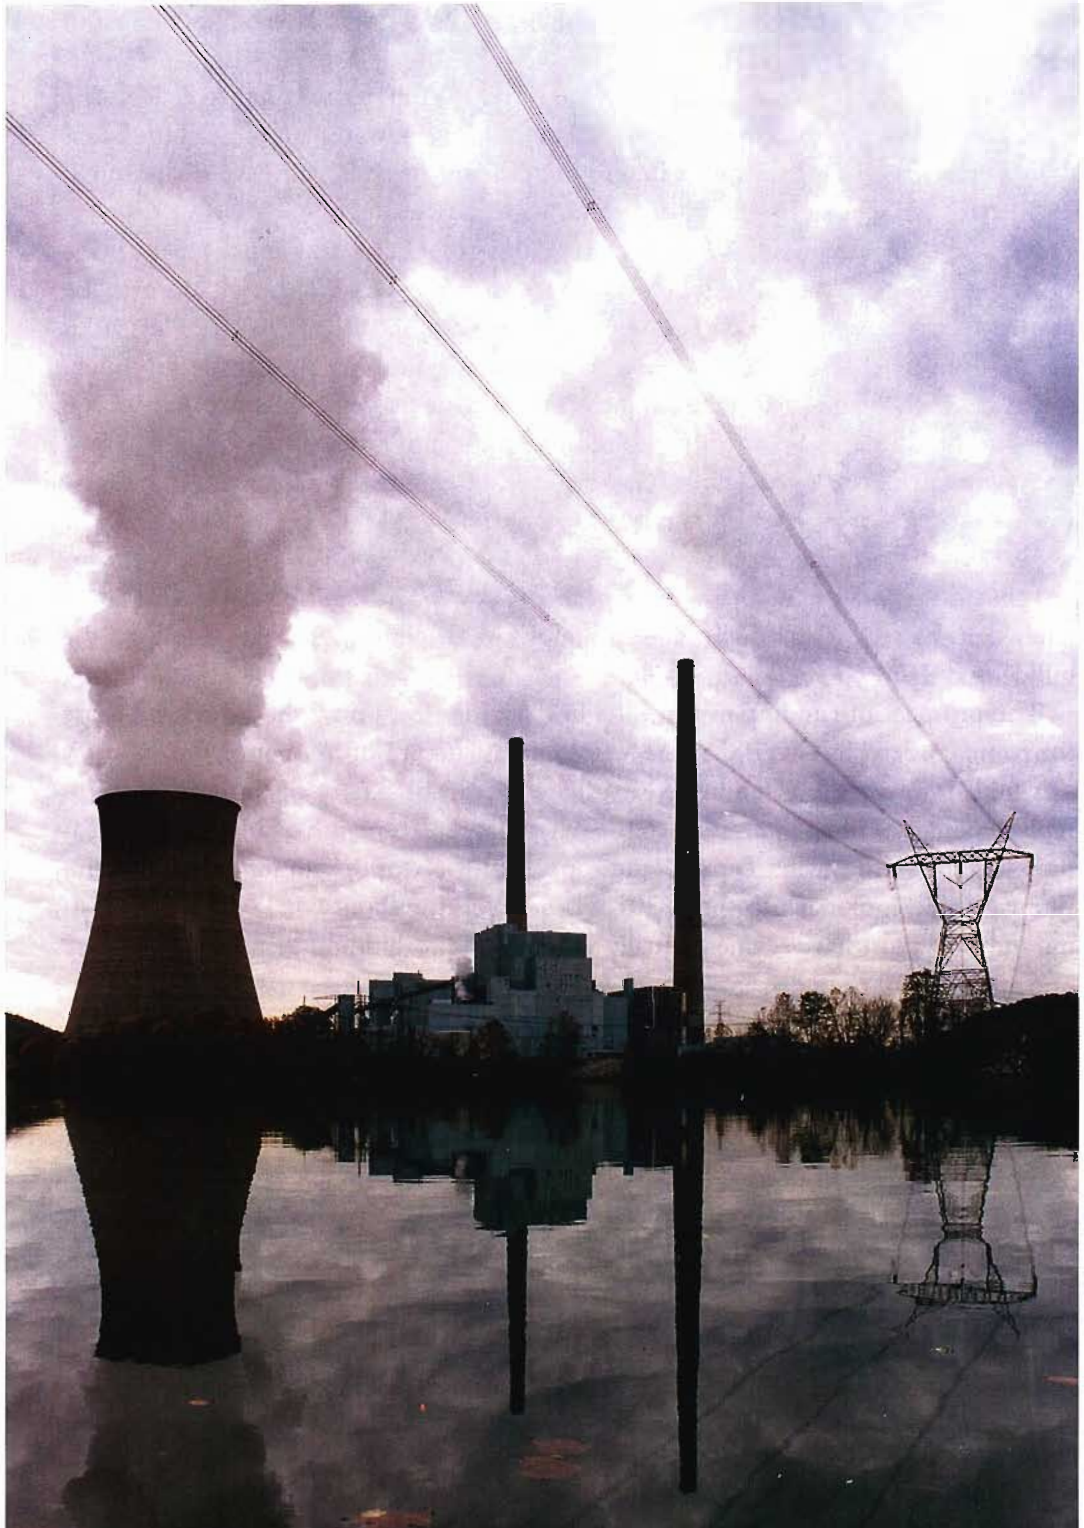

C H A P T E R 5

POWER PLANTS

HEN HENRY ADAMS, the one-of-a-kind American historian, wanted to contrast the medieval with the modern, the emblems he chose for those two eras were the Virgin and the Dynamo. The earlier age, he said, expressed its highest aspirations in building cathedrals consecrated to a spiritual ideal; in our time we exalt the generation of electricity—another invisible but powerful essence. The votaries of the electric cathedral certainly agree with this assessment. A textbook for power-plant operators says of the generating station, "It is like a shrine or source of unfailing light which must be given ceaseless attendance once it is brought into being."

Adams encountered the dynamo at the Universal Exposition in Paris in 1900. A century later the electric power plant is no longer a novel piece of machinery that can attract a crowd at a fair, but it is more central than ever to daily life. And the dynamo (or generator, or alternator—they are all terms for the same machine) still seems an apt symbol of both the hopes and fears invested in industrial progress.

Electricity has become the standard currency of the energy economy. It is not in itself a natural resource that you can dig out of the ground or pump from a well, but other forms of energy are converted into electricity for convenience of distribution and use, just as the body converts a variety of foods into a few simple sugars that circulate to all the tissues. Thus, the power plant doesn't create energy; it merely transforms it. The chemical energy locked up in coal, for example, is captured in the heat and pressure of steam, then passed on to the kinetic energy of a spinning turbine, and finally converted into electric current in the generator.

Three kinds of power plants are scattered around the American landscape. Fossil-fuel plants, which burn coal, oil, or natural gas, make up almost two-thirds of the nation's generating capacity. Nuclear plants tap energy from the disintegration of uranium atoms. Hydroelectric plants are found only where the water is—or more specifically where the water runs downhill.

The John E. Amos power plant *(opposite page)* presides over a moody moment on the banks of the Kanawha River, a few miles from Charleston, West Virginia. From left to right the major structures are cooling towers; the blue-clad buildings that house boilers, turbines, generators, and other machinery; two smokestacks; and a tower for the transmission line that delivers electric power. The coal-fired plant, operated by American Electric Power Company, has three units, each with its own boiler and generator. The total capacity is 2,900 megawatts, which makes the plant one of the 10 largest in the United States.

