P_{rated}/2$ to P_{rated}. The maximum values of all the estimated means are found to occur when $P_{sp} = P_{rated}$, for all wind-speeds. Therefore it is assumed that the maximum numerical value of each estimate across the whole operating space is the maximum of the values for which $P_{sp} = P_{rated}$.

The fraction of 10^{-3} by which the maximum possible rate of damage or energy capture is multiplied when evaluating the threshold in the second heuristic rule is chosen as follows. The second heuristic rule is designed to terminate a simulation run only if $\hat{\dot{d}}$ (all measures) and $\hat{\dot{e}}$ satisfy two conditions. Firstly, $\hat{\dot{d}}$ and $\hat{\dot{e}}$ should both be negligible. The fraction of 10^{-3} ensures that this is the case. Secondly, $\hat{\dot{d}}$ and $\hat{\dot{e}}$ should, when combined with the annual wind-speed probability distribution $[\ p(\overline{v}) \times \hat{\dot{d}}(\overline{v}, P_{sp}),\ p(\overline{v}) \times \hat{\dot{e}}(\overline{v}, P_{sp})\ ]$, make negligible contributions to the cumulative totals. The probability of occurrence of high wind-speeds is much lower than that of low wind-speeds, so relatively low rates of damage and energy capture in low wind-speeds can become significant within the cumulative total. The ratio of values of $p(\overline{v})$ in the range of wind-speeds of interest, $p(8\ \mathrm{m/s})$:$p(23\ \mathrm{m/s})$, is 90:1. Therefore the fraction of 10^{-3} ensures that the second heuristic rule only terminates simulation runs for which negligible contributions are made to the cumulative damage and energy capture.

Appendix E

Overview of the genetic search algorithm

E.1 The basic algorithm

The algorithm that is implemented is Goldberg's 'Simple Genetic Algorithm', Pascal code for which is presented in (Goldberg, 1989). The coding of the power set point curves into strings for use in the algorithm is described in Section 5.2.1 (pages 144*ff*).

An initial group, or population, of strings is chosen at random, and this population is then allowed to evolve from generation to generation. In natural terminology, each string is equivalent to a chromosome, and each bit within a string is equivalent to a gene. The heart of the genetic algorithm is the evolution from one population to the next, which consists of the following steps:

- Reproduction, by roulette wheel selection
- Random mating and crossover
- Random mutation

In the reproduction stage, individual strings are copied according to their fitness. In this application the fitness is the annual financial income produced by the wind turbine. Offspring strings are allocated using a roulette wheel with slots sized according to fitness, so the greater the fitness of an individual, the greater its probability of being chosen for reproduction.

In the mating and crossover stage, two steps take place. Firstly, members of the newly reproduced strings in the mating pool are mated at random. Secondly, each pair of strings undergoes crossing over as follows: an integer position i along the string is selected uniformly at random between 1 and the string length l minus one $[1, l-1]$; then two new strings are created by swapping all bits in positions $i+1$ to l. A crossover probability of 0.6, recommended by Goldberg, is used in this application.

In the mutation stage, bits are chosen by random throughout the population and their values are changed. The probability of mutation of any individual bit should be low and a mutation probability of 0.01 is used in this application. The performance of the search is lower for mutation probabilities that are substantially lower or higher than 0.01.

The bulk of a genetic algorithm's processing power comes from the reproduction and crossover stages. Substrings within each string contain, in effect, various notions of what is important or relevant to the task. In the reproduction and crossover stages, high-quality notions are reproduced according to their performance and are then crossed with many other high-performance notions from other strings. The algorithm's mutation operator has, by contrast, a secondary effect on performance. Mutation is necessary purely to provide insurance against the premature loss of potentially useful genetic material (0's or 1's at particularly locations).

When used for the purposes of searching for ideal power curves, the performance of the basic algorithm is poor. The convergence is so slow that the results appear to be no better than a random search. Various extensions to the basic algorithm suggested by Goldberg have therefore been investigated.

E.2 Extensions to the basic algorithm

E.2.1 'Elite' function

In natural terminology, the 'elite' function involves comparing the 'king' of the old generation to the 'king' of the new generation. If the king of the new generation is weaker than the king of the old generation, then the 'runt' of the new generation is found, killed, and replaced with the 'king' of the old generation.

