Article 5: The Single Wind Turbine: From the Wind to the Blades

In this article, we bring the reader along on a tour of an individual large modern wind turbine up close, introducing the key components that allow it to harness the wind's energy and convert it into mechanical energy. We begin by noting the size of the turbine and the layout of the wind farm in which it is located. We then explain why a turbine looks as it does today: why it has three blades, why the blades taper and twist, what limits how quickly the blades rotate, and how the blades generate power. We also tour the inside of the turbine, looking at the key components and control systems within the nacelle.

5.1 The Turbine's Blades and Tower

Visiting a wind farm is markedly different from visiting a solar farm. Each wind turbine stands tall, separated from its neighbors by several hundred meters or more. Surrounding each turbine is open space - often farmland with animals grazing or crops growing. In some cases other infrastructure (oil and gas wells, for example) shares the land. Figure 5.1 shows one such scene, although many are not as idyllic.

Figure 5.1: Wind turbines in a bucolic setting. Photo: Symbiot/shutterstock.com, https://www.shutterstock.com/video/clip-10171613-stock-footage-summer-countryside-with-wind-turbines-and-agricultural-field-with-grazing-cows-full-hd-p.html.

As you approach an individual wind turbine, its enormity becomes apparent. You realize that the blades and tower must bear the force of the wind pushing them backwards, and they must be very strong to resist this force. For specificity, we assume the rated power of the turbine is three megawatts, which is a typical value for the large turbines sited on land recently. (Offshore turbines are generally larger.) The tower stands 80 meters tall, and that's not including the blades, which make it taller still. It is an upright, cylindrical structure, several meters in diameter, tapering as its height increases. The tower rests on a large concrete foundation.

This is the most common modern tower. It is assembled onsite from a few tubular sections. The towers supporting some older wind turbines have a steel lattice-work structure that requires extensive on-site construction. Some newer towers are beginning to be constructed from concrete, either assembled onsite from modular sections or even cast onsite.

A pod, known as a nacelle, is sitting on top of the tower. At the front of the nacelle is a hub, which is where the blades meet and connect. Together, the hub and blades make up the rotor, so called because it rotates as the wind blows. As with all of the other components in front of you, the blades are enormous too. Each blade is 50 meters long, so the total rotor diameter reaches 100 meters. Looking up at the turbine, you see that there are three blades. The choice of three blades is a signature example of the trade-off between efficiency and cost. A wind turbine's sole purpose is to convert wind energy into electrical energy. To do this effectively, it must capture as much energy as possible from the incoming wind. Having more blades allows the turbine to "sweep" more air per revolution, providing the potential to capture more of the incoming wind energy, but at the expense of increased weight, complexity, and cost.

To reduce costs, a turbine could use fewer blades, perhaps only two. To generate the same amount of energy as a higher number of blades, two blades would need to sweep through the air more quickly. However, other issues arise. The noise and vibrations originating from the blades could quickly exceed acceptable limits, as the blade tip is traveling too quickly. The structural strength of the blade becomes one concern and noise becomes another. Therefore, the rotation speed needs to be limited, which in turn reduces the energy production efficiency. The result is a lower cost turbine, but also one that produces less power. On balance, the power penalty makes the two-blade design more expensive.

In the middle lies the Goldilocks compromise: the ideal balance that maximizes the ratio of energy extracted to cost. Across the wind industry, this compromise is now settled at three blades.

A Damaged Blade is on the Ground

Land is so cheap at this Texas wind farm that a blade that had to be removed because it was damaged was not trucked away but was left resting on a mount on the site, so you can see it up close (see Figure 5.2). What an opportunity to look closely at the blade's shape.

Figure 5.2: A 50-meter blade stored at ground level, being examined by one of the authors (Greg Davies) at the Sherbino 2 wind farm in western Texas. The blade is viewed from its root, looking toward its tip; it is thicker at its base, twisted, and tapered. Photo: Ryan Edwards.

