Wind Energy

Will McLennan
Based on lectures by Dr. C. Anderson

Short Description

Wind energy is the fast growing renewable source for electricity generation. This course presents an introduction to the scientific and operational aspects of harnessing this resource. The objective is to present an overview of wind energy covering all aspects from planning a wind farm to how a wind turbine works. The main emphasis of the course is resource assessment starting with basic meteorology and understanding how this relates to wind farm design. Economics of wind energy are driven by both long and short-term variations in wind speed. This links to wind turbine control and grid issues. Finally other aspects such as environmental impacts will be detailed.

Summary of Intended Learning Outcomes

On completion of this course the student will:

- Be capable of conducting a wind resource estimation and be knowledgeable regarding methodologies for resource assessment and sources of uncertainty.
- Understand the basics of wind turbine design and operation.
- Understand mechanism for integrating wind energy developments into the electricity generating and distribution network.

Contents

1 Evolution of modern wind turbines

The earliest use of windmills appears to have been in China and the Middle East (particularly Mesopotamia and Persia), and windmills first appeared in Europe in mediaeval times. A key step was the introduction of the step-up gearbox to achieve the high rotation speeds needed for electric generators, but low speeds at the wind turbine rotor. One of the great advances in the 20$^{\text{th}}$ century development of wind turbines in Europe and the USA was the refinement of the wind turbine rotor, which moved away from traditional sail-based blades to designs more akin to aircraft wings.

Worldwide interest in wind energy took off in the 1970s and 80s in the light of uncertainty over oil supplies and the safety of nuclear power. The USA introduced a market for wind-generated electricity via a tax credit system, this acted as a stimulus to the progressive development of then-small wind turbine designs from the US and Europe.

Compared with a small wind turbine from 30 years ago, a large modern machine may have benefited from advances in some or all of the following areas:

- Rotor blade structural design: Improved blade root designs have greatly reduced stresses at the blade/hub interface. Aerofoil profiles with greater section thickness have enabled overall lighter blade construction.
- Aerodynamic design: development of aerofoil profiles with improved lift to drag ratios has yielded higher power output and smoother stall properties. Attention to trailing edge thickness and 3D tip design has reduced rotor aerodynamic noise.
- Composite materials: use of composites such as GRP and CFRP instead of metal has brought significantly greater structural efficiency and fatigue resistance. Unidirectional fibre lay-ups optimise blade stiffness and strength while reducing weight.
- Blade pitch control: fixed-speed stall regulation is simple and effective at small scale, but power control and blade aerodynamic damping are inadequate for larger machines, and energy yield is compromised. Blade pitch control enables higher energy yield and superior control of blade loading.
- Generator speed control: single-speed induction generators were predominant on the first small stall-regulated wind turbines, but were soon superseded by 2-speed variants using 4/6 pole generators. More recently, broad-range variable speed control has been enabled by the use of doubly-fed induction generators (DFIGs) or synchronous generators with AC-DC-AC connection.
- Continuous research and development: Evolution of wind turbine designs has benefited from knowledge gained either from research prototypes, or from commercial machines. Often, advances in understanding, and development of better technology, have been driven by problems experienced on new production types.

2 The wind and its characteristics

The fundamental origin of the wind is solar heating of the earth, giving rise to differential pressure distributions in the atmosphere. Due to the earth's rotation and inclination, however, and the effect of surface interaction, the resulting flow patterns are relatively complex. The basic flow model is expansion of heated air at the equator and contraction of cold air at the poles, giving rise to convective currents (air rising at the equator and sinking at the poles). The basic driving force is the pressure gradient between the different regions.

This simple model is however greatly influenced by the rotation of the earth, which has surface speed of 1000mph at the equator but zero at the poles. The resulting Coriolis force causes the airflow to veer strongly, depending on the latitude and direction. In addition inertial and viscous effects impose other forces on the moving air mass that further complicate its motion. Nevertheless the wind can be characterised as behaving in a combination of predictable and random behaviour.

For instance the annual average wind speed at a particular point on the globe can be estimated with a high degree of confidence many years ahead. The actual wind speed on a particular day a few weeks away is, however, much harder to predict. It is the long-term predictability of the wind that makes it a valuable energy resource.

2.1 Balance of atmospheric forces

The principal atmospheric forces that determine wind patterns are:

- Pressure gradients
- Coriolis force
- Centrifugal force
- Friction

Pressure force acts from areas of high pressure to low, at right angles to lines of constant pressure known as isobars. Moving air experiences an apparent deflection, however, when observed in the earth's rotating frame of reference. This is due to the Coriolis force, which is sometimes referred to as a 'fictitious' force as it is really an artefact of observing linear motion from within a rotating frame of reference; nevertheless it is valid and useful when analysing motion in a rotating frame.

The strength of the Coriolis force depends on the speed of the air motion, the rate of rotation of the earth (a constant), and the latitude of the air particle in question. The Coriolis force tends to pull a particle of air at right angles to its motion, hence to the left in the southern hemisphere and to the right in the Northern hemisphere.

Under the influence of the above two forces the wind speed increases until the pressure and Coriolis forces are in balance, at which point the direction of motion is aligned parallel to the isobars. The resulting motion is known as the geostrophic wind. It is typically used to estimate wind speed at the top of the boundary layer, but is not representative of the wind at the earth's surface. As isobars are generally curved the wind is also subject to centrifugal force, acting outwards from the centre of curvature. The resulting wind again flows parallel to the isobars, and is known as the gradient wind. So far, the wind is here defined by ideal fluid mechanics without friction or loss.

The final force to consider is, however, due to friction at the interface between the atmosphere and the earth's surface, known as the boundary layer. In this region the flow is retarded by the roughness of the earth's surface, which gives rise to a wind gradient rising from zero speed for molecules directly in contact with the surface, to the gradient wind velocity at a height of around 2km above the ground. The vertical wind gradient or shear profile is strongly dependent on the scale of surface roughness.

Other influences may determine atmospheric circulation patterns at a more detailed level. Differences in the earth's surface, most notably between sea and land, lead to differential amounts of heating or cooling that can give rise to seasonal or daily wind patterns. Another major factor is orography, the term used to describe undulations on the surface of the earth of any size and shape, from small hills to major mountain ranges.

The influence of orography on local wind conditions may include the following:

- Sheltering effect behind hills or in valleys
- Creation of turbulence due to flow separation
- Speed-up of the airflow over hills
- Deflection of the wind from its prevailing direction

Prospective wind turbine sites may experience some or all of the above effects, emphasising the need for careful site assessment, and usually for detailed local wind measurements. Some sites, e.g., very flat terrain with few features, can be characterised fairly easily. Others, such as complex terrain in mountainous areas, may exhibit large difference in wind conditions within a relatively small area. It is essential to be aware of the local influences on wind conditions.

2.2 Time and space scales

Atmospheric motions are in general characterised by variation in time and scale. The two are related in that large-scale wind patterns also tend to occur on a long time base. Examples would be trade winds, or westerlies, which occur over periods of months and affect large areas. An entire country the size of the UK can be simultaneously subject to SW gales during the winter months. For convenience, timescales can be subdivided into the following categories:

- Inter-annual
- Annual (seasonal)
- Time of day (hourly)
- Short-term (gusts and turbulence)

These categories have different implications for wind turbine site assessment. Knowledge of inter-annual variation is important for economic predictions, where a wind farm may be designed to operate for 20 years or more. It is important to know what the long-term average wind speed is likely to be, so that the success of the project is not compromised.

As a rule of thumb at least five years data are considered necessary to accurately estimate the long-term average wind speed for a site. A single year's data is sufficient to estimate the mean wind speed only to within 10% accuracy and 90% confidence level. Given the non-linear relationship between wind speed and power, such accuracy may not be sufficient: this also depends on the level of financial risk involved in the project the higher the risk, the greater the accuracy needed in a wind assessment.

Annual variation refers to the difference in seasonal or monthly wind speeds. In NW Europe the trend is for the highest wind speeds in the winter months and lowest in the summer. Significant monthly variations may occur from year to year, but the annualised trend is relatively predictable. The difference in mean wind speed between summer and winter months can be a factor of 2, and is therefore significantly greater than the inter-annual variation.

Hourly (or time of day) variation in wind speed may be dominated by random influences, but averaging can reveal underlying trends. The most common of these is diurnal variation due to the temperature change between day and night, caused by the solar radiation cycle. Typically the daytime wind speed tends to be higher, with the lowest speeds between midnight and sunrise. Diurnal variation is noted in both temperate and tropical latitudes, and tends to be greatest during spring and summer, and smallest in winter.

Very long term inter-annual variation in mean wind speeds have been studied and some long-term trends were identified; there are even suggestions that climate change may adversely influence the mean wind speed in some areas, though over a long time scale. Long-term comparisons for any given region must be made with care, however, to ensure consistency: the type and location of the measuring instruments may, for instance, be changed several times during a lengthy measurement period (e.g. 50 years).

2.3 Turbulence, surface roughness and shear

The natural wind consists of turbulent fluctuations superimposed on a quasi-steady mean. A measure of the gustiness of the wind is the longitudinal turbulence intensity I, defined as the ratio of the standard deviation σ to the mean speed U. Although turbulence is 3-dimensional, the horizontal variation in wind speed tends to be the most significant, and of greatest relevance to wind engineering. Turbulence intensity generally decreases with increasing wind speed, ie the ratio s/U is greatest in low winds. Turbulence intensity also decreases with height above the ground, as we move further out of the atmospheric boundary layer and towards the undisturbed freestream.

The degree of turbulence is strongly influenced by local surface roughness, essentially the 'lumpiness' of the terrain over which the wind is blowing: the greater the surface roughness, the higher will be the turbulence. Different landscape types can be categorised by their roughness length z_0, which is related to the vertical dimensions of the terrain features. Very smooth landscapes such as ice or still water have low values ($z_0 = 0.002m$) while trees and buildings have relatively high values ($z_0 = 0.4 - 1.6m$).

Roughness length influences both turbulence and wind shear (the vertical gradient of wind speed in the boundary layer). Wind speed U at height z above ground can be estimated using the Prandtl logarithmic law model[3] (noting that u^* is the friction velocity and$K$ the von Karman constant = 0.4.

$$U = (u^*/K)\ln(z/z_0) \tag{2.1}$$

A more useful form of this allows wind speed to be extrapolated with height:

$$U_1/U_2 = \ln(z_1/z_0)/\ln(z_2/z_0) \tag{2.2}$$

The above equation would be used to extrapolate from a wind speed measured at, say, 10m above ground to the 50m hub height of a proposed wind turbine, taking account of local surface roughness. A frequently-used alternative is the empirical (power law) wind shear equation based on shear index, a:

$$U_1/U_2 = (z_1/z_2)a \tag{2.3}$$

Turbulence intensity I also varies with height z above ground, depending on roughness length according to:

$$I(z) = \{\ln(z/z_0)\}^{-1} \tag{2.4}$$

This allows turbulence intensity measured at one height to be extrapolated to another, by intermediate calculation of the roughness length z_0. Note that all of the above equations are widely used in wind siting and engineering practice. Empirical values for roughness length z_0 and shear index a may be found in tables in most text books.

2.4 Weibull wind distribution

An important characteristic of any wind site is the wind speed frequency distribution, which determines the number of hours per year a particular wind speed will occur. Most (though not all) real sites can be characterised using the Weibull distribution, which is almost universally used in economic predictions. The Weibull probability density $p(U)$ gives the likelihood of a particular mean wind speed U:

$$p(U) = \left(\frac{kU}{C^2}\right)^{k-1} \exp\left(-(U/C)k\right) \tag{2.5}$$

where C is the characteristic wind speed, and k is the shape parameter. The characteristic and average wind speeds are related by the following:

$$U_{av} = C\Gamma(1 + 1/k) \tag{2.6}$$

Where U_{av} is the average wind speed and Γ is the gamma function. A special case of the Weibull distribution is the Rayleigh distribution for which $k = 2$, and $C = 1.13U_{av}$. This is often used as a first guess for site wind assessment, as many real sites have k values in the range 1.5-2.5. The value of k reflects the amount of variation about the mean: higher values indicate a narrower spread of wind speeds. A useful form of the Weibull function is the cumulative probability of exceedance Q, being the probability that the mean speed exceeds a particular value V such that:

$$Q(> V) = \exp(-(V/C)k) \tag{2.7}$$

The number of hours per year the wind speed lies between an arbitrary lower limit V_1 and upper limit V_2 is then:

$$Q(V_2 > V > V_1) = [\exp(-(V1/C)k) - \exp(-(V2/C)k)] \times 8760 \qquad \text{(hours/year)} \tag{2.8}$$

This allows us to calculate eg the length of time that a wind turbine will spend between two wind speeds: used in conjunction with the turbine output power curve the annual energy yield can then be calculated.

2.5 Site wind assessment

Site wind assessment is needed for two principal reasons, namely:

- To estimate the economic outcome of a wind project
- To verify the suitability of the wind turbines being installed

In both cases the underlying aim is to minimise risk. Energy production is very sensitive to wind speed and, for example, over-predicting the annual mean wind speed by 10% could lead to under-production of energy by 20% or more. If the project economics are very risk sensitive this could lead to serious financial difficulties.

Similarly, to avoid technical problems, the site wind characteristics must be verified as being within the design envelope of the proposed wind turbines. Wind turbine types are designed for a range of possible wind conditions, and not all are suitable for sites with very high mean wind speed or high turbulence. For instance, a turbine may have a relatively large diameter rotor in relation to its rating to capture more energy in low wind regions, but will not be designed for the stresses of extreme sites.

The International Electrotechnical Commission (IEC) has characterised sites according to wind class with prescribed limits in regard to mean wind speed, turbulence intensity, shear, and maximum gust. This system has been widely adopted by wind turbine manufacturers, who design their machines to a particular IEC wind class. It is then important that the site wind conditions are assessed according to the same system, usually by on-site anemometry.

The most commonly used instruments for measuring wind speed are cup anemometers, which are accurate and relatively cheap. The output of the anemometer is usually a pulse signal, whose frequency is proportional to wind speed. Some designs require a power supply, with the pulse generated using an optical sensor, while others are effectively mini-generators based on a rotating permanent magnet. The disadvantages of cup anemometers are (a) they need to be supported on a met mast, and (b) they measure only at a single point.

More modern instruments include ultrasonic anemometers, which have no moving parts, and can simultaneously measure wind speed and direction. Their accuracy is very high but they are expensive, and their use is more common as sensors on wind turbines than on met masts. Like cup anemometers they are restricted to measuring the wind at a single point.