This method ensures that the fitness of the fittest individual never drops from generation to generation. It has been shown to significantly improve performance on unimodal surfaces but to degrade performance on multimodal functions (Goldberg, 1989). Therefore elitism improves local search at the expense of global perspective.

Following implementation of the elite function, satisfactory convergence of the fitness of the strongest individual of each population is achieved, although the effect on average fitness is negligible.

It can therefore be concluded that, when using a genetic algorithm to search for ideal power curves, the use of an 'elite' function is necessary.

E.2.2 Fitness scaling

Fitness scaling can help to prevent two potential problems which can occur when using a genetic algorithm for search. The first problem can be experienced at the start of a search, where it is common to have a few very strong individuals in a population of mediocre colleagues. If left to the normal selection rule, the extraordinary individuals would take over a significant proportion of the population in a single generation, thereby giving premature convergence. The second problem can be experienced late on in a search, where the opposite situation takes place. Although there may be significant diversity within the population, the population average fitness may be close to the population best fitness. If this situation remains unchanged, average members and best members get nearly the same number of copies in future generations, and the survival of the fittest necessary for improvement becomes a random walk through the mediocre.

The raw fitness f is scaled using a power law:

$$f' = f^a$$

where a is a constant. It is always necessary for the average scaled fitness $\overline{f'}$ to be equal to the average raw fitness $\overline{f}$, because subsequent use of the selection procedure will ensure that each average population member contributes one expected offspring to the next generation. $\overline{f}$ and $\overline{f'}$ are given by:

$$\overline{f} = \frac{1}{n_{pop}} \sum_{i=1}^{n_{pop}} f_i \qquad\qquad \overline{f'} = \frac{1}{n_{pop}} \sum_{i=1}^{n_{pop}} (f_i)^a$$

where n_{pop} is the number of individuals in the population.

As $a = 1$ is not of interest:

$$\overline{f'} \neq \overline{f}$$

Therefore a constant of multiplication b is inserted to compensate:

$$f' = bf^a$$

where

$$b = \frac{\sum_{i=1}^{n_{pop}} f_i}{\sum_{i=1}^{n_{pop}} (f_i)^a}$$

This achieves the required relationship $\overline{f'} = \overline{f}$.

To evaluate a, various values in the range 1.001 to 16 were tested and $a = 8$ was then selected.

When generating the final results, initial searches are carried out without fitness scaling, using the genetic algorithm's inherently noisy nature to give relatively wide coverage across the search space, and further duplicate searches are then carried out with fitness scaling turned on, the aim of the latter being to home in more precisely on the optimum. The purpose of the initial searches is to check that the later ones do not locate local minima, which could be a problem when using the large scaling exponent $a = 8$.

In practice, the effect of the fitness scaling on the fittest individuals is marginal. The maximum fitness changes little when a is varied across the large range of 1.001 to 16, thereby indicating firstly that the effect of the fitness scaling on maximum fitness is probably dominated by the effect of the 'elite' function, and secondly that the performance functions' surfaces are likely to be unimodal, that is, they do not contain local minima. Although the fitness scaling has a negligible effect upon the maximum fitness, the average fitness converges considerably faster and achieves a higher level of fitness.

It can be concluded that, when using a genetic algorithm to search for ideal power curves, the use of fitness scaling is unlikely to provide performance improvements and need not be implemented.

E.2.3 Parallel populations

The theory of parallel populations is based upon the existence of a number, n, of independent populations which evolve separately. The fittest $1/n$ individuals within each population are then, at some chosen point, brought together to form a single population, which is then used as the initial population for further application of the genetic algorithm. This structure can be further cascaded to more than two levels.

The use of parallel populations in genetic search can help to avoid a search converging to a local minimum. In natural terminology, the cross-fertilisation between the populations protects against 'in-breeding'. A very simple implementation of parallel populations is used in this application. Three populations are started in parallel, and, once the searches have converged, the fittest individual of the three populations is then picked out. In all cases, the results of the three searches differ by negligible amounts, thereby providing further evidence that the surfaces of the performance functions for the searches do not contain strong local minima.

It can be concluded that, when using a genetic algorithm to search for ideal power curves, the use of parallel populations is unlikely to provide performance improvements. It will, however, give greater confidence in the results.