A few other energy sources are also squeezed to produce the juice of electricity. Their contributions are smaller, but the machinery is no less interesting to look at. Wind power, in particular, creates haunting landscapes. This chapter discusses only generating plants; the transmission and distribution lines that carry electricity to the consumer are the subject of Chapter 6.

FOSSIL-FUEL POWER PLANTS

A big coal-fired power plant going full blast burns a thousand tons of coal an hour, or 30,000 pounds of coal a minute. It generates a billion watts of electricity—in round numbers, enough for a million households. These inputs and outputs are conspicuous features of the plant. On one side you'll see long trains of hopper cars unloading coal, which is heaped up in mountainous black stockpiles. On the other side is a high-voltage switchyard, with electrical transmission lines disappearing over the horizon. Another "output" of the plant is also obvious: a tall smokestack, surrounded at its base by pollution-control equipment.

What's not so easy to see—at least from the outside—is everything that happens between the burning of the coal and the generation of the electricity. The paragraphs that follow describe a power plant in terms of three major flows. First, we follow the fuel from the coal pile through the furnace to the exhaust plume and the ash pit. Then, we trace the circulation of water and steam. Finally, we consider the flow of electricity from the generator into the transmission network.

Coal Handling. When you need to move coal at a rate of 30,000 pounds per minute, it's not done with shovels, or even with bulldozers and trucks. Large power plants all have conveyor systems; one of the most characteristic sights at a coal-fired plant is a conveyor slanting across the skyline. Closely related to conveyors are the stackers that

GETTING A LOOK

Few power plants today are ready to welcome the casual visitor, but some will accommodate you if you call ahead. For a long time, nuclear power stations were more readily accessible than fossil-fuel plants. Companies in the nuclear industry were eager to promote and explain their technology, and so they offered tours on a regular schedule; some plants had educational exhibits and even souvenir shops. All that changed in the aftermath of September 11, 2001. Security is now very tight at nuclear power plants, and visitor programs have been shut down, at least for the time being.