There is growing use of remote measurement techniques such as SODAR (sonic detection and ranging) and LIDAR (light detection and ranging). Both are ground-mounted, and measure wind speed and direction to heights of several hundred metres, and over potentially large volumes of air. They remain very expensive, however, and used primarily for short periods to solve specific micrositing issues. Power requirements are also much higher than for conventional anemometers.

2.6 MCP analysis

Accurate characterisation of the wind conditions at a given site would in principle require several years of measurement to account for the time dependencies of the wind. For instance, it takes about 5 years to establish the long-term mean accurately, and 6-12 months to characterise turbulence intensity and shear. The need for long periods of measurement, however, often conflicts with a desire to plan and build projects in a shorter timescale.

The way round this is to use long-term wind data from established met stations in the same geographical area as the proposed site, in combination with short-term measurements from both the met station and the proposed site. The strategy is usually referred to as measure-correlate-predict (MCP), and the basic steps are:

1. Simultaneous wind measurements (hourly averages) are taken at the proposed site and the reference met station over a period of 6-12 months.

2. The two data sets are correlated via a scatter plot of site wind speed V_{site} against reference wind speed V_{ref}, and the relationship between them found by linear regression.

3. The above relationship is applied to the long-term (5+ years) mean wind speed at the reference site, to yield a long-term wind speed estimate for the proposed turbine site

MCP is a highly useful technique, but relies on a reasonably strong correlation existing between the two sites: usually the reference site would be within 50km or so of the turbine site, but the distance is dependent on the influence of intervening terrain. In practice the MCP procedure also incorporates directional weighting to account for differences between the wind direction pattern experienced during the short-term measurement campaign and the long-term distribution recorded at the met station.

3 Wind turbine aerodynamics

A wind turbine rotor extracts kinetic energy from the air and converts it into mechanical power. The resulting aerodynamic loads affect not only the rotor blades but the entire structure including the drivetrain, nacelle, tower, and foundation - and even the airflow conditions at other nearby wind turbines. An understanding of the interaction between the rotor and the airflow, and the forces produced, is essential for two broad reasons:

(a) Commercial: to estimate the annual energy output of the wind turbine.

(b) Technical: to calculate the loads and deflections that the entire structure will experience and thereby ensure it is adequately designed.

Unlike the drivetrain elements (gearbox and generator), the wind turbine rotor has no close parallels in other applications, and requires a fundamental understanding of both the aerodynamics of rotating blades and the structure of the wind. Modern wind turbine theory has been developed over more than a century, drawing on sources including wing and propeller theory, low-speed aerofoil research, and helicopter airflows, but increasingly it has represented a field of its own.

The most widely used analytic tool for wind turbine design and performance prediction is Blade-Element Momentum (BEM) theory, which combines actuator disc theory with analysis of the loads on individual blade aerofoil sections. The approach uses a combination of fluid mechanics theory and empirical data, including measured lift and drag characteristics for real blade profiles.

3.1 The actuator disc

The actuator disc is a hypothetical energy absorber occupying the swept area of the wind turbine rotor, and oriented normal to the incident flow. The disc is comprised of an infinite number of blades, and flow conditions are assumed to be steady and uniform across its surface, creating a uniform thrust loading. The following assumptions also apply:

- Flow through the disc is homogeneous and incompressible
- Frictional drag is neglected
- There is zero wake rotation
- Ambient pressure exists far upstream and downstream

Analysis of the conditions at the disc surface is based on the further assumption that the flow slows down continuously through the disc, but the pressure undergoes a discontinuity: high pressure is developed locally on the upwind side and low pressure on the downwind. Bernoulli's theorem may be applied separately to an upwind and downwind streamline, and the resulting pressure difference at the disc equated to the rate of change of fluid momentum.

The analysis yields expressions for the power and thrust coefficients, C_P and C_T, as functions of the axial induction factor a, which is the fractional decrease in the freestream velocity at the rotor disc. The net decrease far downstream is $2a$, so that half of the velocity deficit caused by the rotor occurs upwind, and the other half downwind. By setting $dC_P/da = 0$ it can also be shown that the maximum value of C_P occurs for $a = 1/3$, when $C_P = 16/27$. This is the'Betz Limit', the theoretical limit of power extraction.

Although the actuator disc gives useful insights into the mechanism of power extraction, the simplifying assumptions in the model limit its applicability. In particular, the aerodynamic conditions across a real wind turbine rotor are highly non-uniform even in a steady wind, due to the difference in blade velocity between the hub and tip: the blade root velocity is almost zero, while the tip may be travelling at 60m/s or more. This leads to proportionately higher loading on the outboard stations.

To account for this the actuator disc model is replaced by a set of concentric stream tubes, which are assumed to be radially independent (no flow between them) and able to move at different axial velocities. Each tube is an annular version of the actuator disc, obeying the same rules for incremental thrust and power. The axial thrust load on a given annulus at the rotor plane is then given by

$$dF = \tfrac{1}{2}\rho U^2 A_d C_T \tag{3.1}$$

in which A_d is the annular area at radius r and C_T is now a local thrust coefficient, applicable to the stream tube in question and defined as previously:

$$C_T = 4a(1-a) \tag{3.2}$$

where a is likewise a local axial induction factor. This representation allows the actuator disc flow to approximate more closely to that of a HAWT rotor in terms of the spanwise distribution of blade loading.

The concentric stream tube model still, however, assumes uniform flow conditions exist at all parts of the annulus, without accounting for the discrete presence of the blade. This limitation is addressed by simultaneously considering the aerodynamic conditions at a real blade element in terms of the airflow geometry, and resulting lift and drag forces, as now described.

3.2 Blade element theory

The forces on a real blade section are determined by the lift and drag coefficients C_L and C_D, which are functions of the angle of attack, α, which is in turn determined solely by the flow angle ϕ and the local blade pitch angle β according to:

$$\phi = \alpha + \beta \tag{3.3}$$

It is important to consider pitch, as the section chord line does not in general lie flat in the rotor plane because (a) blades are generally twisted to compensate for the variation in flow angle along their length, and (b) the complete blade may in addition be set at a non-zero pitch angle for power control, or when starting or stopping. For the present, pitch angle β is assumed to account for both. The lift on a blade section is defined by:

$$L = \tfrac{1}{2}\rho V_R^2 A C_L \tag{3.4}$$

Neglecting drag (which is very small at low α), the axial thrust loading is given by:

$$F_{thr} \approx \tfrac{1}{2}\rho V_R^2 A C_L \cos\phi \tag{3.5}$$

in which the velocity resultant V_R and flow angle ϕ are functions of the known wind speed V, local tangential velocity ωr, and local pitch angle β, as well as the unknown values of axial induction factor a and tangential induction factor a' (the latter is introduced to account for the rotational motion imparted on the wake by the blade, according to Newtons 2nd law).

Following from the above, evaluation of F_{thr} requires us to find a, a' and lift coefficient C_L. The last of these is found from tabulated wind tunnel measurements for the given profile, noting that C_L is determined only by angle of attack α; as noted above, α can be simply deduced from ϕ and β. Ultimately, then, the only unknowns in the thrust load equation are the induction factors a and a'.

This outcome is similar to that of the earlier stream tube analysis, which also yielded an estimate for F_{thr}. We now equate the elemental thrust loads from the two analyses:

$$2\rho A_d V^2 a(1-a) = \tfrac{1}{2}\rho V_R^2 N A C_L \cos\phi \tag{3.6}$$

where the factor$N$ on the right hand side accounts for the total thrust due to N blades (the left hand side already represents the thrust load on the complete annulus). Now defining local solidity $\sigma = NA/A_d$, and noting that $\sin\phi = V(1-a)/V_R$, the following expression is derived:

$$\frac{a}{1-a} = \frac{\sigma C_L}{4\sin\phi\tan\phi} \tag{3.7}$$

Similar analytic procedures lead to the following two expressions, which, together with the above, fully define the flow characteristics at the blade element:

$$\frac{a'}{1-a'} = \frac{\sigma C_L}{4\cos\phi} \tag{3.8}$$

$$\tan\phi = \frac{1-a}{\lambda_{loc}(1+a')} \tag{3.9}$$

Note the use of λ_{loc} for the local velocity ratio, where $\lambda_{loc} = \omega_r/V$, evaluated at blade element radius r. The above three equations cannot be solved analytically, and an iterative process is used to find consistent values of a, a', and C_L. The following is the typical procedure:

1. Assume initial values for a and a', and evaluate flow angle ϕ
2. Calculate angle of attack α, hence evaluate C_L from tabulated data
3. Recalculate induction factors a and a' , and repeat until convergence

The iterative calculation is repeated for every blade station (radial independence is assumed) to yield final values of the lift and drag coefficients, and local flow velocities. These in turn determine the local thrust and tangential forces F_{thr} and F_{tan}, respectively, which act perpendicular and parallel to the rotor plane. Knowledge of the blade element forces then leads to all the important rotor loads, e.g., axial thrust force, net torque, power and bending moment etc.

The dimensionless rotor torque and power coefficients C_Q and C_P also follow from the above. The BEM code can then be used for specific power curve predictions, or to produce non-dimensional performance curves to explore rotor design. The same code is also used to provide aerodynamic load for structural design purposes.

3.2.1 Limitations of BEM theory

- Breakdown of momentum theory

  A limitation of BEM theory is breakdown of the basic momentum assumptions at high tip speed ratio, when the rotor is highly loaded and induction factors a and a' are large. Reversal of the flow in the wake is predicted, with C_T decreasing to zero. In a real flow, however, the wake becomes turbulent under these conditions, while C_T continues to rise above unity. To account for this limitation an empirical adjustment is made to the thrust curve, using C_T values based on wind tunnel measurements. The conditions in question are of limited interest, however, as most practical wind turbines operate at values of a less than 0.6.

- Radial independence of stream tubes

  The assumption of radial independence between concentric stream tubes is also incorrect, and on a real rotor there may be significant radial flow due to centrifugal effects, particularly when the blade is fully or partially stalled. This is again dealt with using empirical corrections to the code, in this case by adjusting the profile lift/drag data in the post-stall region, based on measurements from full-scale wind turbines. More modern prediction methods such as CFD may be more reliable in this regard, but they tend to be more computationally intensive and slow, and despite its limitations BEM theory is still widely retained in the wind industry.

- Tip loss correction

  Conditions near the tip of the blade are not well modelled by stream tube theory. On a real turbine blade, as on an aircraft wing, there is a significant migration of air from the pressure surface to the suction surface at the tip of the blade, with creation of a tip vortex. The lift on the blade decreases to zero at the tip with a consequent loss of performance over the outer stations; basic BEM theory would by contrast predict a high efficiency from the tip section.

  Tip loss has been extensively researched and modelled, and BEM theory includes mathematically rigorous (non-empirical) corrections based on propeller theory. The prediction code includes a loss factor at each blade element taking account of its radial position relative to the tip, ie r/R. A similar loss factor should in practice be applied at the blade root, but is usually neglected, as the loading in this region is very small due to the low tangential velocity.

3.2.2 Application of BEM theory

Application of BEM theory yields important insights into the fundamental aerodynamic properties of rotors, and the influence of important parameters. The key relationships are as follows:

Tip speed ratio and rotor solidity

Tip speed ratio and rotor solidity share an inverse relationship: the higher the tip speed ratio, the lower the solidity required for optimum performance. Faster running rotors therefore have more slender blades and/or fewer of them. Higher rotational speeds also correspond to lower torque loads, thus enabling lighter drivetrains, and there are additional performance benefits due to reduced tip loss and wake rotation.

Chord distribution

The inverse relationship between tip speed ratio and solidity is also seen within a given blade design. The local velocity ratio at a given blade station ($\lambda_{loc} = \omega r/V$) varies from zero at the root to a maximum at the tip, and the corresponding solidity decreases from the root outwards towards the tip. The optimum chord distribution is nonlinear, with a potentially very large value at the root. Real blade planforms tend, however, to be a compromise enabling simplified structural design. There is little power generated at the blade root in any case (due to the low rotational velocity) so the resulting loss is small.

Twist distribution

The angle of the local airfoil chord line relative to a fixed reference in the blade (usually taken at the tip) is known as the twist angle. It is essentially the structural variation in pitch along the blade length (noting that pitch angle usually refers to the rotation of the complete blade about its long axis, eg for power control). Twist angle is necessary to compensate for the change in local flow angle f along the blade length, noting that:

$$\phi \approx V/\omega r \tag{3.10}$$

hence ϕ tends to 90° at the blade root (where $r = 0$). Now the angle of attack α is given by:

$$\alpha = \phi - \beta \tag{3.11}$$

To ensure a low value of α (typically $< 10°$) and hence a high lift/drag ratio, the local pitch angle β (strictly the twist angle) is varied from a large value at the root to almost zero at the tip. An optimal blade design is highly twisted, but may not be practical to manufacture, so real blades tend to be a compromise at the root section. The energy output at this location is, however, small due to the low rotational speed, so the loss is not usually significant.

Lift and drag

A critical parameter for the aerofoil profile is its lift to drag ratio C_L/C_D, which has a profound influence on the achievable rotor efficiency (particularly at high tip speed ratio). For any blade number or tip speed ratio, however, the higher lift to drag ratio, the higher the performance. Modern profiles can achieve L/D ratio of 140, while the older aircraft-based profiles have a maximum ratio of 85.

Rotor drag is also reduced by avoiding the use of the bracing wires or struts that were common on early wind turbines. Attention should also be paid to the blade root/nacelle interface where flow interference can cause a loss of efficiency. The most recent designs are highly optimised in the blade root region, with consequent benefit in energy capture.

Number of blades

The number of blades on a rotor is not as critical as it may seem, as the C_P performance can in principle be maximised by matching the tip speed ratio and solidity for any case. Hence a 2-blade wind turbine is more or less as effective as a 3-blade in terms of $C_{P(max)}$ and energy yield. Increasing the number of blades does have a slight performance benefit, however, due to reduced tip losses, and other factors also tend to favour 3-blade over 2-blade designs, such as lower noise due to slower tip speed, and more favourable dynamics.

3.3 Aerofoils

The aerofoil lift and drag coefficients used in BEM codes are generally taken from accurate wind tunnel measurements on the relevant 2D sections, measured under conditions designed to minimise end effects (tip loss) or wind tunnel blockage. Generally the data are most reliable for angles below stall; good post-stall data is harder to obtain as tunnel blockage and loss of measurement accuracy become significant factors at high α.