E.3 Performance of search

The performance of a search algorithm can be evaluated by considering the number of evaluations of the function that are carried out as a fraction of the total number of points in the search space:

$$f_p = \frac{n_{eval}}{n_{space}}$$

where f_p = performance fraction (low fraction $\equiv$ high performance)

n_{eval} = number of function evaluations

n_{space} = number of points in search space

Consider the case of two tariff seasons per annum, for which $n_{space} = 32^{36} \approx 10^{54}$. The population size is 100 individuals, and if it is assumed that the search can be considered to have converged after 10 000 generations, then $n_{eval} = 10^6$ and so $f_p = 10^{-48}$.

Appendix F

Simple extremum control algorithm

The extremum control algorithm consists of a simple direct search which probes along predetermined directions (Dixon, 1972)[pp. 65-66]. The directions in this application are the power set points of the three turbines that are simulated. Logic is incorporated in the algorithm, to firstly prevent the set points from exceeding their constraints, and secondly stop the same set point from being specified more than once while searching in a particular direction.

Nomenclature:

i	= step number
$k(i)$	= step length
$P_{sp}(i)$	= power set point
P'_{sp}	= best point so far
P_{sp_min}	= lower constraint on power set point
P_{sp_max}	= upper constraint on power set point

For each P_{sp} in turn:

Initial conditions:

$$
\begin{aligned}
i &= 0 \\
k(0) &= -\tfrac{1}{2} \times P_{sp}(0) \\
P_{sp}(0) &= P'_{sp} \\
P'_{sp} &= 300 \text{ kW} \\
P_{sp\_min} &= 0 \text{ kW} \\
P_{sp\_max} &= 300 \text{ kW}
\end{aligned}
$$

$\boxed{\text{At start of controller's step:}}$

1. Increment step counter: $i := i + 1$

2. Choose new P_{sp}:

   (a) Initial choice: $P_{sp}(i) = P'_{sp} + k(i)$

   (b) Limit $P_{sp}(i)$ and $k(i)$ if $P_{sp}(i)$ exceeds a constraint:

`if` $P_{sp}(i) < P_{sp\_min}$ `then`

$$
\begin{aligned}
P_{sp}(i) &= P_{sp\_min} \\
k(i) &= P_{sp\_min} - P_{sp}(i-1)
\end{aligned}
$$

`end`

`if` $P_{sp}(i) > P_{sp\_max}$ `then`

$$
\begin{aligned}
P_{sp}(i) &= P_{sp\_max} \\
k(i) &= P_{sp\_max} - P_{sp}(i-1)
\end{aligned}
$$

`end`

(c) if $P_{sp}(i)$ has been tried before, choose new P_{sp} halfway between $P_{sp}(i-1)$ and previously tried value:

```
for j = 1 to i
        if P_sp(i) = P_sp(j) then
                k(i)     = 1/2 [P_sp(j) - P_sp(i-1)]
                P_sp(i)  = P_sp(i-1) + k(i)
        end
end
```

$$\texttt{for } j = 1 \texttt{ to } i \quad \texttt{if } P_{sp}(i) = P_{sp}(j) \texttt{ then} \quad k(i) = \tfrac{1}{2}\left[P_{sp}(j) - P_{sp}(i-1)\right], \quad P_{sp}(i) = P_{sp}(i-1) + k(i) \quad \texttt{end} \quad \texttt{end}$$

$\boxed{\text{At end of controller's step, once estimate of cost is available:}}$

3. Test new point for improvement:

(a) If new point is an improvement, accept it and double step length:

$$\texttt{if } f\left[P_{sp}(i)\right] < f(P'_{sp}) \texttt{ then}$$
$$P'_{sp} = P_{sp}(i)$$
$$k(i+1) = 2k(i)$$
$$\texttt{end}$$

(b) If new point is not an improvement, leave P'_{sp} as it was, change direction of search, and halve step length:

$$\texttt{if } f\left[P_{sp}(i)\right] > f(P'_{sp}) \texttt{ then}$$
$$k(i+1) = -\tfrac{1}{2}k(i)$$
$$\texttt{end}$$

4. Return to step 1.

After the controller has made four steps in the set point (so $i > 4$), set P_{sp} equal to P'_{sp} for the remainder of the search and move on to the set point of the next turbine.