You see that the blade, in cross-section, has a similar shape to an airplane wing (the shape is called an airfoil). You note that the blade is thick for about the inner third of its length and then tapers down gradually to a smaller size at the tip. You are surprised to see that the blade is twisted along its length: the blade's front edge (called the leading edge, since it contacts the wind first) points in a different direction depending on whether you are looking at the blade root, which is the base of the blade near the hub, or the blade tip, or somewhere in between. Let's understand these three features of the blade one at a time.

Why does a blade have a cross-section like an airplane wing?

A great deal of design effort ensures the turbine blade is shaped so that the energy in the wind is harvested efficiently. A cross-section through the blade reveals that it has a shape called an airfoil (see Figure 5.3). In many respects, a turbine blade is similar to an aircraft wing, except that instead of incoming air producing a lift force that pushes an aircraft upwards, the incoming air creates a lift force that rotates the blades, hub, and shaft. In doing so, the blades extract kinetic energy from the wind and transform it into rotational kinetic energy, which is then harnessed in the turbine's mechanical and electrical systems to generate electricity.

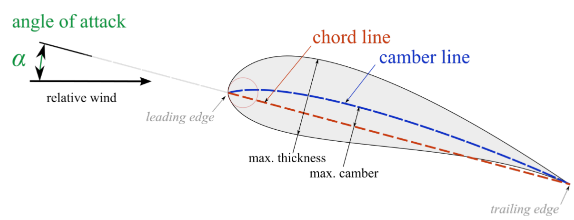

Figure 5.3: Features of an airfoil of both a wind turbine blade and an airplane wing. Source: Olivier Cleynen, https://commons.wikimedia.org/wiki/File:Wing_profile_nomenclature.svg.

The rounded front edge of the blade is the leading edge and the sharper rear edge is the trailing edge. The incoming airflow comes into the leading edge, and the outgoing air leaves from the trailing edge. The chord of the airfoil is the straight line joining the leading and trailing edges. The camber line is a curved line that lies midway between the upper and lower surfaces of the airfoil and tells us how asymmetric or curved the airfoil is.

The orientation of the airfoil with respect to the incoming wind determines how effectively lift is produced. It is quantified by the "angle of attack," which is the angle between the chord line and the incoming airflow (see Figure 5.3). A symmetrically shaped airfoil will produce no lift at a zero angle of attack; however, an airfoil with camber will produce lift even at zero angle of attack, due to the asymmetry in its shape. As the airfoil angle of attack is increased, the lift force increases approximately linearly, up to a certain angle. Tilt too far, and the amount of lift will drop dramatically. This is called stall, and it occurs when the air flowing around the airfoil is no longer hugging the airfoil surface closely but begins to separate from the airfoil.

Why are the blades tapered toward the tip?

Two reasons: first, the taper equalizes the energy generation along the entire blade. If a blade were not tapered but rather had the same chord length from the hub to the tip, more energy would be extracted from the wind at the tip than at the root, and the result (not obvious) would be a reduction in the total rotational energy produced by the same incoming wind. Second, the taper reduces the bending of the blade by reducing the wind forces near the tip. A larger force on the tip would create structural requirements to carry that force all the way to the root, so the taper helps to avoid excessive blade bending, fatigue, and failure.

Why are the blades twisted?

An oversimplified answer is that the blades are twisted because when the blades are spinning, the air hits the tip of a blade and the base of the blade from very different directions. This is because the blade tip is traveling far faster than the blade root.

Imagine that you're strapped to the tip of the rotating blade, with your face pointing in the direction of rotation. From which direction does the wind appear to be coming? As the velocity of the blade tip is much faster than the incoming wind, the apparent wind (known as the relative wind) is moving almost directly toward your face, and only a small component of the wind velocity is hitting your face from the side. For a tip speed ratio of 6, the angle of the relative wind into your face is about 10 degrees. The relative wind direction is almost entirely in the plane of the rotating blade, and hardly matches the direction of the distant incoming wind at all. Said another way, the wind is coming at you from a completely different direction than if you were standing stationary in front of the turbine. Thus, not only is the relative wind speed much greater at the tip, but it also has a different direction than the relative wind direction near the blade root.