Aerofoil profiles used in wind turbine design are generally designed for low speed (<100m/s) performance and high lift/drag ratio. Requirements for blades may include:

- Reduced peak lift for stall regulated rotors
- Thicker sections for increased structural rigidity near the blade root
- Profiles tailored to accommodate certain construction materials (e.g. wood laminate)

3.4 Non uniform flow

In its simplest form the BEM calculation assumes steady and uniform wind speeds across the rotor, and is typically used to generate a rotor power curve for commercial purposes. The more detailed analyses required for turbine structural design require consideration of nonuniform wind conditions, which give rise to cyclic variation of blade loading. The following are the principal sources of non-uniform inflow:

- Wind shear (boundary layer velocity profile)
- Tower shadow
- Yaw misalignment

Wind shear leads to higher effective wind speeds on the blades during the upper part of their rotation, and a roughly sinusoidal variation in blade loading. Tower shadow is a more impulsive wind speed variation arising when the blade passes in front of the tower, and through a region of locally slower flow. Yaw misalignment results in a distortion of the airflow vector triangle, with again a roughly sinusoidal variation in loading.

As noted above, these influences are important when evaluating blade loading, and in particular fatigue loads: a once-per-rev (1P) variation in blade bending load may occur 250 million times over the lifetime of a large wind turbine. In addition some cyclic blade loads may appear in the output rotor torque, though the larger amplitude cycles due to shear or yaw error tend to cancel out when aggregated at the hub (rather as blade gravity loads cancel out on a balanced rotor).

To account for non-uniform wind effects the blade element (BEM) code must be run for a number of blade positions throughout the rotor cycle, with loads evaluated at each. Typically the assessment may be carried out at $10°$ increments of rotor rotation. This is more time consuming than evaluation of the power curve, which assumes uniform wind across the rotor.

3.5 Vortex wake codes

As noted above, BEM theory has a number of limitations that are dealt with by empirical methods in performance codes. This approach has proved fairly satisfactory in industrial applications to date, where the computational speed of BEM codes is a great advantage.

Such codes represent the blade as either a line vortex, or (more accurately) as a spanwise distribution of vortex segments, that generate the same lift as the blade aerofoil profiles when acted on by the incident wind. The so-called bound vorticity does not terminate on the blade, but is shed into the downstream wake, creating a helical vortex pattern behind the rotor. The wake vorticity influences the inflow characteristics in the rotor plane, and determines the change in wind velocity at the rotor. In this way the induction factors can be calculated along the blade length and the blade loads determined, as previously.

The advantage of vortex wake codes is that they account for the complex 3-dimensional nature of the flow behind the rotor, including radial flows, in a fundamental manner. They are also capable of modelling the time-dependency of the rotor loading: the inflow velocities at the rotor plane are found by integrating the effect of the entire helical wake, which takes time to respond to changes in e.g. blade pitch angle or incident wind velocity. This capability is useful in developing or analysing power control strategies.

The disadvantage of vortex codes is their computational expense, and they require a significantly higher number of calculations than BEM codes to calculate the rotor conditions, even in the steady state. In addition vortex codes are less useful for stalled flow, as their underlying theory assumes ideal flow characteristics. As a result BEM codes tend to be retained for industrial use, despite their limitations. Vortex codes have, however, been successful in explaining the wake characteristics observed in experimental rotor studies.

3.6 Wake effects

The downwind wake of a wind turbine is characterised by a region of reduced wind velocity and increased turbulence. Both factors are important when planning a wind-farm array, in order to avoid excessive loss of output, or introduce high fatigue loading. Generally an optimisation is carried out, whereby wind turbines spacing is as close as possible to minimise infrastructure and land costs, but sufficiently great to avoid the above issues.

Wake interference has been measured on full-scale turbines and scale models, leading to empirical expressions for wake deficit and turbulence intensity for use in optimisation codes. In the near field the wake diameter is slightly larger than the rotor itself, and its effect on adjacent machines will depend strongly on wind direction. At a distance of 5 rotor diameters downstream the velocity deficit on the centreline may be as much as 40% leading to a loss of production on the downwind machine.

The overall array efficiency depends, however, on the annual wind directional distribution with the above figure being factored by the proportion of time spent in the wake, which may be only a few percent of the total. As a rule of thumb wind turbines in arrays should be spaced at least 5 rotor diameters apart, assuming the wind is roughly omnidirectional. If there is a strongly prevailing wind direction then downwind spacing of turbines may have to be increased, whereas crosswind spacing may safely be reduced to as little as 3 diameters.

Array optimisation may also have to account for prevailing levels of turbulence e.g. in complex terrain, where the overall turbulence experienced by a particular turbine will be the aggregate of the ambient turbulence and the wake-induced turbulence from other machines. In such cases more detailed studies may be needed.

4 Structural design and analysis

This section examines the issues involved in the structural design and analysis of horizontal axis wind turbines.

4.1 Sources of loading on a wind turbine structure

Aerodynamic forces originate at the rotor (refer to previous lecture) and are the major source of loading on a wind turbine.

Gravity acts on all parts of the structure, and gives rise to both steady loads, such as the combined weight of the rotor and nacelle acting on the tower top, and cyclic loads, e.g. the blade root gravity bending moment, which reverses once per rotor revolution.

Gyroscopic loading arises principally when the rotor yaws to face the wind, with the magnitude of the resulting loads proportional to the yaw angular velocity. This can be a significant concern for free-yaw wind turbines, which may have high yaw rates.

Centrifugal force affects the rotor, and contributes e.g. to the blade radial loading. Centrifugal force is proportional to $\Omega^2 R$ (where Ω is the rotor angular velocity and R the radius), and it may be a major factor in the design of smaller wind turbines. Centrifugal force is less of a concern on very large wind turbines - though it can be a source of cyclic loading on an imbalanced rotor.

Electromechanical forces originate with the generator. Normally the generator reaction torque is in equilibrium with the rotor input loading, but a major consideration is the short-circuit fault torque, which may be an order of magnitude higher than the rated value. This can lead to an extreme load on the drive train and rotor blades.

4.2 Steady and fluctuating loads

Steady loads are those with constant value or very slow time dependency. The principal examples are mean aerodynamic loads in a constant wind speed, and gravitational loads on the stationary parts of the structure. In practice, however, the wind loading is always a combination of steady and fluctuating (or cyclic) components.

Cyclic loads vary continuously with a short time dependency. They can be subdivided into deterministic and stochastic loads. Deterministic loads are relatively predictable in terms of their magnitude and frequency, and are caused by the interaction of the rotating blades with non-uniform airflows and/or gravity. Deterministic loads occur even in steady wind conditions, as their cyclic nature is due to the rotation of the blades.

Stochastic loads are caused by atmospheric turbulence and vary in a quasi-random fashion, depending on the magnitude and spatial distribution of turbulence at the rotor plane. They also give rise to loading on the rotor with harmonic content at multiples of the rotation frequency this is due to rotational sampling. Examples of deterministic and stochastic load inputs are as follows:

Deterministic

- Once-per rev (1P) blade flapwise load due to wind shear
- 1P blade flapwise load due to tower shadow
- 1P gravity bending moment at blade root
- 3P tower-top thrust load due to tower shadow (3-blade rotor)

Stochastic

- Random variation in blade bending moments
- Random variation in rotor torque and power output

Rotational sampling of the turbulent inflow by the rotating blades gives rise to quasi-random variations in the blade loading, which appear in the blade load spectrum at frequencies that are multiples of the rotor rotational frequency. The effect can be visualised as the blade 'chopping' through a turbulent eddy and generating periodic spikes in the load input.

Rotationally sampled loads appear in the blade load spectrum at frequencies equal to 1P, 2P, 3P, etc with diminishing amplitude. The aggregate loads at the rotor, however, are predominantly at nP where n is blade number (e.g. 3P on a 3-blade rotor) as the other harmonics tend to cancel out at the hub.

4.3 Static and dynamic loading

Steady load inputs give rise to static loading of the structure, whereby the resulting deflections and stresses are dependent only on its stiffness, and there is no time dependency. For instance the mean bending moment at the blade root at a particular wind speed, and the corresponding blade tip deflection, would be analysed as a static load case.

Cyclic load inputs, however, give rise to dynamic loading, whereby the deflection and stress of components are dependent on both stiffness and inertia, and the time-dependent response of the structure to the fluctuating loads. Analysing dynamic loads is of particular importance in regard to fatigue design; for instance the blade root flapwise bending moment during normal operation would be assessed using a dynamic model to extract a load spectrum on which fatigue evaluation could be based.

4.4 Load cases

The load cases to be considered in a wind turbine design can broadly be subdivided into extreme, operational, or fault load conditions. Examples of each are:

Extreme loads

- Maximum gust on the stationary rotor
- Extreme change of wind direction

Operational loads

- Normal power production in a range of wind speeds
- Rotor starting and stopping cycles
- Normal yawing

Fault conditions

- Generator short-circuit
- Emergency braking
- Rotor overspeed

Load predictions in each case may require static or dynamic analysis. In addition the number of occurrences of each load condition may be relevant to the fatigue design of a particular component. For instance, the blade design will include an assumption of a certain number of start-up and shutdown sequences per year, as each cycle adds to the cumulative fatigue damage at key areas such as the blade root, or the attachment bolts.

4.5 Load predictions and modelling

Static and dynamic analyses are generally based on aerodynamic loads generated using blade element momentum (BEM) theory, due to its computational efficiency and reasonable degree of realism (see previous lecture). Static loads are relatively simple to evaluate, by applying the relevant aerodynamic loads to components whose stiffness and strength have been previously estimated.

For dynamic analyses the BEM aerodynamic loads are used as input to a modal model of the complete wind turbine. This is a simplified structural model in which the main components are represented as multi-element beams with limited degrees of freedom and prescribed values of mass, stiffness, and damping. A structure comprised of n elements will in general be characterised by n modes of vibration, whose response to prescribed cyclic load inputs can be explored by time-domain computer simulation.

Modal analysis is a relatively complex procedure, but key to evaluating the deflections and stresses experienced by the complete wind turbine under a wide range of operating conditions. The mode shapes and their frequencies (also known as eigenfrequencies) for the structure are first found from free vibration analysis, which neglects damping effects. The model is then usually simplified in terms of only the most important structural modes, i.e. which primarily dictate the response of the structure. Sometimes further simplification is possible for key components.

For instance, the blade flapwise motion is generally contained in one of the coupled rotor modes, whose frequencies are similar to that of an isolated blade. The blade may then be treated in isolation for some analyses and modelled as a simple single degree-of freedom (DoF) system with mass, stiffness, and damping based on design or measured values. Such a representation is useful for basic engineering purposes and can be helpful, for instance to gain insights into structural and aerodynamic damping effects. The response of a single DoF system with undamped natural frequency ω_n to a harmonic forcing function at frequency ω is given by:

$$Q = \sqrt{(1-r^2)^2 + (2\zeta r)^2} = \frac{\text{Actual vibration amplitude}}{\text{Amplitude under static loading}} \tag{4.1}$$

where r is the frequency ratio, $r = \omega/\omega_n$ and ζ is the non dimensional damping coefficient, $\zeta = c/c_c$.

The blade flapwise response to a 1P cyclic load input such as wind shear or tower can be examined using the single DoF model, where the forcing function is a harmonic load applied at a frequency equal to the rotor angular velocity. The magnitude of the response will be dictated by the frequency ratio r and modal damping coefficient ζ.

4.6 Damping

Components such as the blades and tower have low inherent (structural) damping, with damping coefficient, ζ ,(ratio of actual to critical damping) of 0.5% or less; this is typical for large rigid structures made from steel or composite materials. Low damping can be seen by the length of time e.g. the vibration of a blade on a test stand takes to decay, with the damping coefficient found using:

$$\zeta = \frac{\ln(A_0/A_n)}{2\pi n} \tag{4.2}$$

where A_0 and A_n are the vibration amplitudes at the start of the measurement period, and after n cycles respectively.

When a wind turbine blade is operating in attached flow, however, there is significant added damping due to the action of the airflow. The flapping motion of the blade the predominant dynamic response causes variation in angle of attack, α, proportional to the out-of-plane velocity, and a corresponding variation in lift tending to damp the motion out. The effective damping coefficient may then be of the order 16%, and aerodynamic damping is a valuable influence in limiting dynamic response.

Aerodynamic damping becomes less significant on a stalled blade due to the loss of the linear relationship between α and the lift force. This can be a problem for blade flapwise vibration, but of more concern is the lack of edgewise aerodynamic damping, or occurrence of negative damping - on large stall-regulated rotors, which has required some large HAWT blades to be fitted with internal mechanical dampers, and is one reason that stall regulation is now rarely used on wind turbines larger than a few hundred kW.

4.7 Resonance

Referring again to the response of a single DoF system, the phenomenon of resonance occurs when the cyclic load input frequency is too close to a modal frequency. The response amplitude can then give rise to loads many times higher than the equivalent static load. This is a major consideration in wind turbine design, where there are a variety of cyclic load inputs at different frequencies, and many structural modes of response. The possibility of resonance, and consequent fatigue damage, must be carefully considered at the design stage.

One example is the response of the tower due to the cyclic loads caused by the blades passing, i.e. tower shadow. On a 3-bladed wind turbine this gives rise to a 3P variation in tower-top thrust loading, and it is important that this does not coincide with the lower tower eigen frequencies. This is a major consideration in wind turbine design, as resonance could quickly lead to fatigue at the tower base connection.

Most large HAWTs are designed with 'soft' towers, whose lowest modal frequencies lie well below the 3P excitation frequency, and even below 1P, so that they are not susceptible to tower shadow excitation (at 3P) or unbalanced gravitational loading on the rotor (1P). Careful consideration, however, is needed for variable speed wind turbines, where the rotor excitation frequencies 1P, 2P, 3P etc. are not fixed, but vary with operating conditions. In this case the control system may be programmed to prevent the rotor from dwelling at critical rotation speeds for long, but instead accelerate through them.

5 Electrical aspects

This section is an overview of some of the electrical issues surrounding grid connected wind turbine generators. It is not intended to be too deeply technical, but includes a short refresher on some key electrical concepts. The main topics thereafter are the different types of wind turbine generator and their characteristics, a brief description of distribution and transmission networks, and discussion of various issues in regard to grid connection, and compliance with utility requirements for embedded generators.

5.1 Generator types: fixed speed

- Permanent magnet (off-grid)

  The simplest type of wind turbine generator uses permanent magnets on its rotor, with a conventional wire-wound stator. This arrangement is common on small stand-alone wind turbines, i.e. those not connected to the grid. Such machines generally operate at variable speed and, consequently, with variable output voltage. Control of the rotor speed and the output power are achieved by varying the load, for example by switching in more heating elements as the voltage rises. Large permanent magnet generators are starting to be used in grid connect applications, in combination with variable frequency power converters. In this respect they operate in essentially the same way as synchronous generators (see below).