References

ACSL (1991). *ACSL: Advanced Continuous Simulation Language.* Mitchell and Gauthier Associates Inc., 200 Baker Ave., Concord MA 01742-2100, USA, 10.0 edition.

Ancona, D., Loose, R., and Cadogan, J. (1993). The United States wind energy program in a decade of change. In *Proc. European Community Wind Energy Conference*, pages 74–77. H. S. Stephens & Associates.

Antoniou, I. (1991). Turbulence measurements on top of a steep hill. In *Proc. European Wind Energy Conference*, pages 755–761. Elsevier.

Bendat, J. S. and Piersol, A. G. (1986). *Random Data.* Wiley, 2nd edition.

BLADES (1988). *BLADES Wind Turbine Aerodynamic Analysis.* Garrad Hassan and Partners, Coach House, Folleigh Lane, Long Ashton, Bristol BS18 9JB, 3b edition.

Bossanyi, E. A. (1979). Wind characteristics and the output of wind turbines. In *Proc. 1st British Wind Energy Association Conference, Cranfield.* Multiscience.

Bossanyi, E. A. (1981). Wind turbines in a turbulent wind: Energy output and the frequency of shut-downs. *Wind Engineering*, 5(1):12–28.

Bossanyi, E. A. (1985). Short-term wind prediction using Kalman filters. *Wind Engineering*, 9(1):1–8.

Bossanyi, E. A. (1992). Private communication. Wind Energy Group Ltd, Greenford House, 309 Ruislip Road East, Greenford, Middlesex UB6 9BQ, UK.

Bossanyi, E. A., Smith, G. J., Leithead, W. E., and Agius, P. (1992). 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.

Bouwmeester, R. J. B., Schlueter, R. A., and Park, G. L. (1985). Simulation and assessment of power variations from wind arrays during local storms. *Journal of Wind Engineering and Industrial Aerodynamics*, 21:207–221.

BS 436:Part 3 (1986). *BS 436:Part 3:1986, Spur and Helical Gears*. British Standards Institution, 2 Park St., London.

Compact Orbital Gears (1991). *Uprating of MS-3 to MS-3A specification*. Compact Orbital Gears Ltd., Brynberth Industrial Estate, Rhayader, Powys.

de Boer, W. W., Mortier, M. W., Romar, P. A., and Smithuis, G. J. F. (1991). Measurement results of the wind turbines of the experimental wind farm (SEP) in Friesland. In *Proc. European Wind Energy Conference*, pages 63–67. Elsevier.

de la Salle, S. A., Reardon, D., Leithead, W., and Grimble, M. J. (1990). Review of wind turbine control. *International Journal of Control*, 52(6):1295–1310.

De Wilde, M., Lauriks, L., Vanderborght, P., Gabriels, M., Van Melckebeke, W., Van Steenkiste, S., Van Leuven, J., Cosaert, D., Stevens, D., Van den Poel, R., De Ceulaerde, J., and Verbrugghe, R. (1988). The windfarm in the outer harbour of Zeebrugge – 9 months of experience. *Journal of Wind Engineering and Industrial Aerodynamics*, 27:199–211.

Dekens, M. (1990). Private communication. Administration for Electricity and Electromechanics, Ministry of Public Works, Wetstraat 155, 1040 Brussels.

Dixon, L. C. W. (1972). *Nonlinear Optimisation.* English Universities Press.

Downing, S. D. and Socie, D. F. (1982). Simple rainflow counting algorithms. *International Journal of Fatigue*, pages 31–40.

Draper, C. S. and Li, Y. T. (1951). *Principles of Optimalizing Control Systems and an Application to the Internal Combustion Engine.* ASME.

Eggleston, D. M. and Stoddard, F. S. (1987). *Wind Turbine Engineering Design.* Van Nostrand Reinhold.

Ekelund, T. and Schmidtbauer, B. (1993). Trade-off between energy capture and dynamic loads in variable speed wind turbines. In *Proc. International Federation of Automatic Control Congress*, pages 521–524 (vol. 7).