Each airfoil section of the turbine blade must have a small angle of attack with respect to the relative wind velocity. At the tip of the blade, where the relative wind velocity is almost completely in the plane of the blades, the leading edge of the blade must point almost in the direction of rotation. Close to the root of the blade, where the relative wind velocity is coming largely from the wind direction, the leading edge of the blade must point nearly into the incoming wind. In short, the blade must twist. The optimal angle of attack for each section of the blade takes into account drag forces as well, which oppose the motion of the blades much like drag acting on a car. There is typically an optimal angle of attack that maximizes the ratio of the lift force to the drag force.

Compromises involving aerodynamics, structural forces, materials, and costs

The details of the shape of a blade are the outcome of a balancing of the needs for adequate structural strength and the need for aerodynamically efficient power generation, all the while taking into account the limitations of the materials from which the blades can be made and minimizing the total cost. In turn, the cost incorporates issues of safety, longevity, and maintainability. The outcome of this optimization, oversimplifying, is a design of the inner third of the blade (which carries almost the entire blade load) that is dominated by structural considerations, resulting in a thick airfoil shape. In contrast, the optimized design of the outer third of the blade, where almost all of the turbine power is generated, is driven by aerodynamics. The shape tapers significantly to optimize the distribution of the loads along the length of the blade and prevent a concentration of forces at the tip. The middle third is where the trade-off of aerodynamics and structure allows for a variety of solutions.

The components of a wind turbine must be able to withstand immense stresses over their design lifetime, which is at least 20 years. A particularly important stress is the "fatigue load" that results from the blades moving through turbulent air, where the local angles of attack, and therefore the loads, are constantly changing. A second stress is the cyclic gravitational load, which is the force that opposes the blade's rotation when the blade is ascending and augments the blade's rotation when it is descending; the weight of the blade is important here. A turbine blade may go through forty million rotations over its lifetime.[1]

The Turbine Blades are Spinning

The blades on the turbine in front of you are spinning. You hold an anemometer up, a tool that is used to determine wind speed, and measure the wind high above you blowing at a steady but moderate speed of 9 meters per second (about 20 miles per hour). You count the amount of time it takes for a blade to complete a full revolution: about six seconds. Therefore, the rotor spins at 10 revolutions per minute (rpm). You notice that the tip of the blade is moving much more quickly than the section near the root, as it has to travel the full circumference of the rotor with each revolution. On this 50 meter blade, the tip is traveling at 52 meters per second, or 120 miles per hour. This is much faster than the speed of the incoming wind. The ratio of the tip

[1] Authors' estimate: A typical rotational speed for a wind turbine producing electricity at its maximum rate is six seconds per rotation; a blade rotating at that speed will complete five million rotations each year. Because much of the time the wind is not strong enough to produce maximum power (and sometimes not strong enough to produce any power), the number of rotations of the blades in a year is considerably lower, perhaps two million rotations per year, in which case in twenty years the blades will experience forty million gravitational load cycles.

speed to the speed of the incoming wind is called the tip speed ratio; it is about 6 in this case (52 meters per second divided by 9 meters per second is about 6).

This is an important observation for large, modern wind turbines. Even with the rotor spinning relatively slowly, the tips of the blades are still traveling very quickly, due to their sheer length. What does it mean for the design of a turbine that the tip speed is much higher than the incoming wind speed? First, to maintain noise within strict regulatory limits, the tip of the blade must travel below a maximum speed, typically about 80 meters per second. This is one reason modern turbines with long blades turn slowly - increasing the blade length means the rotation speed must decrease accordingly. For the largest turbines, with blades around 80 meters in length, typical rotation speeds might not exceed 10 rpm.

You Walk Halfway Around the Turbine

Turbine efficiency

With the wind still blowing hard, you now walk from the front of the turbine around to the rear and stand behind the turbine. You hold up your anemometer again and see that the wind is still blowing quite strongly. However, it is blowing less strongly than in the front. You conclude that some of the energy in the incoming wind has been extracted by the turbine, but not all of it.