- Synchronous (fixed speed) In a synchronous generator the magnetic field on the rotor is created by current-fed windings, with electromagnetic poles effectively taking the place of the permanent magnets of a small machine. DC current is fed to the rotor windings via slip rings, or may be generated by a small generator mounted on the rotating shaft.

  The stator has fixed windings, which are connected in parallel with the AC grid. These give rise to a rotating magnetic field, whose angular speed is determined by the number of poles in the stator and the grid frequency.

  The interaction of the rotating magnetic fields due to the rotor and stator give the synchronous machine the characteristic that torque (T_G) is proportional to the rotation angle between the physical rotor and the rotating magnetic field of the stator, also known as the power angle δ. The synchronous generator behaves like a very stiff torsional spring, with torque proportional to the power angle d. When connected, the rotor always rotates at the same speed as the grid connected stator magnetic field, i.e., the synchronous speed. The generator torque, and hence output power, rises from zero when $\delta = 0$ to fully rated when δ equals (typically) 30 degrees. The stiffness of the torque characteristic is an issue for the drivetrain dynamics, and one reason that synchronous generators are not now used on fixed-speed wind turbines.

One very powerful characteristic of the synchronous machine is the ability to control the power factor and electrical phase angle ϕ, by varying the field current fed to the rotor. For this reason synchronous machines are favoured by utilities for large-scale power production, e.g. in coal-fired or nuclear power stations. Synchronous machines are most suitable in cases where the input load and generator speed can be closely controlled. Some of the early large wind turbines in the UK used synchronous generators in conjunction with fixed-speed operation, and output power limiting in high winds via blade pitch control. The choice of synchronous machines was mainly utility-driven, in keeping with traditional grid-management requirements. These turbines were often, however, problematic due to the high torsional rigidity conferred by the generator, which tended to amplify power swing caused by turbulent wind variation. The problem was generally worst during power controlled operation.

As a result fixed-speed synchronous generators are nowadays uncommon on large wind turbines. In recent years, however, the synchronous machine has made a significant return in conjunction with variable speed operation.

- Induction (fixed speed)

  Asynchronous, also known as induction, generators have a similar stator arrangement to synchronous machines, with grid-connected armature windings to create a rotating magnetic field. The rotation speed of the field is again dictated by the grid frequency and number of poles. The rotor of the induction set, however, has no windings and is not supplied with current: instead it comprises a ring of parallel conducting bars arranged like a circular cage (hence the common name squirrel cage) and short-circuited at the ends.

  The magnetic field of the stator rotates at synchronous speed and induces currents, and hence magnetic fields, in the rotor bars. This causes the rotor to turn, just as a permanent magnet would if located within the rotating magnetic field. The torque developed between the stator and rotor is proportional to the difference in their angular speeds, also known as the slip. The rotational speeds are usually quoted in r.p.m., but as slip is non-dimensional any consistent units may be used.

  When the rotor turns at exactly synchronous speed there is no torque developed at the generator, and full load torque is achieved at a slip usually in the range 1-2%. Hence the induction generator rotates very nearly at constant speed and is often referred to as a fixed speed machine; the small difference in speed represented by the slip is, however, essential to its operation.

  Because the generator torque T_G is proportional to the relative speed of the rotor, rather than its angular position, the generator has the characteristic of a viscous damper. Note that this does not mean that electrical energy is lost, but it does give the generator favourable dynamic characteristics: the drive train torsional response to variable torque input is much less with an induction generator than with a synchronous machine, i.e., it is more 'forgiving'.

Combined with the above, the simplicity of construction of the induction machine, and its general ruggedness (without windings on the rotor) lend it well to wind energy applications, and for these reasons the induction generator was the predominant design choice for many years, from far back as the 1950s. The main disadvantage of the induction set is its requirement for reactive power to magnetise the rotor, and an induction generator may run at an inductive power factor as low as 0.8 at full power, and lower still on part load. In addition the transient or inrush current on energisation is very high, typically up to 8 times the rated current value. Thus induction generators in general require both power factor correction, and 'soft-start' mechanisms.

5.2 Power factor correction

The inductive power factor of the asynchronous generator would, if uncorrected, lead to higher reactive power flows in the electricity network. This tendency can be counteracted by power factor correction using capacitors. Capacitive reactance acts in the opposite sense to inductive reactance, and can in principle enable unity power factor (f=0) to be achieved, with zero net reactive power flow.

The reactive power demand of an induction set varies with the load, however, so a fixed value capacitor will only provide correction at one particular operating point. In practice, many stall regulated wind turbines have staged capacitor banks so that power factor correction is applied incrementally depending on the load. Normally it is desirable that the power factor remains slightly inductive, to assist system stability.

5.3 Generator types: variable speed

Most modern wind turbine types are capable of operating with broad range variable speed, so that they can rotate at any speed between 15 and 25 r.p.m. with the generator continuously connected to the grid. The electrical power remains at fixed voltage and frequency despite the speed variation of the generator. This characteristic offers great advantages including:

- Operation of the rotor at high aerodynamic efficiency by tracking C_P^{max}
- Accurate limiting of peak power levels
- Reduced noise in low winds, due to reduced rotor speed

Variable speed generators achieve their capability through the use of power electronics, semiconductor-switching devices that enable essentially digital control of the generator supply characteristics. There are two main types of variable speed generator, namely synchronous and doubly-fed induction machines. Variable-speed synchronous generators are conventional synchronous machines connected to the grid via an AC-DC-AC link: the generator operates at variable AC frequency, which is rectified to DC and fed via an inverter to provide fixed frequency output on the grid. Such systems offer immense flexibility not only in the operation of the wind turbine, but also in the power conditioning at the point of grid connection.

There are a few manufacturers of wind turbines using synchronous variable speed These manufacturers produce direct drive machines, with multi-pole low speed generators, which are directly attached to the wind turbine rotor, dispensing with the need for a gearbox. This is a further advantage, as gearboxes can be a source of unreliability.

The doubly-fed induction generator (DFIG) differs from the variable-speed synchronous machine in that the stator windings are directly connected to the grid in the conventional manner, while the DFIG rotor (unlike a conventional induction set) is wound, and supplied with external power via a variable-frequency power converter. This enables it to operate at limited-range variable speed, and facilitates control of the power factor without the need for capacitor correction.

Broadly speaking the DFIG offers the same operational advantages as the synchronous machine in terms of aerodynamic performance, power control, and noise reduction, but slightly less versatility in terms of grid management. It is, however, more cost effective than the synchronous machine, as the power converter does not require to be fully rated.

5.4 Starting and stopping

Conventional fixed speed generators (both synchronous and induction) must be run up to synchronous speed prior to energisation. On stall-regulated wind turbines with fixed pitch blades this is a potentially delicate operation, as there is no means of speed control while the turbine is offline. Under the action of the wind the free rotor accelerates rapidly and the generator contactor must be closed at just the right moment to prevent high torque transients.

The inrush current to the induction set can be largely suppressed using a ‘soft start’ module, a solid-state device that momentarily limits the generator excitation at the time of connection to the grid. The soft start module is switched out shortly after connection to avoid losses. Resistor banks can also be used in this way, and were a common soft start component in early stall regulated machines. Wind turbines with blade pitch control allow the rotor to be brought to synchronous speed offline, with the induction generator switched in only when the speed is correct and stable. This is a great advantage compared with the simple stall-regulated case, though soft start modules are still required to limit magnetising inrush.

5.5 Distribution networks

Distribution networks are used to supply electricity to individual large customers, or to population centres of small-medium size. Distribution is carried out at medium voltage (typically 6kV to 33kV depending on country), and most wind turbines feed in at this level. Wind turbines generate at low voltage, however, so their output requires transformation up to MV. There is normally one transformer per wind turbine, either external to the tower, within the tower base, or for very large machines in the nacelle.

5.6 Transmission networks

Transmission networks are at high (HV) or extreme-high (EHV) voltage, typically 132kV to 400kV in the UK, and are used to transmit large-scale power between regions. The output of a nuclear power station will be transformed up to HV for transmission to large population centres, possibly over hundreds of miles away. Wind turbines are not usually connected directly to the transmission network, although it may be appropriate for very large (e.g. offshore) wind farms.

The UK transmission system was essentially designed around centralised generation using large thermal power stations, and is not entirely suitable for distributed renewable energy production. In addition wind energy projects are often in areas distant from population centres, so future expansion of wind capacity may require transmission line upgrades.

5.7 Grid issues

The use of widespread, decentralised generation raises some important issues in regard to the stability and safe operation of a distribution network. These include:

- Voltage rise

  Generators cause local voltage rise on distribution networks which have high line or cable impedance; this limits the amount of capacity that can be installed on weak networks and may require the use of additional equipment such as automatic voltage regulators (AVRs).

- Flicker

  Flicker is the term given to voltage fluctuations on the distribution network caused by variations in the current supplied or consumed by an embedded generator. The inrush (magnetising) current due to connection of an induction generator or transformer can produce a significant voltage dip, and cause nuisance to domestic customers in the area. Flicker is less of an issue with more advanced generators (variable speed synchronous or DFIG). Utilities impose statutory limits on flicker, with the severity of allowable voltage dip dependent also on the frequency of occurrence.

- Harmonics

  Can potentially be caused by the high frequency switching operations of thyristor based soft start devices (which limit the inrush current to induction generators), or by the solid state power converters used by modern variable speed machines (both synchronous and DFIG). Typical switching frequencies are 2-6kHz, but the resulting harmonics may be filtered by suitable circuit design and are rarely a problem.

- Power factor

  Generally speaking utilities prefer to operate distribution networks at unity power factor, i.e. with minimal reactive power Q, so that phase angle $\phi = 0$ and $\cos\phi = 1$. High reactive power flows contribute to network losses, and high values of Q (in particular capacitive) can led to instability. Induction generators require power factor correction to compensate for their high inductive demand,and this is supplied either by switched capacitor banks or static VAr compensators (SVCs).

- Fault protection

  Safety is of paramount importance on electricity supply networks, which must remain safe under fault conditions such as short-circuits on transformers or generators, or failure of overhead lines, or underground cables. Under these circumstances embedded generators must disconnect immediately on detection of the fault.

6 Control and operation

This lecture examines aspects of wind turbine control including supervisory control (e.g. yaw alignment), continuous power control, and starting and stopping sequences. The differences between pitch and stall regulated rotors are explained, and the implications of fixed and variable speed operation. Other issues include regulation of electrical power factor and turbine noise.

6.1 Levels of control

The various levels of control on a wind turbine can be divided into the following categories:

- Supervisory control
- Power limiting
- Starting and stopping
- Power quality
- Noise limiting

Supervisory control refers to functions with relatively low time dependency, where the controller samples inputs such as wind speed, yaw offset (directional error), air temperature, etc, and takes decisions regarding the operation of the machine. For instance the wind speed will be averaged on a 20-30s basis, and if the value exceeds the shutdown speed the turbine will be taken offline. Likewise yaw error will be corrected on a fairly slow time base. The supervisory controller may also, however, respond to an instantaneous wind gust and shut the turbine down if necessary.

Power limiting is a real-time control function, where the controller aims to maintain the electrical power output at the rated value in high winds. This may be achieved through the use of blade pitch control and/or regulation of the generator torque. Power limiting is required continuously when the turbine is online, so high sampling rates and actuator control response are required.

Starting and stopping requires the controller to take the wind turbine from rest up to synchronous speed and connect the generator, or the reverse process, i.e. disconnect and bring the turbine to a stop. In doing so it is desirable to avoid high load transients at the instant of connection, or the possibility of rotor overspeed.

Fault and emergency control covers situations such as loss of the grid, or detection of out-oflimits conditions such as tower vibration, extreme generator temperature, or grid voltage imbalance. The wind turbine is then either brought to a normal, or an emergency stop, depending on the severity of the condition.

Power quality control relates to the instantaneous voltage and current at the point of connection to the grid. The most modern wind turbines types continuously monitor the grid and using solid-state technology can achieve real time control, e.g. maintaining constant or variable power factor, and limiting inrush current on synchronisation. Older, fixed-speed, types have more limited control of power quality.

Noise limiting is essentially a subset of rotor speed control, available only on variable speed machines. Rotor aerodynamic noise is sensitive to tip speed, which may be reduced in low winds to limit environmental impact. Not all wind turbines require all of these control functions. Small off-grid machines, for instance, may be installed without supervisory controllers and be able to operate in all wind conditions, with passive mechanisms for power limiting and yaw control. The larger and more sophisticated the wind turbine, however, the more control functions will be necessary.

6.2 Yaw control

Upwind HAWT rotors must be steered to face the wind, to prevent loss of energy capture or the introduction of cyclic aerodynamic loading. Yaw control is achieved using a relatively slow acting control loop that takes its input from a wind vane or ultrasonic anemometer on the nacelle, and periodically drives the nacelle in yaw to minimise the offset. The wind vane signal is averaged, after filtering to remove high frequency content. In addition it may require to be corrected for the directional offset imparted by wake rotation, so that the final signal fed to the yaw controller is an accurate estimate of yaw orientation. Yaw drives are usually small geared electric motors in the nacelle base, acting on a toothed ring on the tower. The nacelle turns on a large bearing, either of rolling-element type, or a plain annular ring on which the nacelle rests via friction pads that also act to damp the yaw motion.

Fast yaw response is generally undesirable, as it leads to gyroscopic loading, fatigue, and general wear. It is also unnecessary as the rotor can tolerate reasonable yaw misalignment without severe performance loss (energy capture falls away roughly as $\cos^2 \gamma$ where γ is yaw angle).

6.3 Aerodynamic torque control

Aerodynamic torque control is commonly achieved by varying the blade pitch, so that the entire blade is turned about its long axis via a hydraulic or electric actuator at the hub; the blade is mounted on a rolling element bearing at the root. On some early machines only the outer blade span was pitchable, with the actuator and bearing assembly mounted outboard in the blade. The effect of pitch control is to alter the C_P, λ characteristics of the rotor, with increased pitch angle effectively shifting the C_P curve to higher tip ratio λ. When the wind speed reaches the rated value blade pitch is progressively increased so as to keep the power output at the generator limit.

Pitch control effectively means continually shifting from one power curve to another to suit the changing wind conditions. On a fixed speed wind turbine the power control loop compares the actual and reference power output, and sends a demand to the pitch actuator based on the measured difference.

Generally the controller is based on PI (proportional + integral) response, though higher order controllers may be used to filter out harmonic terms in the power signal (e.g. due to tower shadow) that would cause unnecessary pitch activity. Blade pitch control can also be used via a speed control loop, when the generator is disconnected from the grid: this is the basis of the start-up and shutdown strategy for a pitch regulated wind turbine. A measurement of actual rotor speed is compared with a reference command signal, and the difference used as input to the pitch actuator. In this way the rotor can be accelerated smoothly up to synchronous speed, then held there stably while the generator is connected.