Estanqueiro, A. I., Aguiar, R., Gil Saraiva, J., Rui M. G. Castro, and Ferreira de Jesus, J. M. (1993). The development and application of a model for power output fluctuations in a wind park. In *Proc. European Community Wind Energy Conference*, pages 798–801. H. S. Stephens & Associates.

Föllings, F. J. (1989). Economic optimization of wind power plants. In *Proc. European Wind Energy Conference, Glasgow*, pages 983–987. Peter Peregrinus.

Föllings, F. J. and Pfeiffer, A. E. (1988). Economics of wind power plants. *Journal of Wind Engineering and Industrial Aerodynamics*, 27:263–274.

Fordham, E. J. (1985). The spatial structure of turbulence in the atmospheric boundary layer. *Wind Engineering*, 9(2):95–133.

Frost, W., Long, B. H., and Turner, R. E. (1978). Engineering handbook on the atmospheric environmental guidelines for use in wind turbine generator development. Technical Report TP-1359, NASA.

Gamble, C. R. (1990). Private communication. Wind Energy Group Ltd, Greenford House, 309 Ruislip Road East, Greenford, Middlesex UB6 9BQ, UK.

Garrad, A. D. and Hassan, U. (1986). The dynamic response of wind turbines for fatigue life and extreme load prediction. In *Proc. European Wind Energy Conference, Rome*, pages 401–406.

Garrad, A. D. and Hassan, U. (1988). What is power quality? In *Proc. 10th British Wind Energy Association Conference*, pages 137–141, London, U.K.

Garrad Hassan (1989). Characterisation of wind turbine turbulence and its implications on wind farm spacing. Technical Report ETSU WN 5096, ETSU.

Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley.

Haines, R. S. (1993). Private communication. Wind Energy Group Ltd, Greenford House, 309 Ruislip Road East, Greenford, Middlesex UB6 9BQ, UK.

Halliday, J. A. (1983). A study of wind speed statistics of 14 dispersed UK meteorological stations with special regard to wind energy. Technical Report RL-83-124, Rutherford Appleton Laboratory, Energy Research Support Unit, Chilton, Didcot, Oxon OX11 OQX.

Harris, A. (1992). A sensitivity analysis of windfarm energy production. In *Proc. 14th British Wind Energy Association Conference*. Mechanical Engineering Publications.

Hassan, U., Glendinning, A. G., and Morgan, C. A. (1990). A wind tunnel investigation of the wake structure and machine loads within small wind turbine farms. In *Proc. 12th British Wind Energy Association Conference*, pages 47–52. Mechanical Engineering Publications.

Hausfeld, T., Hock, S. M., and Thresher, R. W. (1987). Preliminary results from the dynamic response testing of the Westinghouse 600-kW wind turbine. In *Proc. Windpower '87*, pages 7–13.

Hicks, R. (1992). Private communication. Compact Orbital Gears (1989) Ltd., Brynberth Industrial Estate, Rhayader, Powys.

Holley, W. E. and Madsen, P. H. (1985). A method for estimating fatigue damage rates for wind turbine structures. In *Proc. 4th. ASME Wind Energy Symposium*, pages 55–64. ASME, New York.

Hu, J., Beans, J. P., and Auslander, D. M. (1991). Start/stop control of fixed-pitch wind energy turbines. *Solar Energy*, 46(1):29–40.

Jacobs, O. L. R. (1993). *Introduction to Control Theory*. Oxford University Press, 2nd edition.

Javid, S. H., Hauth, R. L., Younkins, T. D., Reddoch, T. W., and Barnes, P. R. (1985). A method for determining how to operate and control wind turbine arrays in utility systems. *IEEE Trans. on Power Apparatus and Systems*, PAS-104(6):1335–1341.

Justus, C. G. (1985). Wind energy. In Houghton, D. D., editor, *Handbook of Applied Meteorology*, chapter 33, pages 915–944. Wiley.

Katić, I., Højstrup, J., and Jensen, N. O. (1986). A simple model for cluster efficiency. In *Proc. European Wind Energy Conference, Rome*, pages 407–410.

Kreyzig, E. (1983). *Advanced Engineering Mathematics*. Wiley, 5th edition.

Landman, H. (1991). Private communication. N. V. Energiebedrijf IJsselmij (energy distribution company), Zwolle, The Netherlands.