In fact, it is impossible for a wind turbine to convert all the wind energy that hits the blades into electrical energy. The slower the speed of the wind behind the turbine, the more energy the turbine has extracted from the incoming wind. However, with quite general assumptions, there is a limit on the maximum amount of energy that a wind turbine can extract from the wind, known as the Betz limit. At the Betz limit, a turbine has an efficiency of 59 percent. (In more technical language, at the Betz limit the electric energy output is 59 percent of the kinetic energy in an incoming wind that, if the blades were not turning, would strike a disk of equivalent size to the swept area of the turbine blades.) At the Betz limit, it turns out, the air downwind of the turbine has slowed down to exactly one-third of the wind speed far upwind of the turbine. (At the turbine itself, at the Betz limit, the wind speed has dropped to two-thirds of the upstream wind speed). So, the wind speed behind a real turbine must be more than one-third of the speed of the incoming wind.

To gain some insight into why an optimal turbine efficiency exists, it is helpful to consider two limiting cases: no slowing down of the wind, and complete slowing down of the wind. At one limit, if the incoming wind travels straight past the turbine without slowing at all, the wind won't be losing any energy, and therefore the turbine cannot be producing any energy. Clearly the turbine needs to slow the wind in order to extract some energy from it. At the other limit, if the turbine brings the air to a complete standstill, then no wind would pass through the turbine (think of it like a grid-locked highway). With no wind traveling through the turbine, it would produce no energy, as before.

In between lies the optimum: the turbine must slow the air in order to extract energy from the wind, but there must still be sufficient wind speed behind the turbine to allow the "used" wind to escape and make room for a continual flow of "new," energetic wind through the turbine. With careful design it turns out that modern wind turbines are able to get close to the Betz limit: typical efficiencies range from 40 to 50 percent – that is, from 70 to 85 percent of the ideal Betz efficiency.

An enormous amount of wind energy influences the turbine. For our three-megawatt turbine in a nine meter per second wind, almost 90 metric tons of air flow past the turbine every second, from which roughly half of the kinetic energy is extracted and converted to electricity.

Wakes

Still standing behind the turbine, turn your back to the turbine and imagine the air going downwind. Air that flowed around the turbine is mixing with the air exiting from the turbine blades, creating a "wake." The mixing helps restore the speed of the wind behind the turbine – more completely as the distance downwind increases – until eventually the wind recovers its initial strength. You notice a second wind turbine in the distance; if it is close enough, the wake has lowered its output, compared to what that turbine would have produced from unaltered wind. When the wind farm was laid out, this negative wake-turbine effect was taken into account and the spacing between the turbines was made large enough to minimize losses. Indeed, the wake-turbine effect can limit the number of turbines that are sited on a wind farm.

The Wind Speed is Varying

You return to the front of the turbine and watch how the turbine performs as the wind speed varies. You get out your anemometer again. The wind has died down, all the way to zero, and the blades are not spinning. But then the wind slowly picks up, increasing to 1, then 2 meters per second, but the turbine blades are still not spinning. However, as the wind speed increases slightly more, to 3 meters per second, the blades change their orientation to the incoming wind (their pitch), and the turbine slowly begins to spin up. The speed where the blades first start to rotate is called the "cut-in" wind speed; it is the minimum wind speed at which a turbine has been designed to produce power.

The wind speed continues to increase, and you see that the rotational speed of the blades is also increasing, approximately in proportion to the increase in wind speed. As the wind speed increases, the turbine produces more and more power. You notice that, compared to the fully stopped position, the leading edge of the blades points more towards their plane of rotation and less straight toward the oncoming wind. The pitch of the blades is being adjusted by its control system to provide a near-optimal angle of attack for the incoming air.

The wind now is really blowing: its speed has reached 12 meters per second (27 miles per hour) and is climbing higher. You notice that once the wind speed exceeds 12 meters per second, the blades turn no faster than at 12 meters per second. Their rotation speed is about 16 rpm. The turbine has reached its power limit (the "rated" power). You notice, however, that the blades continue to change their pitch as the wind speed increases. This pitching keeps the power output constant, by reducing the angle of attack, and therefore the lift force, on the blades.