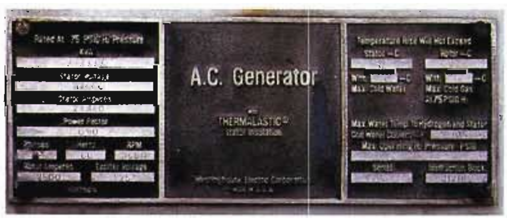

Hydroelectric power stations are more likely to be open to visitors, especially plants owned by public agencies. Similarly, the operators of wind power and solar power installations are accustomed to getting queries from the public and requests for tours.

In areas where the climate isn't too severe, many coal-fired power plants are built on an open, unsheathed framework, so that the furnace and boiler are visible. Even when the boiler is enclosed, the coal-handling equipment is exposed, and so are smokestacks, cooling towers and some pollution-control devices.

The Homer City power plant, 40 miles east of Pittsburgh, is a startling apparition when it suddenly comes into view as you come over the crest of a hill. The buildings housing the boilers and generators are dwarfed and partly hidden by ancillary equipment: three smokestacks in the center, three cooling towers at right, and a series of storage bunkers and conveyors for coal, receding into the distance at the left. The plant has its own coal-cleaning facilities. It was designed to burn coal mined in the immediate vicinity but now gets some of its fuel from more distant sources. Owned by a subsidiary of Southern California Edison, the plant's three units can produce more than 1,800 megawatts of electricity.

pile coal into storage heaps and the unstackers that eventually retrieve it from those heaps. These machines have immense pivoting booms that nod up and down and sway from side to side like long-necked dinosaurs ponderously grazing on the coal.

The coal can't just be piled up and forgotten. It needs watching because of a process called weathering, which is just oxidation, or, in other words, slow-motion burning. If the process gets out of hand, the burning will no longer be slow. The largest piles are inspected daily for hot spots. A worker probes the coal with a long steel rod, inserted like a cake tester. If a section of the rod comes out too hot to hold in the bare hand, the pile needs attention.

From the storage pile the coal moves by conveyor to a bunker or silo, and then another conveyor carries it into the plant. You might imagine lumps of coal being shoveled into a furnace by sweating, shirtless men, but that's not how it's done. All the big plants burn pulverized coal, which is blown rather than shoveled into the furnace. Big lumps are broken into smaller lumps by tumbling in a crusher, which looks like the drum of a giant clothes drier. Then the smaller lumps are pulverized in another rotating drum, this one with steel balls inside or hammers on hinges. The result is a powder as fine as beach sand.

Oil-Fired and Gas-Fired Plants. Coal is the fuel of choice at most U.S. power plants, but there are exceptions, especially in the New England states and some areas of the West Coast. The obvious difference at an oil-burning plant is that you'll see a tank farm instead of a coal pile.

The oil burned by utilities is nothing like the light fuel oils used for home heating. You'll have a more accurate impression if you think of roofing tar. Thick, black, and corrosive, it is the "bottom" product of the refinery, in both the literal sense (it

Lights glimmer within the "exoskeleton" of the Mayo plant, operated by Progress Energy near Roxboro, North Carolina. The plant has two coal-fired boilers supplying steam to a single turbine and generator. Most of the photographs on the next 10 pages were made at the Mayo plant.

The coal pile at the Mayo plant could keep the boilers fired for a few weeks if supplies were interrupted. The two conveyors in the background carry coal to and from the stockpile; it is piled up by a device called a stacker, with a swiveling arm. The coal is retrieved from the pile by an underground auger and then is brought into the plant by the upward-sloping conveyor line in the foreground.

comes out of the lowest tap in the distillation tower) and in the sense that it commands the lowest price (see Chapter 4).

The major problem in handling heavy oil is that it won't flow in cold weather; it needs to be melted. To run through the fuel pipeline, the oil has to be above 100 degrees Fahrenheit; to be sprayed as tiny droplets from a burner nozzle, it needs to be about 200 degrees. Thus, steam heating coils are installed in the storage tank. The fuel pipeline may also have steam tracer lines to keep the oil fluid.

Natural gas is a less troublesome fuel and much cleaner, but also more expensive. It is burned mainly in urban power plants where air pollution levels allow no alternative. A utility-scale plant takes its gas not from the municipal gas mains, which operate at low pressure, but directly from high-pressure transmission lines.