Leithead, W. E. (1989). Private communication. Industrial Control Unit, University of Strathclyde, Marland House, 50 George St., Glasgow G1 1QE.

Leithead, W. E., de la Salle, S., and Reardon, D. (1991). Role and objectives of control for wind turbines. *IEE Proceedings Part C*, 138(2):135–148.

Leithead, W. E., de la Salle, S. A., Reardon, D., and Grimble, M. J. (1992). Wind turbine control system modelling and design (phases I and II). Technical Report ETSU WN 5108, Industrial Control Unit, University of Strathclyde, Marland House, 50 George St., Glasgow G1 1QE. Prepared for ETSU.

Leithead, W. E., Rogers, M. C. M., and Agius, P. R. D. (1993). Independent blade control of HAWT's. In *Proc. European Community Wind Energy Conference*, pages 522–525. H. S. Stephens & Associates.

Leithead, W. E., Wilkie, J., and Anderson, C. (1989). Simulation of wind turbines by simple models. In *Proc. European Wind Energy Conference*, pages 336–340. Peter Peregrinus.

Lindley, D. (1993). Wind energy – where are we now? In *Proc. European Community Wind Energy Conference*, pages 1–5. H. S. Stephens & Associates.

Lissaman, P. B. S., Foster, D. R., and Rumbaugh, J. H. (1989). Technical description of AVENU (a comprehensive computer software system for wind farm design). In *Proc. SOLAR 89 – Annual Conference of the American Solar Energy Society*, pages 290–295.

Madsen, P. H. (1988). Power regulation options for wind turbines. In *Proc. European Community Wind Energy Conference, Herning, Denmark*, pages 1–7. H.S. Stephens & Associates.

Madsen, P. H. and Frandsen, S. (1984a). Pitch angle control for power limitation. In *Proc. European Wind Energy Association Conference, Hamburg*, pages 612–619. CEC, UNESCO and EWEA.

Madsen, P. H. and Frandsen, S. (1984b). Wind-induced failure of wind turbines. *Engineering Structures*, 6:281–287.

MATLAB (1992). *MATLAB Numeric Computation and Visualization Software.* Mathworks Inc., Cochituate Place, 24 Prime Park Way, Natick, Mass. 01760, USA, 4.0a edition.

Mattsson, S. E. (1984). *Modelling and Control of Large Horizontal Axis Wind Power Plants.* PhD thesis, Lund University, Dept. of Automatic Control, Box 118, S 221 00 Lund, Sweden.

McNerney, G. M. (1980). Control algorithm investigations. In *Proc. Vertical Axis Wind Turbine (VAWT) Design Technology Seminar*, pages 124–136.

Merritt, H. E. (1971). *Gear Engineering*, pages 234–289; 303–308. Pitman.

Milborrow, D. J. (1988). Technical issues in windfarms and supporting activities within the International Energy Agency (IEA). *Journal of Wind Engineering and Industrial Aerodynamics*, 27:309–318.

Miner, M. A. (1945). Cumulative damage in fatigue. *ASME Journal of Applied Mechanics*, 12(3):A–159 to A–164.

Mutone, G. A., Klein, F. F., and Milnes, R. D. (1986). Control of the Westinghouse WWG-0600 wind turbine farm. In *Proc. 5th. ASME Wind Energy Symposium*, pages 153–157. ASME, New York.

Nascimento, C. L., Zarrop, M. B., and Muir, A. (1992). A neural network extremum controller. Internal Report 774, U.M.I.S.T. Control Systems Centre, P.O. Box 88, Sackville Street, Manchester M60 1QD, UK.

Nierenberg, R. (1989). Terrain effects, wake deficits and macro-scale wake effects. In *Proc. European Wind Energy Conference, Glasgow*, pages 978–982. Peter Peregrinus.

Pacific Gas & Electric Co./US Windpower (1992). *Altamont Pass Wind Power Plant.* US Windpower Inc., 6952 Preston Ave., Livermore, CA 94550, USA.

Paes, H. F. M. and Nauta, Y. (1988). IJsselmij realises Europe's largest windfarm (7.5 MW) on the dike of the Noordoostpolder. *Journal of Wind Engineering and Industrial Aerodynamics*, 27:455–467.

Pal Arya, S. (1988). *Introduction to Micrometeorology.* Academic Press.