Finally, as you brace yourself against the wind, its speed reaches 25 meters per second (56 miles per hour). The blades change their pitch until the leading edge of each blade is pointing directly into the oncoming wind, and the turbine comes to a standstill. The turbine's "cut-out" speed has been reached, above which a turbine control system shuts the turbine down to assure that it is not damaged in sustained high winds.

The turbine's power curve and the wind speed

All this information is summarized in the "power curve" for a wind turbine. An idealized but representative power curve for a turbine is shown in Figure 5.4, left panel. It shows the cut-in speed, the rapid increase in output power between the cut-in speed and the rated speed, the plateau in output from the rated speed to the cut-out speed, and the fall of output power to zero above the cut-out speed.

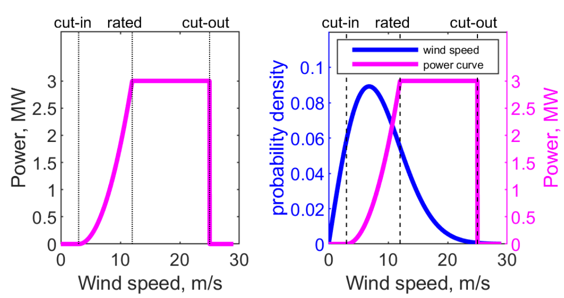

Figure 5.4: *Left:* An idealized power curve for a wind turbine. *Right:* The same power curve is superimposed on a representative wind-speed distribution.

You might wonder what fraction of the turbine's total energy production over a year is produced at wind speeds between the speeds for cut-in and rated power, and what fraction at wind speeds between the rated speed and the cut-out speed. The right panel of Figure 5.4 superimposes the wind probability at a particular site (blue curve) onto the turbine's power curve (pink curve, identical to the curve in the left panel). The relative positions of the two curves are typical: the site's most common wind speed (peak of the blue curve) is lower than the turbine's rated speed. Most of the time, the speed of the wind at the site is between the turbine's cut-in and rated wind speeds.

Some of the energy in the winds that blow faster than the turbine's rated speed (the wind speeds to the right of the dashed vertical line at 12 meters per second) is thrown away. That is the meaning of the fact that the turbine's power curve stops rising at its rated velocity. The choice of the turbine's rated speed for a given site is a compromise. Consider choosing a lower rated speed for the same site. In Figure 5.4, right panel, this would mean squeezing the turbine's power curve toward the left (perhaps, squeezing it so much that the rated speed coincides with the maximum of the wind-speed distribution). Such a "derated" turbine would be smaller and less expensive, but now more of the turbine's energy would be produced from winds above the rated speed, inefficiently. On the other hand, choosing a higher rated speed would result in a more expensive turbine producing only a limited amount of additional electricity from the high winds.

Figure 5.5 shows the electricity generated at various wind speeds for a specific turbine. This is a three-megawatt wind turbine, located in Colorado. In both panels, the leftmost bar represents the turbine not producing any power, and the rightmost bar represents the turbine producing at its full rated power, three megawatts. The ten bars in between represent the wind power grouped into power increments of three-tenths of a megawatt: i.e. 0 - 0.3, 0.3 - 0.6, and so on.

The left panel shows the fraction of time that this wind turbine generated various levels of power. Much of the time very little power is produced; for example, the first bar tells us that almost 10 percent of the time no power is produced at all, and the second bar tells us that just over 20 percent of the time the turbine is producing between zero and three tenths of a megawatt. Combining the first four bars, almost half the time the turbine is producing at less than one third of its maximum output power. From right-most bar, we learn that the turbine produces its full three megawatts of output power 20 percent of the time. This is the power generated when the wind speed is greater than the rated speed – in this case, 12 meters per second. The right panel conveys the same information in a different way. It shows the fraction of the total energy produced by the turbine, for the same wind speeds as in the left panel. Very little of the turbine's total energy production is at low wind speeds. Almost half of the total output of the turbine occurs when the turbine is operating at three megawatts, its rated power, even though (left panel) the turbine operates at its rated power only 20 percent of the time.