Partial span blade pitch achieves essentially the same function as full-span pitch, though is rarely used now due to the practical difficulties associated with locating the pitch actuator out along the blade. The method has, however, some aerodynamic advantages over full span pitch in that peak power overshoots are more limited as the fixed part of the blade behaves as a stall-regulated rotor. Power overshoots are, however, far less of an issue on variable speed wind turbines, and full span pitch control is now standard.

Note that some wind turbines operate by 'pitching into stall', i.e. reducing blade lift by increasing the angle of attack to beyond the stall point, rather than pitching to a lower angle. This method, also known as 'active stall', is used on turbines such as the Bonus 1.3MW.Active stall tends to operate on a slow control loop, with periodic change in pitch setting to suit different wind conditions. It is a simpler control method than positive blade pitch, though less sensitive, and subject to higher aerodynamic noise and lower aerodynamic damping. Positive pitch, with the blade in attached flow, tends to be used on variable speed machines.

Pitch control is also the primary means of braking the rotor, and providing a fail-safe emergency stopping capability. On loss of the grid, or on an emergency command signal, the blades will feather to near 90o pitch, in which condition they provide a high braking torque and bring the rotor to a standstill very quickly, generally within a few rotor revolutions. The advantage of pitch-based braking as opposed to the application of shaft brakes is that the drivetrain remains completely unloaded in the former case.

Though blade pitch is by far the most common aerodynamic control method, others used in the past have included spoilers: these are small movable strips embedded lengthwise in the suction surface of the blade, which can be raised into the airflow by mechanical actuators and thus 'spoil' the lift, hence controlling the power output. They are relatively crude control devices, lacking the strong braking capability of pitch control, and little used now.

Movable end-plates at the blade tips have also been used as control and/or braking devices. The AOC15/50 stall-regulated turbine has hinged aluminium tip brakes held in line (low drag configuration) by electromagnets. On loss of the grid the magnetic holding force disappears and the tip brakes open under centrifugal force, causing a high tangential braking load that quickly decelerates the rotor. On the 15/50 these fail-safe brakes are not, however, used for either power or speed control.

The Dutch Polenko 60kW stall-regulated wind turbine was equipped with similar tip brakes to the AOC machine, but in this case actuated via push rods linked through the hub to an actuator in the nacelle. Although their main function was air braking, they were also used to control rotor speed during the run up to synchronisation. Without speed control, most stall regulated turbines must connect the generator while accelerating through synchronous speed, which can give rise to severe torque transients if the slip value is high.

6.4 Generator torque control - fixed speed

When a fixed-speed wind turbine is online the generator torque always matches the rotor input torque: with T_G and T_R in balance there is no acceleration. This can be considered as a very basic, and in many ways effective, form of torque control. Also, referring to the previous section, the rotor of a fixed speed generator is always rotating at grid synchronous speed (or very near it in the case of the induction machine): as a consequence the rotor of the wind turbine is directly locked to the grid frequency, via the gearbox reduction ratio.

Fixed speed operation works best in conjunction with stall-regulated rotors, which have no active means of aerodynamic control, but the power automatically levels off in high winds due to rotor stall. This has been the preferred configuration for small (<200kW) wind turbines for many years. Some stall regulated machines incorporate 2-speed induction generators with the capability to switch from 6-pole to 4-pole operation in high winds, essentially a crude (though effective) form of speed control.

Fixed speed generators were used for the first generation of pitch controlled wind turbines, with considerable success. Although some had synchronous generators (e.g. WEG LS-1, Howden HWP45) the induction set became more or less standard due to its superior dynamic characteristics. Examples of fixed-speed wind turbines with pitch control include the Vestas V27, and V47 (Denmark), HMZ Windmaster 300kW (Belgium), and WEG MS-3 (UK).

One of the disadvantages of the fixed-speed variable pitch configuration, however, is the difficulty in controlling the power accurately in high winds: large power excursions were common on many types. This is due to the sensitivity of the rotor aerodynamic torque to sudden changes in wind speed, and the inability of pitch mechanisms to responds fast enough to gusts. Large torque transients at the rotor are then mirrored by the generator reaction torque, causing high load extremes on the drivetrain components.

In some cases the dynamic response of the wind turbine drivetrain compounded the problem, amplifying the input torque loading, and causing the blade pitch system to over-respond. A great deal of R & D went into improvements to controller algorithms and pitch control mechanisms1, but without ever completely eliminating this issue. Some turbines, such as the Vestas V27, were less affected due to a combination of robust design, and modest rotor size in relation to the generator rating.

6.5 Generator torque control - variable speed

The ultimate solution to limiting drivetrain torque transients was the introduction of variable speed, achieved by active torque control at the generator. Referring again to the drivetrain equation of motion, by holding generator torque TG constant, variation in rotor input torque TR is translated into a change in rotor speed with excess input energy stored as angular momentum. This has several great advantages, namely:

- Drivetrain (generator and gearbox) loading is strictly capped without overshoots
- The mismatch between input and output torque is 'flywheeled' at the rotor
- Rotor pitch can be regulated on a relatively slow speed control loop

These advantages were first demonstrated on otherwise conventional turbines such as the WEG LS-1, which had a variable-speed 'reaction drive' gearbox that decoupled the rotor speed from the rest of the drivetrain, which included a conventional fixed-speed synchronous generator. The output power of the wind turbine was then relatively smooth under power limiting conditions (compared with the Howden HWP45, for instance, also with synchronous generator but lacking variable speed capability).

Mechanical means of variable speed operation were pursued for a time, including torque limiting gearboxes, but were superceded by developments in variable speed power electronics, which introduced the capability for broad-range variable speed generators. These brought all the advantages of variable speed operation with the mechanical simplicity of the conventional wind turbine drivetrain. The first wind turbines with variable speed generators retained a gearbox, i.e. with the generator running at high speed. Examples included the Enercon E33-300 and the Tacke (later GE) 1.5MW machine, both made in Germany in the early-mid 1990s. More recently low-speed direct drive synchronous generators have been introduced, most notably by Enercon with the E48 and E70 series of machines.

Although these direct drive generators are formally synchronous machines, they can operate at broad range variable speed due to the action of the solid state AC-DC-AC power converter, which effectively severs the link between grid and generator frequency. The generator operates at a locally synthesised AC frequency, with output power first rectified to DC then inverted to grid-compliant AC. The doubly fed-induction generator (DFIG) has very similar characteristics to the synchronous variable-speed machine in terms of operation, but differs in the detail of the power conversion system (see previous lecture). Modern variable-speed turbines with DFIGs include the Vestas V52, V80 and V90, RePower MM70, and Nordex N80.

6.6 Power control loops

Power control on pitch-regulated machines is achieved by closed loop control. There are, however, subtle differences between the control regimes for fixed and variable speed wind turbines. In the former case there is no active control of generator torque TG, and the purpose of the pitch controller is to limit the rotor torque TR to the rated level in high winds. As noted above this is achieved using a power control loop, comparing actual power with the reference level, and sending a demand to the pitch controller based on the difference.

On the variable speed machine, once power limiting conditions are reached the generator torque T_G is controlled by a fast acting electronic control loop, which limits the output power to the required level without any mechanical actuation. The generator torque controller acts more or less instantaneously. Any mismatch between aerodynamic torque TR and generator torque T_G then leads to acceleration of the rotor ($d\Omega/dt \neq 0$).

In rated power conditions the pitch controller acts under a speed, rather than power, control loop. Pitch response is not sufficiently fast to match the variations in a turbulent wind, but the error is simply translated into variation in rotor speed: the rotor acts as a flywheel, storing and releasing energy.

Power control on the variable speed machine is thus more complex, but more effective, than on the fixed speed wind turbine. In low winds (below rated power) the pitch angle is held nearly constant, with the rotor speed varied to maintain optimum CP. In high winds the pitch angle is progressively increased to limit the power to Prated, while the rotor speed is held approximately constant: the speed is, however, allowed to vary on a short-term basis to smooth out transient input loads.

The pitch mechanisms for fixed and variable speed wind turbines are similar. With Vestas machines, hydraulic actuators are used to pitch the blades, and both the fixed speed V27 and variable speed V52 have a single hydraulic ram in the nacelle, acting through the hub. Enercon wind turbines use electric pitch drives, with one hub-mounted pitch motor per blade.

6.7 Starting and stopping

In addition to power regulation, the wind turbine control system must also manage the startup and shutdown operations, including connection and disconnection of the generator. In the general case, during the start-up sequence when the rotor accelerates from a parked stationary state to synchronous speed, the generator is not connected and applies no torque T_G to the drivetrain. Thus the behaviour of the system is dictated by the rotor aerodynamic torque T_R acting on the rotational inertia I.

6.7.1 Fixed pitch stall-regulated turbines

In the case of a stall regulated wind turbine the rotor torque TR depends on the rotation speed and the wind speed, with no means of active control during start-up. The rotor simply accelerates freely under the action of the wind, and the controller closes the generator contactor when synchronous speed is achieved. This must be done accurately, or the rotor will accelerate rapidly through synchronous speed to a potentially high runaway speed.

At the instant of connection there is inevitably some mismatch between the actual rotor speed and synchronous speed, and this is minimised through use of an electronic 'soft-start' mechanism. In the unlikely event that a stall-regulated turbine failed to synchronise, the controller would detect an overspeed condition and deploy the blade tip air brakes, to bring the turbine back to rest, or at least a safe speed. The same method is used when disconnecting the turbine from the grid in extreme winds, or under any fault condition.

The fixed-pitch rotor of a stall-regulated machine must be capable of developing enough torque to start up in the first instance, when the stationary blades are almost flat on to the wind (pitch angle $\approx$ zero) in a low-lift configuration. This is achieved by incorporating sufficient twist in the blade to generate a modest starting torque, enough to start the rotor gently rotating in a light wind (formally speaking the rotor aerodynamic torque coefficient C_Q requires to have a positive value at zero tip speed ratio λ).

As the rotor speeds up the flow angle ϕ on the blades decreases and the lift increases, until at a certain speed the rotor becomes unstalled and the acceleration picks up dramatically: this explains the need for accurate timing of the generator contactor, to prevent overspeed. Once the generator is synchronised, however, constant speed operation is automatically maintained.

The stopping sequence for a stall-regulated machine generally involves simultaneously opening the generator contactor, and releasing the blade tip air brakes. The rotor decelerates quickly under the action of the air brakes, until a low speed is achieved, when a shaft brake may be applied. Shaft brakes rarely applied at high rotation speeds, to prevent excessive wear. Some smaller wind turbines such as the AOC15/50 (see below) may incorporate a dump load in the electrical circuit: this is a resistor bank that can absorb the output of the generator during the braking cycle and effect electromagnetic braking (capacitive excitation of the generator is needed under these circumstances).

6.7.2 Variable pitch wind turbines

Variable pitch introduces the capability to control the rotor speed when the machine is offline; this is a powerful advantage in regard to starting and stopping. The start-up sequence begins with the release of the shaft brake, with the rotor pitched to an advantageous angle to generate starting torque. As the rotor accelerates the pitch controller operates on a speed-control loop to bring the rotor to synchronous speed smoothly, and without the high acceleration seen in the stall-regulated case.

The pitch controller holds the rotor near synchronous speed until satisfactory values of speed and acceleration are achieved, when the contactor is closed. From that point, the generator automatically maintains synchronous speed. Note that the blade pitch angle at the point of synchronisation is such as to achieve zero acceleration prior to synchronisation, hence T_G is zero and T_R is just enough to overcome rotational friction in the drivetrain. Once the turbine synchronises, the pitch angle is brought back to the 'run' position to generate positive aerodynamic torque according to the wind conditions.

The stopping sequence is essentially the reverse of the above. The blade pitch controller first reduces the rotor power by pitching to achieve zero kW output with the generator still online; the contactor is then opened, and the pitch controller feathers the rotor blades to bring the rotor to a standstill. Starting and stopping is far more controlled on a pitch regulated wind turbine than on a stall-regulated machine, and fatigue loading is consequently lower.

6.8 Mechanical brakes

Most large HAWTs are equipped with shaft brakes, usually on the high-speed shaft out of the gearbox. Conventional disk-type brakes with calliper mounted brake pads are the most common: these are designed to be fail safe, with the brake callipers held off by hydraulic force, such that a loss of power results in rapid deployment of the brake. Some design codes require a wind turbine to have two fully independent failsafe braking mechanisms, of which the shaft brake is one, and the blade pitch system the other. Shaft brakes are sometimes mounted on the low speed shaft, in which case a far higher torque rating is required to supply the same braking effect, and a large disk with many callipers may be required. The advantages of the low speed over the high speed brake are (a) there is negligible heating in the former case, and (b) the low speed brake acts directly on the rotor without imposing any torque on the gearbox.

6.9 Power conditioning

As noted in the previous section, power quality at the point of grid connection is defined by several electrical parameters including power factor and voltage stability. The wind turbine controller has some control over these factors, with different levels of sophistication achieved depending on the type of generator.

Fixed-speed wind turbines with induction generators have limited control over the terminal voltage and are generally equipped with fairly crude capacitor-based power factor correction systems. These supply fixed capacitive reactance to offset the inductive reactance of the generator. In some cases (e.g. the V27) the capacitors are provided in banks, so that the power factor correction can be roughly matched to the operating conditions.

The high inrush current on connection of the induction set is suppressed using a thyristor based soft start module, although the peak current may still be higher than the rated value. Utilities in the UK are increasingly keen to see more sophisticated power conditioning on wind turbines, and traditional fixed-speed induction generators are becoming less common.

Variable speed generators incorporating solid-state power converters enable much better control of power factor and voltage than traditional induction machines. The power converter allows precise control of reactive current, enabling positive or negative power factor to be set depending on the network requirements; this is, furthermore, achieved without the need for capacitors or inductors.

A common control algorithm on such machines is 'voltage dependent power factor control', whereby reactive current is drawn from the grid so as to limit voltage rise due to the export of real power. The modern DFIG or synchronous machine can do this on a dynamic basis, with the power factor set point periodically changed according to the local steady-state voltage. Enercon machines are also capable of reducing their peak power output in order to limit the local voltage rise, though this is a less common requirement.

6.10 Sector management

A further type of control available on modern wind turbines is sector management, a supervisory control function whereby the wind turbine is operated at a different power level, or stopped altogether, depending on the wind direction. Sector management is used in the Vestas Sound Reduction System (SRS) in order to protect nearby properties from undue noise when they are downwind of the turbine: the SRS selects a noise-reduced power curve for given pre-programmed wind directions. The energy yield penalty is usually small.