Peltola, E., Neuvonen, H., and Lemström, B. (1992). Simulations and measurements of the effect of a wind farm on the local grid. In *Proc. European Wind Energy Association Special Topic Conference '92*, pages D6–1 to D6–7. The Association of Danish Windmill Manufacturers, DK-7400 Herning, Denmark.

Poularikas, A. D. and Seely, S. (1991). *Signals and Systems.* PWS-KENT, 2nd edition.

Quarton, D. C. (1990). The impact of pitch control on wind turbine fatigue. In *Proc. 12th British Wind Energy Association Conference*, pages 77–83. Mechanical Engineering Publications.

Quarton, D. C. and Morgan, C. A. (1989). Wind turbine loading and fatigue damage in wake flow. In *Proc. European Wind Energy Conference*, pages 503–508. Peter Peregrinus.

Rizzi, M. P. and Auslander, D. M. (1990). Effect of control algorithms on fixed pitch wind turbine generators. *Trans. ASME; Journal of Solar Energy Engineering*, 112:147–152.

Roberts, P. D. and Lin, J. (1991). Potential for hierarchical optimising control in the process industry. In *Proc. European Control Conference, Grenoble*, pages 213–217.

Rossen, T. (1990). Optimierung der Energieabgabe eines Windenergieparks [Optimisation of the energy output of a windfarm]. Diplomarbeit, Institut für Regelungstechnik, Technische Universität Braunschweig, Hans-Sommer-Straße 66, Postfach 33 29, 3300 Braunschweig, Germany.

Schlueter, R. A., Park, G. L., Bouwmeester, R., Shu, L., Lotfalian, M., Rastgoufard, P., and Shayanfar, A. (1984). Simulation and assessment of wind array power variations based on simultaneous wind speed measurements. *IEEE Trans. on Power Apparatus and Systems*, PAS-103(5):1008–1016.

Schlueter, R. A., Park, G. L., Lotfalian, M., Shayanfar, H., and Dorsey, J. (1983a). Methods of reducing wind power changes from large wind turbine arrays. *IEEE Trans. on Power Apparatus and Systems*, PAS-102(6):1642–1650.

Schlueter, R. A., Park, G. L., Modir, H., Dorsey, J., and Lotfalian, M. (1983b). Analysis and simulation of storm induced WECS generation changes on system operation. *IEEE Trans. on Power Apparatus and Systems*, PAS-102(1):144–152.

Schlueter, R. A., Sigari, G., and Costi, A. (1986). Wind array power prediction for improved operating economics and reliability. *IEEE Trans. on Power Systems*, PWRS-1(1):137–142.

Simburger, E. J. and Cretcher, C. K. (1983). Load following impacts of a large wind farm on an interconnected electric utility system. *IEEE Trans. on Power Apparatus and Systems*, PAS-102(3):687–692.

Sisouw de Zilwa, L. G. and Bakema, G. F. (1991). Experience with a 15 MWe wind-farm in The Netherlands. In *Proc. European Wind Energy Conference, Amsterdam*, pages 316–320. Elsevier.

Smith, D. R. (1993). Private communication. Research and Development, Pacific Gas and Electric Company, 3400 Crow Canyon Road, San Ramon, CA 94583, USA.

Smith, G. J. (1988). MS–3 wind turbine generator power ramp notes. Draft Technical Report WEG 33-4514, Wind Energy Group Ltd., Greenford House, 309 Ruislip Road East, Greenford, Middlesex UB6 9BQ, UK.

Smith, G. J. (1990). Private communication. Wind Energy Group Ltd., Greenford House, 309 Ruislip Road East, Greenford, Middlesex UB6 9BQ, UK.

South Western Electricity Board (1988). *Use of System Tariff.*

Spruce, C. J. (1988a). Introduction of PC-MATLAB for control system analysis and design at WEG. Technical Report WEG T-4118, Wind Energy Group Ltd., Greenford House, 309 Ruislip Road East, Greenford, Middlesex UB6 9BQ, UK.

Spruce, C. J. (1988b). Machine Features: Wind Energy Group MS-3 prototype wind turbine. *Windirections*, VIII(3):20–22. ISSN 0950-0642.