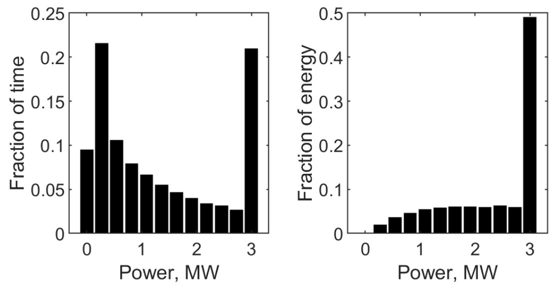

Figure 5.5: *Left:* The fraction of the time that a three-megawatt turbine produced power at various narrow ranges of rates, in megawatts (MW). *Right:* The fraction of the total power that the same turbine produced at these same rates. In both cases the first and twelfth bars show the fractions for no power at all and for full power, respectively. The ten intermediate bars show intervals of three tenths of a megawatt.

The Wind Twists and Turns. So the Turbine Moves like an Owl's Head.

While you are observing the wind turbine, the wind direction changes several times. You notice that the nacelle pivots on the tower so that the plane of the blades always faces the wind. This movement of the entire nacelle and blades is called yaw motion, and it is accomplished by yaw motors in the nacelle. If it were not for this adjustment to track the wind direction, the incoming wind would strike the turbine at an angle, and the turbine power output would drop.

Tracking the wind direction is an important form of wind turbine control. Turbines have a small weather station mounted on top of the nacelle, with an anemometer to measure wind speed and a wind vane to measure wind direction. These data are fed to a turbine-control system that engages a yaw motor to rotate the nacelle into the wind. Thus, the performance of the wind turbine is critically dependent on the wind vane.

The turbine cannot track the wind direction instantaneously, because the nacelle is heavy and the rotor has large loads acting on it. Instead, the yaw control system moves slowly. If the wind direction fluctuates rapidly, the wind turbine designer does not want the nacelle to chase after every change in the wind direction – to move toward the east from the north for ten seconds, for example, and then immediately back toward north. The majority of turbines have a very simple control strategy: only after the wind direction (or heading) has deviated by two to four degrees will the turbine yaw to the new wind heading. This assures that the turbine is always oriented nearly optimally, while avoiding too frequent excursions. With this control strategy, the losses from misalignment can be kept below one percent.

5.2 Inside the Tower

The person in charge of maintaining the wind turbine joins you and invites you to head into the tower to see what's inside. You enter the tower through a small door. Immediately you realize that the tower is a hollow steel shell. At the entrance level, the tower floor is crowded with electronic equipment, evidence that wind turbines have become "smart" devices that integrate the signals from dozens of sensors to lower overall costs. Some controllers adjust the rotational speed and the pitch of the blades to assure optimal aerodynamic efficiency. Other controllers minimize the accelerations and decelerations of the blades and the mechanical vibration of the blades, gearbox, and towers produced, for example, by gusts of wind.

You look upward and notice vertical ladders ascending the structure, with intermediate platforms located periodically at various heights. There is also a service elevator that lets you travel from the bottom to the top of the structure. It can carry personnel as well as smaller parts and tools for turbine servicing and repairs. You see a significant amount of wiring routed down the tower. These wires carry the power generated up in the nacelle, as well as control signals and information, to the tower base. The power is then fed into a collection system that gathers power from the whole farm.

You ride the service elevator up almost 80 meters and then climb the final section on a ladder into the nacelle. The nacelle is large, about the size of a school bus. You are able to walk around on narrow platforms, although the space is filled with a lot of equipment and is very tight. Particularly prominent are electric power generators and also the yaw motors that move the entire nacelle on a large ring gear. The layout of the equipment in Figure 5.6 is typical of what you might see.