Sector management may also be used to curtail wind turbine operation in turbulent wind conditions, where the turbulence is due to local topography. Again, the supervisory controller will be programmed with the relevant wind directions in a lookup table, and the turbine may be operated at reduced shut down wind speed, or stopped altogether, with the wind from an unfavourable direction.

7 Materials and manufacture

This section gives a broad overview of the materials that comprise the primary structure of a wind turbine, including steel (various grades), SG cast iron, advanced composites including glass and carbon fibre-reinforced plastics, and wood-epoxy laminate. Some fundamental properties of composites are explained, including fibre volume fraction, material safety factors, and allowable fatigue strength. A comparison of specific (per unit weight) material properties is used to illustrate why composites are superior to steel for blade construction.

Blade manufacture is described for both wood-epoxy laminate technology, and resin infusion moulding (RIM), in the latter case with glass-reinforced epoxy composite. The advantages of the two technologies are compared. Wood-epoxy construction is illustrated with reference to hand lamination of wood veneers, and vacuum forming in the mould. RIM manufacture is illustrated with respect to a typical GRE blade made in a temperature-controlled mould. Different blade root attachment methods are described, including bonded and keyed studs, and T-bolts. The advantage of preloading the T-bolt to reduce fatigue loading is described.

7.1 Materials used in turbine construction

The principal materials used in construction of a large wind turbine are as follows:

Steel (various grades): this comprises the main material (by weight) in the structure. Tower sections and foundation inserts are manufactured from steel plate, rolled and welded by automated machinery. Typical yield strength is 350-400MPa. Towers are generally made in several parts, depending on the final height. Each section has rolled and machined bolt flanges at the ends, made to higher dimensional tolerance than the main tubes.

Spheroidal graphite (SG) cast iron: SG iron combines relatively high strength (UTS in the range 370-726MPa) with toughness and good wear resistance, with the ability to cast relatively large and complex components. Its favourable properties are due to microscopic graphite nodules embedded in a ferrite matrix. SG iron is used to cast load-bearing components such as the rotor hub, nacelle bedplate, or gearbox casing. Finished components can be machined or surface ground to high tolerance (e.g. for taper fit shafts or bolt flanges).

Glass-fibre reinforced plastic (GFRP): GFRP is widely used in wind turbines, both in general purpose applications with light loading, such as the nacelle cover and rotor hub spinner, and for the rotor blades, which are required to carry very large and fluctuating loads. Nacelle covers and spinners are generally made from thin laminates with multi-axial fabrics (or chopped stand mat) in epoxy or polyester resin. Rotor blades are nowadays usually made in glass-epoxy, with high performance fibre-oriented fabrics.

Carbon fibre reinforced plastic (CFRP): carbon fibre composites are also used to some extent, though not as widely as GFRP due to their much higher cost. CFRP tends to be used in key areas of the blade composite structure (e.g. at the blade root) to achieve very high strength and/or stiffness.

Wood epoxy laminate: wood-epoxy has been used for some years as the primary structural material for large rotor blades due to its favourable specific properties (i.e. strength/weight and stiffness/weight), and because wood (from managed plantations) is a sustainable raw material. The preferred species are birch, poplar, and Douglas fir; previously mahogany (khaya) was used, but fell out of fashion due to issues of sustainability. In conventional wood-epoxy blades the veneers (3-4mm) are coated with epoxy resin and vacuum formed in moulds. The resin requirement is much lower than with GFRP/CFRP, as the epoxy forms a glue line between the wood laminates, rather than a matrix as in GFRP. Larger blades (40-50m) are now being made using a wood and CFRP composite, with vacuum infusion moulding techniques.

7.2 Properties of fibre composites

GFRP and CFRP are composites of glass or carbon fibre in a thermosetting resin matrix. The properties of the resulting composite are anisotropic, i.e. strongly dependent on the fibre orientation, with maximum strength and stiffness developed in the fibre direction, and very little at right angles to it. Composite properties can therefore be 'tuned' by combining fibres of different orientation in the mix. The volume fraction V_f for component of volume V_i in a composite of total volume V_{total} is defined by:

$$V_f = \frac{V_i}{V_{total}} \tag{7.1}$$

The elastic modulus E_c of a fibre-based composite may be approximated by the rule of mixtures, or Voigt estimate:

$$E_c = E_m V_m + E_f V_f \tag{7.2}$$

where E_m and E_f are the respective moduli, and V_m and V_f the volume fractions, for the resin matrix and the fibres. From this it is seen that the composite properties will be dominated by the stiffer component, i.e. the fibres (for example the tensile modulus E_f for glass fibres is $\sim$ 70GPa, with E_m for the laminating resin typically 3GPa, hence $E_f >> E_m$). The rule of mixtures can also be used to calculate the modulus of a composite containing a mix of different fibre types, and/or fibres with different orientations. The volume fraction V_f of glass in GFRP depends on the method of manufacture. Traditional hand layup may achieve V_f of 30-40%; more modern vacuum infusion technologies enable values of over 50% to be achieved, with almost no air trapped in the laminate.

The resin matrix used for GFRP is a thermosetting plastic, most commonly polyester, epoxy, or vinyl acetate. Polyester and vinyl acetate require solvent sand are therefore less environmentally friendly in production (fume extraction is required). Two-part epoxies, without solvent, are generally preferred for infusion moulding, requiring no solvent, and also giving lower shrinkage in the final product.

7.3 Material safety factors

The allowable strength for a composite is the value used in the design process: it is based on experimental test results (sometimes known as characteristic strength values) modified by a material safety factor γ. This factor itself comprises a number of partial safety factors; for example the Germanischer Lloyd recommendations define the material safety factor γ by:

$$\gamma = C_1 \times C_2 \times C_3 \times C_4 \times C_5 \tag{7.3}$$

where C_{1-5} are partial safety factors as follows:

C1 : General safety factor (1.35)

C2 : Creep strength factor (1.5)

C3 : Temperature effect factor (1.1)

C4 : Production factor (1.2 for hand layup, 1.1 for infusion)

C5 : heat treatment factor (1.0 controlled cure, 1.1 uncontrolled)

Hence, for a high volume fraction glass-epoxy laminate produced by resin infusion, $\gamma = 2.45$.

7.4 Fatigue strength

The fatigue strength of composites and steels may be defined via a further partial safety factor that takes account of the number of load cycles and the S/N curve of the material in question. The S/N curve is of the following general form:

$$\log S_n = \log S_0 - (1/x) \log N \tag{7.4}$$

where S_n is the fatigue strength at N cycles, S_0 is the allowable static strength, and x is an empirical factor dependent on the material in question. The above expression can be used to derive a strength reduction factor, which Germanischer Lloyd define as partial safety factor C_{2b} such that:

$$C_{2b} = N^{1/x} \tag{7.5}$$

The fatigue properties of the material are then defined by the exponent x, which determines the slope of the S/N curve. The higher the value of x, the shallower the curve, and hence the more fatigue-resistant the material. Typical values include Steel (3-5), GFRP (10) and CFRP (40). From this it may be seen that steel is relatively poor in fatigue compared with either GFRP, while carbon fibre composites (CFRP) have extremely high fatigue resistance. This largely explains the use of composites in blade construction, where fatigue loading is a major driver.

7.5 Comparison of material properties

The key material properties are:

- E-modulus
- Tensile and compressive strength
- Shear strength
- Fatigue strength
- Environmental resistance (corrosion, moisture ingress)
- Ease of manufacture
- Cost

Properties for the main materials under consideration are compared below in Tables 1-3. Table 1 gives tensile strength and E-modulus. Table 2 shows specific properties, i.e. per unit weight, in this case normalised with respect to structural steel (350MPa allowable strength). Table 3 compares the fatigue strength of the different materials assuming a cycle count of 107, with reduction of strength dependent on the S/N curve exponent x (see above). Referring to the data in the tables, the following points are noted:

7.5.1 GFRP

Unidirectional glass-reinforced epoxy laminate (GRE) has a superior strength/weight ratio to structural steel by a factor of almost 3. Its stiffness/weight ratio is lower than steel, though by a smaller margin than the above. The fatigue strength of GRE is, however, markedly superior to steel (by a factor of 5.8). These favourable properties, combined with the great manufacturing versatility of GRE and its reasonable cost (albeit higher than steel), make it a natural choice for wind turbine blades.

7.5.2 Wood-epoxy laminate

The specific strength of wood epoxy is more than twice that of steel, while its specific stiffness is only 4% lower than steel. Specific fatigue strength is similar to that of GRE, at 6 times that of steel. With reasonable costs, wood epoxy is therefore a very good alternative to GRE for the primary structure of wind turbine blades. Wood laminate is arguably more difficult to manufacture than GRE, but has some advantages, namely:

- The low density of wood results in inherently thick blade skins with high buckling resistance: GFRP and CFRP skins require sandwich panel construction to resist buckling.
- Wood veneers require only a thin layer of epoxy to bond them, essentially a glue line. By comparison around half the volume of a GFRP composite is epoxy resin. This makes wood laminate a potentially cheaper, and more environmentally sustainable, technology.

7.5.3 CFRP

Carbon-fibre reinforced plastics are the most technically advanced composites in wind turbine use, with specific strength an order of magnitude higher than steel, and specific stiffness around 3 times greater. Their fatigue performance is likewise far superior to steel. The current drawback with CFRP is its very high cost, which makes it less attractive than glass or wood composites for blade manufacture. CFRP is, however, used in combination with these materials where higher stiffness or strength is required, for instance:

- As an interlaminate layer in the root shell of large GFRP blades, where locally high stresses may occur, e.g. the compressive stress in the laminate around the T-bolt fixture.
- As a stiffening element in large wood laminate blades, where it may be advantageous to concentrate high-modulus material near the skin: although wood laminate blade skins are naturally thick (see above), this reduces their material efficiency in regard to bending stiffness.

7.6 Comments on the use of steel

As noted above, advanced composites are now universally used in the manufacture of large rotor blades. Some early wind turbines had steel blade spars, or in some cases steel blade skins (with buckling resistance provided by PU plastic foam) but high weight, and fatigue, were almost always attendant. In addition, large steel components require welded construction, which makes this material far less attractive in manufacture than composites, which can effectively be moulded as a single piece.

It is, however, clear that steel is still the material of choice for most wind turbine towers. The main reasons for this are that weight is much less of an issue for the tower than for the blades, and steel is still significantly cheaper than any of the composites. Note that self-weight does not contribute significantly to the main static or fatigue loads on the tower, which are mainly due to the rotor thrust loads. By contrast, the self-weight of a blade is a significant bending load, and because of the blades rotation it is also a major source of fatigue loading.

7.7 Production techniques

7.7.1 Wood-epoxy laminate

Wood-epoxy rotor blade technology originated in the US, where it was evolved from commercial boat building technology, and first applied to wind turbine blades. This process involves laminating wood veneers under vacuum in a 2-part mould using cold-curing epoxy resins. The veneers are typically 3-4mm thick, produced by rotary peeling of cut logs. The wood requires low moisture content ($< 10\%$) to ensure high strength, and the veneers are laid with grain parallel to the blade long axis for maximum bending strength and stiffness. Once the blade halves have cured in the moulds the veneers are trimmed to size using a router or circular saw. Internal components such as the shear web are bonded into one blade shell at this stage, prior to joining the two halves. The laminate design leaves a wide glue-line at the blade leading and trailing edges, and the two blade shells are bonded together using thickened epoxy paste applied along these areas. The two mould halves are clamped together until the epoxy paste cures sufficiently for blade release.

Wood laminate blades were traditionally surfaces with polyester gel coat: this was applied in the mould prior to the veneer layup, so that the blade was ultimately demoulded with a high gloss finish. More modern wood laminate blades produced by infusion techniques are not gelcoated, and are demoulded prior to sanding and painting, in the same way as GFRP blades.

7.7.2 Glass-epoxy blades (RIM production)

Large glass-epoxy blades are nowadays manufactured by vacuum resin infusion moulding (RIM) rather than the wet layup techniques previously used in GRP shops. With RIM the component materials are laid dry into the mould, starting with the outer blade skin, and then adding successive layers of reinforcing fabric and foam sandwich panel material. The spar cap is laid down using unidirectional (UD) glass fabric oriented along the blade span. The thickness of the laminate layers decreases from blade root to tip, in accordance with the decrease in chord length, and the spanwise load distribution. The root laminate thickness is also dictated by the need to accommodate the steel root studs. Depending on the design these may either be encapsulated in dry cloth and bonded in place during the moulding process, or located in holes bored in the finished blade.

The dry laminate layup is covered by a transport mesh, to aid the flow of resin under vacuum, overlaid with a resin flow channel. Finally a vacuum bag is laid over the complete assemblage, and sealed around the mould edges. When vacuum is drawn the bag compresses the laminate in the mould under 1bar pressure. The vacuum also draws the resin into the mould, where it flows into the dry laminate. The infusion process takes 1-2h before the laminate is fully wet out. The mould halves are heated to speed the resin flow (via lower viscosity) and accelerate the epoxy cure. Electric panel heating is typically used, with 20-30 heating elements per mould, each under closed loop control. The controller takes account of the resin exotherm, which can otherwise lead to elevated temperatures; some moulds incorporate water cooling for this reason.

Once the resin has cured the two blade moulds are stripped of vacuum bags and other disposables, and internal components are fitted to the blade: as well as the shear web these may include lightning conductor, tip brake components, and various prefabricated joining pieces to facilitate the bonding of the two halves. The half-moulds are then brought together and the blade joined in much the same manner as the wood epoxy blade (see above). Thickened epoxy paste is again used for bonding the blade halves.

Blades up to 50m long are now manufactured by the RIM process, with a 24h mould cycle possible in principle. Once the blade is demoulded it may be post-cured in an oven for maximum strength: this involves soaking at 70oC for 6-7 hours. Some areas of the blade may have external laminate added first, e.g. at the trailing edge and/or blade root regions. The finished blades are surface dressed, smoothed, and sanded before painting.

7.7.3 T-Bolts

A common root attachment method for GFRP blades is the T-bolt (often known as the 'IKEA' fitting for its similarity to a furniture joint). With this fixture the circular blade root shell is drilled with intersecting radial and circumferential holes. The T-bolt is a threaded stud located in the circumferential hole, and held by a captive circular nut in the radial hole. In service the T-bolt is heavily pre-tensioned, with the blade laminate in compression resisting the preload. The applied load is then shared according to the relative stiffness of the bolts and the laminate, with the latter dominating due to its greater cross sectional area. As a result the fatigue loading on the joint is carried mostly by the laminate, with the bolt remaining under almost constant tensile loading, and hence highly resistant to fatigue.