Spruce, C. J. (1990). Machine Features: Ecotéchnià 20/150. *Windirections*, IX(4):15–18. ISSN 0950-0642.

Spruce, C. J. and Dexter, A. L. (1991). Windfarm simulation for testing supervisory control algorithms. In *Proc. European Wind Energy Conference, Amsterdam*, pages 638–642. Elsevier.

Spruce, C. J. and Dexter, A. L. (1992). Supervisory control of windfarms. In *Proc. European Wind Energy Association Special Topic Conference '92*, pages D10–1 to D10–7. The Association of Danish Windmill Manufacturers, DK-7400 Herning, Denmark.

Steinbuch, M., de Boer, W. W., Bosgra, O. H., Peters, S. A. W. M., and Ploeg, J. (1988). Optimal control of wind power plants. *Journal of Wind Engineering and Industrial Aerodynamics*, 27:237–246.

Stig Øye, A. F. M. (1986). Unsteady wake effects caused by pitch-angle changes. IEA Wind Section Expert Group Meeting.

Symanski, D. P. (1983). The operating experiences and performance characteristics during the first year of operation of the Crotched Mt. New Hampshire windfarm. *IEEE Trans. on Power Apparatus and Systems*, PAS-102(6):1637–1641.

Thomsen, K., Petersen, S. M., Lading, P., and Sangill, O. (1992). Analysis of loads on wind turbines in complex terrain wind farms. In *Proc. European Wind Energy Association Special Topic Conference '92*, pages D2–1 to D2–8. The Association of Danish Windmill Manufacturers, DK-7400 Herning, Denmark.

Toussaint, P., Hutting, H. K., Cleijne, J. W., and Mortier, M. (1992). Results from operation and research of the experimental wind farm of the Dutch electricity generating board. In *Proc. European Wind Energy Association Special Topic Conference '92*, pages B6–1 to B6–7. The Association of Danish Windmill Manufacturers, DK-7400 Herning, Denmark.

Valpy, B. A. (1992). Private communication. Wind Energy Group Ltd, Greenford House, 309 Ruislip Road East, Greenford, Middlesex UB6 9BQ, UK.

Van der Hoven, I. (1957). Power spectrum of horizontal wind speed in the frequency range from 0.0007 to 900 cycles per hour. *Journal of Meteorology*, 14:160–164.

van der Snoek, L. (1988). Development and validation of wind farm models. *Journal of Wind Engineering and Industrial Aerodynamics*, 27:225–236.

Van Leuven, J., Stevens, D., and Van der Snoek, L. (1989). Conclusions from power measurements at the wind farm in Zeebrugge. In *Proc. European Wind Energy Conference, Glasgow*, pages 289–293. Peter Peregrinus.

van Oort, H. (1988). A numerical model for calculating the power output fluctuations from wind farms. *Journal of Wind Engineering and Industrial Aerodynamics*, 27:127–137.

Vermeulen, P. E. J., Builtjes, P. J. H., and Vijge, J. B. A. (1981). Mathematical modelling of wake interaction in wind turbine arrays (parts I and II). Technical Report 81-01473 & 81-02834, MT-TNO, P.O. Box 342, 7300 AH Apeldoorn, The Netherlands.

Vestas-DWT (1992). Private communication. Velling Mærsk windfarm, Ringkøbing, Jutland, Denmark.

Warren, J. G. (1990). Private communication. Wind Energy Group Ltd, Greenford House, 309 Ruislip Road East, Greenford, Middlesex UB6 9BQ, UK.

Wefelmeier, K. (1993). Supervisory control of a wind park. In *Proc. European Community Wind Energy Conference*, pages 571–574. H. S. Stephens & Associates.

WEG 33-4115 (1991). *MS-3 Operational and Shutdown Hub, Nacelle and Tower Loads (Appendix A).* Wind Energy Group Ltd., Greenford House, 309 Ruislip Rd. East, Greenford, Middx., UB6 9BQ.

Wellstead, P. E. and Scotson, P. G. (1990). Self-tuning extremum control. *IEE Proceedings Part D*, 137(3):165–175.

Zarrop, M. B. and Rommens, M. J. J. J. (1993). Convergence of a multi-input adaptive extremum controller. *IEE Proceedings Part D*, 140(2):65–69.