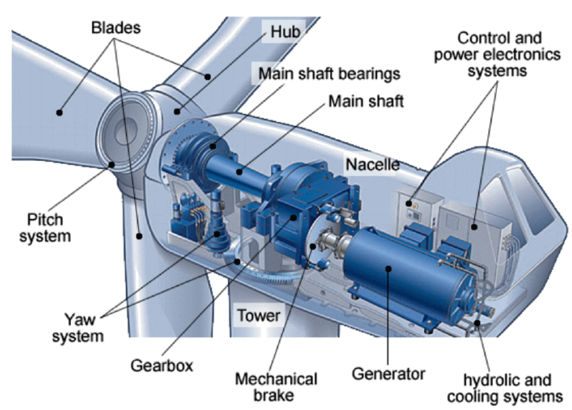

Figure 5.6: A typical layout of equipment inside the nacelle of a modern wind turbine. Source: Tchakoua et al., https://ieeexplore.ieee.org/abstract/document/6618706.

Inside the nacelle it is quite noisy, with air flowing rapidly even though, to keep you safe, your host stopped the turbine before you entered the tower. Towards the back of the nacelle, fans are drawing outside air through a heat exchanger to cool the various electronic and hydraulic systems. Beneath the coolers are the transformer and some power electronics. Moving forward, the next large piece of equipment is the generator that converts the energy in its rotating shaft into electrical energy. The shaft of the generator is connected to a gearbox, which changes the slow rotational speed of the blades into the much higher rotational speed of the generator shaft. Directly in front of the gearbox is the slowly rotating hub to which the blades are attached.

Also within the nacelle is a hydraulic power unit that provides the power to drive the turbine control systems – the yaw motors that head the turbine into the wind control and the motors that control the pitch of the blade (the angle between the blade and the incoming wind). You can see the large bearings at the root of the blades, which change the pitch. Some turbines use electrically driven systems instead of hydraulic systems.

5.3 Back at the Blade on the Ground

You are not able to see inside the blades from the nacelle, but now you have new questions about the blade, so you descend to ground level, go outside, and revisit the spare turbine blade resting on a mount (recall Figure 5.2). Down the surface of the blade, you see small metal fasteners attached periodically to a conductive pathway that extends through the blade all the way back to the hub and eventually to the nacelle. This is a protection system in the event that a turbine blade is struck by lightning, a risk for any large object standing in a wide open space. Damage to a turbine blade from a lightning strike can result in a costly repair.

Standing at the root of the blade, you note that the diameter is so large that you could stand upright inside. You can look inside, down the length of the blade, because it is open at the root (see Figure 5.7, top panel). You see a hollow construction with two spars, called shear webs, which connect the top and bottom surfaces of the blade. These make a three-compartment structure, which is very strong and lightweight.

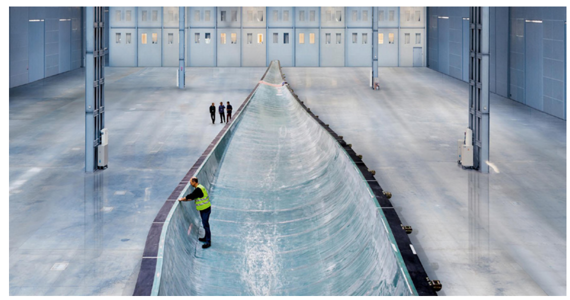

Figure 5.7: ***Top:*** **Looking inside a turbine blade (the BP Sherbino Mesa II turbine blade) at its two shear webs and its balsa/fiberglass construction. Photo: Greg Davies.** ***Bottom:*** **A facility for the fabrication of turbine blades, showing the mold for a half-blade. Photo: Siemens AG, Munich/Berlin, https://www.siemens.com/press/en/presspicture/?press=/en/presspicture/pictures-photonews/2012/pn201204.php.**

The choice of materials for the blade may surprise you: it is a mixture of balsa wood and resin-impregnated fiberglass. (Some large modern blades also incorporate carbon fibers, although this adds to cost.) These materials are preferred over steel and aluminum because of their lighter weight, because they can be shaped into complex forms at less cost, and because they are better at withstanding fatigue. Figure 5.7, bottom panel, shows a typical fabrication facility for turbine blades. Two halves are made separately, and each half is fabricated as a single piece from root to tip. The manufacturing method enables the combination of complex aerodynamic shape and blade strength.

Your tour is over.