8 Planning and siting

This section discusses the steps leading up to windfarm construction, and in particular selection of appropriate sites, and the subsequent planning process. The main attributes of an attractive wind site can be summarised as:

- Good wind resource
- Suitable terrain and infrastructure
- Available grid export capacity
- High likelihood of achieving planning consent

The main content of the section addresses the last point, i.e. planning consent, and the many issues that may need to be addressed within the planning process. Firstly, however, some brief notes follow regarding site selection.

8.1 Wind resource

In some countries the regional wind resource is already mapped at a high level, giving developers a useful guide to suitable project areas. In the UK the BERR (DTI) database estimates the mean wind speed at heights of 10m, 25m, and 45m above ground level with horizontal resolution of 1km. The data are based on long-term records from weather stations across the UK with interpolation for locations in between, and correction for terrain height; the method is described by Burch.

The BERR database should be used with caution, as the estimated wind speeds relate to the highest elevation per 1km square: in areas of complex or undulating terrain there may be a very wide variation in speeds, and the turbine site may not be at the highest point. In addition the accuracy of the wind speeds is of the order $\pm 10\%$, and should only be treated as a rough guide to the local resource. With these caveats, however, the BERR database is quite useful.

For economic purposes the wind resource needs to be known more accurately, which usually involves measurement, preferably at hub height. Long-term estimates of wind speed are made using the Measure Correlate Predict (MCP) technique, where a correlation is established between simultaneous wind readings at the site and at a reference weather station for which long-term records are available.

Where windfarms of several machines are planned, micrositing analysis is required in order to estimate the conditions at each location, as it is rarely economic to erect a separate met mast per position (though large developments may well use several full-height masts to reduce uncertainty). There are now commercial software packages for site analysis: these use measured wind data to seed the analysis, and incorporate wind flow models in conjunction with digital terrain data, and calculations to account for array wake loss.

The most common analysis packages used in the UK are RESoft Windfarm, and Garrad Hassan Windfarmer. The latter is used in conjunction with the Danish WAsP model for terrain flow modelling, while the RESoft code incorporates inbuilt routines for this purpose. Both codes have been developed to model other aspects of wind farm siting such as noise and visual impact (see below). Elsewhere the Danish package WindPro is widely used.

The minimum economic wind speed for a site will depend on overall project cost: with low infrastructure costs sites can be viable at lower wind speeds. In general, however, it is usual to develop projects at hub-height wind speeds above 6m/s. The wind assessment exercise is also used to verify that the wind characteristics (mean speed, turbulence and shear) are within the design limits for commercial wind turbines. The IEC site classification is widely used in this regard.

8.2 Terrain and infrastructure

A high wind site may be of little interest for wind development if the terrain is unsuitable for delivery or construction of wind turbines. This is often the case in remote or mountainous areas. Among the important criteria for project development are therefore:

Terrain: the best wind sites may be on elevated ground, but smooth rounded hills are more suitable than steep sided ones; developing in mountainous or complex terrain can lead to high infrastructure costs, and these sites may also be subject to turbulent or extreme winds. Local site gradients can also be a limitation in regard to wind flow and transport access.

Proximity to trunk roads: infrastructure costs are kept down if the development is within reasonable proximity to a major road network that does not require upgrading to accommodate delivery and construction vehicles, including concrete lorries, cranes, and the turbine components. Transport surveys are usually carried out as part of the wind farm development process.

Access to ports and ferries: the load capacity of scheduled ferries may be a constraint, where projects are not large enough to justify the cost of chartered vessels. Likewise the vehicle length limitations of delivery ports can be critical. These are usually checked as part of the transport route survey (see above).

Land ownership: legal and commercial agreements need to be in place with the owners of the land on which the development is to take place. This is more of an issue for large wind farms where the landowner is often not the owner of the project, but instead leases the site. In addition several landowners may be involved, depending on the size of the windfarm.

8.3 Electrical grid capacity

The available grid capacity in an area is dependent on the following factors:

- The strength of the local distribution and/or transmission network, in terms of (a) the power-carrying capacity of the lines under worst-case conditions, generally in the summer when their resistance is highest, and (b) the level of voltage rise that can be tolerated.
- Proximity to major load centres, e.g. towns or cities. This is really a re-statement of the above, as grid networks tend to be stronger and more widely developed in the vicinity of large population centres. Otherwise, major grid reinforcement is needed to carry power from generation sites to load centres.
- Other generation in the area. Where there are many projects proposed, but only limited grid capacity, some developments will either be prevented or may have to operate on a 'constrained' basis.

The cost of grid reinforcement may be significant and can dictate whether some projects are viable. For instance, here in the UK a number of very large wind projects - some with planning consent - are awaiting transmission line upgrades that may take several years to complete (e.g. Western Isles interconnector, Beauly Denny transmission upgrade).

8.4 Environmental constraints

There may be a number of environmental constraints facing a wind turbine development. While these may differ in detail from one country to another, the following list is fairly typical:

8.4.1 Ecology - flora & fauna

Developments are generally not permitted where destruction of habitat will occur, especially where protected species are involved. The pre-development process usually includes a survey of local flora and fauna, and assessment of likely impact. Mitigation measures may be accepted as part of the development.

8.4.2 Birds

The above comments apply particularly to birds, and a great deal of research has been conducted into the impact of wind turbine developments on bird life. The two main concerns are (a) direct mortality due to collision with rotor blades, and (b) loss of habitat due to construction. Of these (b) is probably now regarded as the greater threat, though there have been specific instances where inappropriately sited wind arms have caused serious bird mortality, e.g. eagle mortality at some of the early Altamont windfarm developments.

In general, however, birds and wind turbines can safely co-exist so long as sites are appropriately chosen and/or suitable mitigation measures applied: in some cases these may include preservation of alternative habitat. Pre and post-construction surveys are usually mandatory in areas of important bird activity; these surveys may include:

- Compiling a species list (all birds) for the immediate area
- A survey of breeding bird numbers (key species) throughout a season
- Recording flightlines, often between roosting and feeding areas
- Recording bird mortality (post construction), if any

A good example of mitigation is Beinn an Tuirc wind farm, in Kintyre. The site forms part of the territory for a pair of golden eagles, a species protected under UK and EU law. When the wind farm was being developed it was found that the eagles occupied marginal territory where food resources were scarce. Low breeding rates were believed to be due to the increase in forestry in the area, which had displaced the birds' natural prey.

To mitigate against the possibility of an eagle collision and improve the habitat for the birds, the developer undertook a habit management plan that increased the availability of prey within the eagles' territory. This involved large-scale removal of immature forestry, together with management of the heather moorland and the creation of prey 'hotspots', suitably separated from the wind farm.

8.4.3 Visual impact

Visual impact of wind turbines, particularly wind farms, is a particularly subjective issue and can lead to strongly polarised views. Independent research in the UK tends to indicate broad support for wind turbines in the landscape7, but individual developments can nevertheless run into problems so sensitive design, together with early engagement with the public, are always recommended.

The ownership of wind farms also plays a large part in how they are perceived visually. A development owned by a farmer or community may be regarded as an attractive addition to the working landscape, if the local population are seen to gain from its presence. A development that brings little income to an area may be less popular. Visual impact is probably the most important public issue regarding wind turbines, and often lies at the root of more wide-ranging objections. In the UK the planning authorities place great stress on visual impact. Independent guidelines for landscape and visual impact assessment are incorporated in the planning process, including the cumulative impact of multiple developments within a given area8. The procedure for a full landscape and visual impact assessment (LVIA) generally includes the following steps:

- Establishment of zones of theoretical visibility (ZTVs) using digital map data
- Identification of important viewpoints from which the development may be seen
- Preparation of wire frame computer images to indicate visibility and scale
- Preparation of photomontages to give representative images of the final development

Preparation of wire frame views and photomontages for particular viewpoints is carried out using standard procedures to ensure that the resulting images neither over- nor underestimate the visual impact of the development. Standard camera focal length and viewing angles are therefore recommended, with focal length of 50mm for a 35mm format camera, in conjunction with 90o field of view.

The LVIA must also consider the quality of the landscape, and its capacity to accept wind energy developments. Clearly some landscapes are more suited than others, and certain areas (such as National Parks) may simply be off limits to wind energy developments, although the scale of the proposal should always be taken into account. Visual impact is a large subject in itself, suffice to say that it requires very careful consideration for most developments.

8.4.4 Noise

Modern wind turbines are relatively quiet due to advances in the aerodynamic design of blades (particularly tip shape), improved noise isolation of the gearbox and generator, and the increasing use of variable speed operation. Nonetheless, developers must always consider potential for noise impact on the public.

Approximate estimates of wind turbine noise as a function of observer distance may be made using the National Physical Laboratory model10, which assumes hemispherical spreading of sound from a point source at the wind turbine hub, with an allowance for atmospheric absorption. Thus, for a wind turbine with sound power level LW the perceived sound pressure level is given by:

$$L_P = L_W - 10\log_{10}(2R^2) - \alpha R \tag{8.1}$$

where: L_P is the perceived sound pressure level in dB(A), L_W is the sound power level of source in dB(A), R is the observer distance from source in m and α is the atmospheric absorption coefficient, dB/m.

Typical values of L_W for large wind turbines are 90-105dB(A), depending on wind speed and turbine size; certified values of source noise LW can be obtained from the wind turbine manufacturer. The atmospheric absorption coefficient α may range from 0.002-0.005dB/m, again depending on rotor size. Larger wind turbines give rise to lower frequency noise, which tends to propagate further, hence lower values of α.

Wind turbine noise tends to increase with wind speed due to higher loading, and in some cases higher tip speed. Background noise also increases with wind speed, however, and the most sensitive operating conditions tend to be in moderate winds with the rotor operating at medium power: in high winds the rotor noise may be masked by background noise.

Allowable noise levels vary from one country to another, but the following guidelines, as used by most UK planning authorities (based on recommendations previously made by the DTI) are fairly typical:

- Noise from a wind farm should be limited to 5dB(A) above background for both day- and night-time, with wind farm noise and background both measured as LA90,10min.
- In low noise environments the daytime level of the wind farm noise should be limited to an absolute level within the range of 35-40dB(A). This may be increased to 45dB(A) in cases where the occupier of the affected property has a financial interest in the wind farm.
- For single turbines or wind farms well separated from the nearest properties, a simplified condition applies with maximum allowable 35dB(A) in wind speeds up to 10m/s.

The noise due to multiple wind turbines is calculated by summing individual noise contributions at a specific location, following the appropriate mathematical procedures. This is performed automatically by commercial software packages such as those referred to above, whose typical output includes noise contours in the vicinity of the planned windfarm. Optimisation routines may also be included, enabling detailed wind turbine micrositing to be cried out to reduce noise at a particular location without significant loss of energy capture.

8.4.5 Shadow flicker

Shadow flicker may occur at certain times of the year when the sun passes behind the rotor of a wind turbine and cast an intermittent shadow over neighbouring properties. The likelihood of shadow flicker is also dependent on geographical latitude, time of day, and the degree of cloud cover. The duration of the effect at a given property can be calculated from the geometry and latitude of the potential site, with the effect quantified in terms of annual hours of flicker nuisance. Graphical methods are often used to illustrate the results.

In general shadow flicker is unlikely to be a problem where the separation between rotor blades and the nearest properties is greater than 10 rotor diameters. Where closer distances are specified, detailed flicker duration calculations may be required. In extreme cases the developer may be required to constrain operation of a wind turbine under some circumstances, though the loss of output is generally low.

8.4.6 Public safety

Blade loss or other catastrophic damage to wind turbines is very uncommon, but can happen, so must be taken into account in the development process. Planning authorities may specify particular setback distances from public roads and/or buildings, with typical values ranging from 1.5 $\times$ tip height to 500m depending on the country and authority.

The probability of impact at a particular location is a function of the distance from the turbine, the kinetic energy of the blade fragment, and the position of the rotor at the time of blade failure. According to US sources '...where information is available, the majority of the major components (rotor, tower, and nacelle) have fallen to within 1 to 2 hub-height distances from the base'.

Ice throw requires similar consideration when siting turbines in cold climates. Fragments of ice can be thrown from blades, usually during rotor startup, and risk assessment is again necessary at the planning stage. Experimental studies indicate a 'risk circle' around the turbine based on:

$$d = (D + H) \times 1.5 \tag{8.2}$$

where d is the maximum throwing distance (m), D is the rotor diameter (m) and H is the hub height (m).

Modern wind turbines are equipped with temperature sensors and employ power condition monitoring to detect ice build-up, and curtail operation if necessary.

8.5 Technical constraints

8.5.1 Aviation collision risk

Wind turbines can pose two main problems for civil and military aircraft, namely collision risk and interference to radar. Collision risk is highest in proximity to airfields or in areas of military low flying. In the UK safeguarding of civil aerodromes is governed by the Civil Aviation Authority (CAA) whose publications CAP168 and CAP174 are relevant. In particular CAP168 gives guidelines for the maximum allowable height of obstructions near airfields: height restrictions are defined by zones called the inner and outer horizontal surfaces (IHS and OHS) centred on the aerodrome runways, and by the runway take-off and landing flight paths.

Typically the IHS may extend to 4km from the aerodrome, within which the maximum obstacle height is 45m above runway elevation; the OHS extends to 10km radius, with corresponding height limit of 150m. If a runway take-off or landing path is lower than these limits, it takes precedence, i.e. a lower height restriction applies. Generally speaking the more major an airfield the more severe the restrictions; runways with instrument landing systems (ILS) are subject to stricter obstacle height limits than those with visual approach only.

Military low flying areas may simply be designated as unsuitable for wind turbine installation, though in some cases developments are permitted below a certain tip height. The defence authorities must be consulted during the planning process and usually have the power of veto on projects, though are sometimes prepared to compromise.

8.5.2 Radar

Interference to radar is a potentially significant impediment to wind energy development, as large land areas may be excluded on a precautionary basis. In the UK at one time around 30% of consented wind energy schemes were being held back due to radar concerns, either civil or military. There are different types of radar, with different implications for wind turbines.

Primary surveillance radar (PSR) operates by echo-locating: the scanning radar antenna detects aircraft range and bearing, but cannot discriminate height above ground level. As a result wind turbines may be mistaken for targets. PSRs incorporate a moving target indicator function to filter out stationary objects based on Doppler returns, but wind turbine rotors still appear due to the movement of the rotor blades.

Standard techniques for removing unwanted 'clutter' include range-azimuth gating (RAG), whereby radar returns from certain areas are ignored, but with a penalty in terms of loss of coverage; as a result RAG is rarely used by UK radar operators. Large windfarms can therefore cause problems for PSRs, as the aircraft flying over the windfarm cannot be reliably tracked due to false plot initiation or track 'seduction'.

Potential solutions to primary radar interference include:

- Terrain shielding, i.e. siting the wind turbines out of view of the radar if possible
- Re-locating radar sets or providing alternative units (expensive, but has been done)
- Use of improved filtering algorithms
- Reduced wind turbine radar cross section (RCS), i.e. 'stealth'

The last two options are being actively pursued in the UK. New civilian radars by Raytheon are currently being introduced with greatly improved clutter discrimination, particularly in regard to wind turbines; in addition research into radar absorbing materials (RAM) has led to trials of a large wind turbine with stealth blades of low radar cross section.

Interference to secondary surveillance radars (SSRs) is less common, but has been reported. SSRs are semi-passive, sending out an interrogating beam that triggers a response from the aircraft, giving its identity and height. False plots on the SSR at Kastrup

(Copenhagen) Airport were attributed to the presence of the nearby Middelgrunden offshore windfarm: This large array (20 x 2MW turbines) was reflecting the SSR outgoing ('uplink') signals, causing responses from aircraft in unexpected directions.

8.5.3 RF and microwave communications

Wind turbines can potentially interfere with radio frequency (RF) or microwave link transmissions if sited too close to the centerline of the beam-path, transmitter, or receiver. The link operators therefore apply conservative rules for link avoidance, e.g. 100-500m clearance zones around the beam centreline, within which no development is permitted. This can sometimes exclude useful wind energy sites.

The principle interference mechanisms are scattering and reflection, which are described in detail by Bacon in the standard reference used by the UK windfarm industry. Scattering may reliably be avoided if no part of the wind turbine intrudes within the 2nd Fresnel zone (RF2) of the link; the radius RF2 is dependent on link frequency f and the distances d1 and d2 from the link ends according to:

$$R_{F2} = \sqrt{\frac{600 d_1 d_2}{f(d_1 + d_2)}} \tag{8.3}$$

where R_{F2} is the 2^{nd} Fresnel zone radius (m), d_1 and d_2 are the distances from the two ends of the link (km) and f is the link frequency (GHz).

If the positions of the link ends and the turbine can be located with sufficient accuracy then the above equation gives a good guarantee of link avoidance, and may be used in preference to the more conservative clearances specified above. Link operators may nevertheless seek mitigation and/or guarantees as a condition of consent. Reflected signals can also cause interference, due to the phase lag between the direct transmission and the signal reflected off the interfering object (i.e. the wind turbine). This gives rise to the familiar 'ghosting' sometimes seen on analogue TV screens. Avoidance of reflection interference is achieved by specifying a minimum distance between the wind turbine and the link ends (transmitter and receiver) and/or use of antennae with high directional discrimination.

8.5.4 TV interference

Television interference is essentially a special case of RF interference (above) and amenable to the same analysis. Planning authorities may require guarantees that a developer will make good any TV interference caused by the wind turbines, but this is rarely an issue nowadays due to the proliferation of digital TV services (via satellite, cable, or terrestrial transmission): digital TV is generally less susceptible than analogue to the kind of interference described.

8.6 Other issues

Pollution risk: this refers mainly to the risk of surface water run-off from access roads, which may cause silting or pollution of natural watercourses. It can be avoided by appropriate design of access tracks, which are typically constructed on a base of compacted stone, without tarmac or concrete surfaces, and so remain porous and well drained. Pollution risk is normally assessed in the environmental impact assessment EIA.

In some areas peat slide risk is also an issue, requiring careful design of access roads and construction areas, avoiding where possible deep peat deposits (these may also be protected under Annex 1 of the EC Habitats Directive). The Scottish Executive has recently published details of peat slide risk and methods for its avoidance.

Archaeological assessments are frequently required where wind turbines are proposed in close proximity to historic sites, which may be protected by law. UK planning guidelines are covered by NPPG1821, which describes the protection for listed buildings, conservation areas, world heritage sites, historic gardens, and designed landscapes and their settings. Developers may be required to carry out pre-construction impact assessments, sometimes including exploratory digs.

9 Construction and offshore

This section summarises various aspects of onshore wind construction, followed by an overview of current offshore wind developments in Europe.

9.1 Foundations (onshore)

Some aspects of civil works are considered, beginning with a comparison of different foundation types for onshore wind turbines. The basic types are (a) standard slab foundations, i.e. gravity bases, (b) rock anchors, and (c) monopiles. Most windfarm foundations tend to be gravity bases or rock anchor solutions.

Both foundation types require a comprehensive steel reinforcing cage to be pre-built into an excavated pit, approximately 2m deep. The steel reinforcement may weigh 30-40t for a foundation this size. Different methods are used to construct the tower interface: either to integrate a tubular steel 'can' in the reinforcement or to incorporate long bolts into the centre of the reinforcing mesh.

Other foundation types include post-tensioned rock anchors, either directly into bedrock, or in combination with small concrete plinths at the surface. Rock anchors eliminate the need for a heavy gravity base and use far less concrete, hence are useful in remote or inaccessible areas. They are also potentially more cost effective, and with lower environmental footprint. Disadvantages are the need for specific rock characteristics capable of supporting the loads, and the need to use specialised contractors.

9.2 Access roads

Wind farm access roads are usually constructed from compacted stone, topped with finer crushed gravel. For cost and environmental reasons they are rarely tarmac or concrete. The track must carry the load of turbine delivery vehicles and cranes, where the latter may weigh up to 200t. Tracks must be designed with adequate drainage, and avoiding peat slide risk.

Typical track dimensions are a running width of 4-5m, and depth of stone/gravel of 300-400mm assuming relatively even underlying soil. The minimum bend radii are dictated by vehicle and load length (particularly the blade and tower transport vehicles), and will be specified by the turbine manufacturer and/or transport company. The track load-bearing capacity is dependent on crane weight. Maximum track gradients of 6-8° are typical, though steeper gradients can be allowed for short stretches or where assisted traction is used.

Public road access must also be considered well before construction, via route surveys and desktop 'swept path' analyses. As with the site access tracks the main concerns are vehicle turning radii, road width, load overhang at corners, maximum gradients, and load bearing capacity (especially over bridges or culverts). Where necessary, temporary road improvements can be made to facilitate delivery.

9.3 Cranes

The requirements for cranes during windfarm construction include lifting capacity, suitability for the terrain, design of site access and crane hardstanding areas, crane size and weight in relation to transport infrastructure, and wind speed limits for the main lifting operations. Crane lifting capacity is a function of jib height and load radius.

Most medium and large wind turbines are erected using two cranes, one large machine capable of the main lifts (tower sections, nacelle, and rotor) plus a smaller 'handling' crane to steady and align the loads. Some manufacturers have devised lifting arrangements that allow the rotor to be installed using only the main crane without need for a steadying crane: a special rig is attached to the hub that allows the rotor to be rotated from the horizontal to vertical plane during the lift.

The possibility of erecting wind turbines without crane is possible with rock anchor foundations. In this case a hand-winch set-up is used, involving a steel A-frame attached to the hinged tower base, to facilitate the initial stages of the winching operation (when winch cable loads are highest).

9.4 Offshore wind

The second part of the lecture deals with offshore wind developments in Europe at a relatively broad level. The overall advantages and disadvantages of offshore, as opposed to onshore, wind projects may be summarised as follows:

Advantages

- Higher wind speeds offshore, yielding more energy
- Lower turbulence, resulting in reduced fatigue
- Visual impact less than with onshore arrays
- Noise impact also less than onshore
- Large arrays bring economies of scale in development and construction

Disadvantages

- Construction and installation cost are higher
- Projects are less accessible for maintenance
- Greater potential for corrosion

- Higher financial risks involved

Current foundation types include:

Gravity bases/caissons: used on near-shore projects in shallow water depth (<10m); the foundation consists of a large concrete or steel base which rests on the seabed; cost is approximately proportional to water depth squared. For installations into sand or soft clay suction caissons have more recently been proposed.

Monopiles: widely used in shallow-medium water depths (5-20m); steel piles 3.5-4.5m in diameter are driven 10-20m into the seabed with the wind turbine tower installed on top.

Subsea jacket structures: suitable for deeper waters (30-50m) these are steel jacket designs based on offshore oil and gas technology; they are unsuitable in shallow waters due to restricted vessel access. Tripod jackets have been proposed.

Water depth and distance offshore are critical factors in determining cost. The superior wind characteristics of offshore wind projects are almost ideal wind conditions which is used to justify the additional cost.

10 Economics and case studies

This section begins with a brief assessment of the production cost of wind energy from both onshore and offshore sources, including the additional costs of balancing (intermittency). Some simple calculations are given based on existing projects. Intermittency costs are assessed with reference to recently published work from EWEA and National Grid. The total cost of wind power is then compared with other generation sources.

The price (as opposed to cost) of generated electricity includes various subsidies or incentives, which vary from country to country. The purpose of these is to make wind energy a sufficiently attractive possibility for short term investment, and hence ensure sufficient capacity is brought on stream within government timescales. The lecture considers the effectiveness of recent UK policy in this regard.

The second half of the lecture considers two project case studies, some technical details of which were covered in earlier lectures. The first is the Mackies wind farm in Aberdeenshire, and example of parallel onsite generation where the wind farm is owned by a business and supplies a local load demand. The second example is the Isle of Gigha wind farm in Argyll, a small grid-connected wind farm owned entirely by the local community.

The section concludes with some thoughts on the different ownership possibilities for wind energy projects, their advantages and disadvantages.

10.1 Cost of energy

10.1.1 Generation

The lifetime unit cost L of energy generated by a wind farm may be simply estimated using the following expression:

$$L = \frac{C}{E} + O \tag{10.1}$$

where C is the capital cost of development, E is the energy produced over the project's lifetime (MWh) and O is the unit operating cost e.g. as £/MWh.

The unit cost L (£/MWh) is then dependent on the amount of energy produced at the site, which is a function of project capacity factor and lifetime. The calculation ignores financing, interest rates, and inflation, but is nevertheless useful for 'ballpark' estimates (the capital cost is fixed at construction, but in a real case O&M costs are likely to inflate roughly in line with energy costs).

Operating costs O can be estimated from current experience, for instance using the price quoted for full-cover warranties from the leading wind turbine manufacturers, plus additional costs for all-risks such as insurance, business rates, and other miscellaneous charges. For onshore wind projects the total O&M cost is usually in the range 10-15% of annual turnover.

More representative calculations must take into account an assumed discount rate to account for inflation, plus an allowance for and cost of balancing power, i.e. the additional generation capacity required to allow for the intermittency of wind energy.

10.2 Intermittency

Wind energy is an intermittent source, and other generation plant must be available to take up the load demand when wind conditions are insufficient. The resulting additional cost is known as the intermittency or balancing cost. Studies from Germany, Denmark and Finland show that for 90% of the time, power swings of dispersed wind generation within one hour are less than 5% of the installed wind capacity, and the additional reserves required on the system are only a few percent of the installed wind capacity.

It should also be noted that a requirement for backup generating plant is not confined to wind or other renewables. Due to the inherent variability of national electricity demand there must always be a percentage of spare capacity or 'spinning reserve' available for rapid despatch. In addition large thermal power stations are not 100% reliable, and backup stations are needed to cover when generation is lost through unscheduled faults or trips.

10.2.1 Total cost of generation

The total cost of wind energy at the point of generation is calculated taking account of capital expenditure, O&M, and balancing (intermittency) costs. Rigorous estimates include discounted cash flow adjustment using appropriate interest rates. The overall cost of energy from wind power is assessed in detail in a 2008 Risø report by Jørgen Lemming et al. These authors assume a project lifetime of 20 years, 7.5% discount rates, and average installation costs of average of 1175 €/kW. Operation and maintenance costs of 1.45 c€/kWh are assumed over the lifetime of the turbine.

The above report concludes that the total generation cost "...ranges from approximately 7-10 c€/kWh at sites with low average wind speeds to approximately 5-6.5 c€/kWh at good coastal positions, with an average of approximately 7c€/kWh at a medium wind site". In UK currency these figures translate into £63-90/MWh for low wind sites, £45-59/MWh for good coastal sites, and an average of £63/MWh.

10.3 Case study: Onsite generation - Mackies

Mackies is an established family farming business producing ice-cream, based in Aberdeenshire, north-east Scotland. Ice cream production involves relatively high electricity usage associated with refrigeration, for both the initial freezing process and subsequent cold storage. In 1993 Mackies consumed around 1.8GWh electricity per annum at a cost of over £100k; today the consumption figure has grown to 3.6GWh. This increase would normally have created a significant added cost for the business, but in 2005 Mackies installed its own medium-sized wind turbine to feed power into the dairy, and in 2007 added a further two machines. Today Mackies is a net exporter of electricity.

10.4 Case study: Community ownership - Isle of Gigha

The Island of Gigha is found in Argyll, on the west coast of Scotland; its area is 1395 ha, and present population around 130. During the late 20th century Gighas population was in serious decline: in the early 1970s there were 200 inhabitants, but by 2002 this had fallen to fewer than 100. In order to stem the decline and provide a long-term economic future for the island, the inhabitants of Gigha in 2002 took the island into public community ownership.

The Isle of Gigha Heritage Trust was founded to manage and run the island, and one of its first priorities was to identify viable economic activities to help sustain the Islands future. A feasibility study commissioned in 2002 identified wind energy as an attractive possibility for Gigha, and indicated suitable areas of the island where a single wind turbine might be sited. This led to further, more detailed studies, focussing on key issues including:

- Wind exposure
- Road access for transport
- Proximity to 11kV grid
- Separation from dwellings
- Noise avoidance
- Visual impact
- Wildlife
- Aviation

The Gigha Heritage Trust purchased three used Vestas V27s from Windcluster, the owners of Haverigg windfarm in Cumbria. The Haverigg machines were inspected in situ and service records made available.

The total capital cost of the Gigha project was £440k (£652/kW); this is significantly lower than average for the time, mainly due to use of second-hand turbines. Development costs were nonetheless kept low, and proportionate to project size and budget. Planning was made much easier by the strength of community support, and lack of objection to the project.

Operating costs are currently of the order 10% of gross turnover; this is again relatively low, but note that there is no manufacturer's warranty on the wind turbines. The simple cost of energy from the windfarm, assuming an 8-year lifetime on Gigha will be around £45/MWh; if the turbines are kept running for 10 years the figure drops to around £40/MWh.

Energy from the windfarm, and a percentage of ROC benefits, is sold via a commercial power purchase agreement (PPA) with a licensed supplier. Current annual income to the island is of the order £75k after all financing and O&M costs are met; when 100% ownership of the windfarm is achieved with all financing costs paid off the income to the island will roughly double.