A pulverizer on one of the lower levels of the Mayo plant grinds coal to a powder fine enough that it can be transported by blowing it through a conduit. Almost all of the coal-handling equipment inside the plant is tightly sealed, so you can walk through galleries of machinery that process thousands of tons of coal each day and never see any sign of the coal itself.

The Firebox. When you think of a coal-burning furnace, it's natural to imagine a bed of coals glowing on a hearth, but that's the wrong image for an industrial-scale burner. The powdered coal is treated like a fluid, not a solid. It is sprayed into the firebox.

The ideal is to burn it all, every last particle. Complete combustion not only gets full value out of the fuel; equally important, it minimizes waste-disposal and pollution problems. Anything that doesn't burn will eventually have to be hauled away.

The key to full combustion is maintaining the right ratio of fuel to air, and making sure they are mixed thoroughly. In most large furnaces there are two air streams. A big fan blows the powdered coal into the firebox and starts the combustion process. Then an even bigger fan adds secondary air, with an effect somewhat like that of the afterburner on a jet engine. The fans are usually of the centrifugal type, with a snail shape, like the blower in a hand-held hair drier. But the power-plant fans are built on a totally different scale. They are as big as a two-story house, and the duct work that carries their output is big enough to drive a truck through.

Plants of this kind run 24 hours a day, not so much because there's always demand for electricity but because shutting them down and starting them up again takes hours.

The startup process—called lighting off—is not just a matter of striking a match. Burning a special ignition fuel, usually kerosene, warms up the firebox enough to establish a stable flame pattern before the primary fuel is switched on. It's a difficult process to manage. If the flame goes out, the furnace has to be purged with air to get rid of unburned fuel, which otherwise might explode on reignition. Once the furnace is lit, another 10 or 12 hours may pass before the plant comes up to full power.

Flue Gases. Inside the furnace a pillar of fire rises 100 feet or more. Even where the flame zone ends, the gases remain extremely hot—up to 3,000 degrees Fahrenheit. The idea guiding the design of the plant is to let none of this heat go to waste.

The pathway of the combustion gases is arch-shaped: up through the boiler, then horizontally across the top of the furnace building, then partway back down again. All along this route the gases pass through devices that extract heat in various ways. First is the boiler, where—obviously enough—the heat is used to boil water and make steam. Then, near the top of the arch is a superheater, which raises the temperature of the steam far above the boiling point. Farther along, on the downward arc, are reheaters, which pour more heat into steam that has already passed through the first stages of the turbine. Then comes the economizer, where water is preheated on its way to the boiler. And even after all this superheating and reheating and preheating, the flue gases are still not quite done with their day's work. Their last task is to preheat the air that blows the fuel into the firebox.

Air Pollution Control. Once upon a time, the spent flue gases would have gone straight up the smokestack, carrying a substantial load of unpleasantness with them. The problem is not smoke, as it would be from a malfunctioning fireplace. Smoke

Coal burners are lined up along one wall of the firebox. The coal-and-air mixture ascends through silver-colored conduits from the pulverizers three floors below; it passes through a scroll-shaped fan housing and is injected into the furnace. The green pipe carries liquid fuel used only in starting up the boiler.

Air-handling equipment at the Mayo plant includes a maze of immense ducts. The green machine at lower left is a fan with intake filters that reach a height of about 20 feet. This "forced-draft" fan pushes air into the furnace. An "induced-draft" fan—the smaller green device at ground level right of center—sucks air and flue gases out of the furnace. The effects of the two fans are balanced in such a way that the pressure in the firebox is slightly negative; thus, any small leaks draw air inward instead of spewing fumes outward. The path through the maze of ducts proceeds upward from the forced-draft fan, through a regenerator, then to the left and downward into the boiler building. Flue gases emerge from the furnace in the uppermost duct. They can make a circuit through a selective catalytic reactor at the upper right, or they can bypass this device and descend directly through the regenerator to the induced-draft fan. From there they flow on to the stack, just visible in the background. The regenerator, which sits at the intersection of the two streams of gases, is an air-to-air heat exchanger. The housing, just above the forced-draft fan, is an octagon; inside, a honeycomb disk some 30 feet in diameter rotates slowly. The disk absorbs heat from the exhaust gases and gives it up to the combustion air.

comes mainly from incomplete burning of carbon, which isn't tolerated in utility-scale power plants. But coal is not pure carbon. For one thing, it includes a mineral residue that just won't burn. The heavier part of this residue, called bottom ash, winds up in a pan at the bottom of the firebox. The lighter part is fly ash, and it is carried along by the flue gases as a fine gray dust. Coal also includes at least a little sulfur, which burns to produce sulfur dioxide, the precursor of sulfuric acid and acid rain.

Power plants today are required to capture nearly all the fly ash before it escapes up the stack. The two main technologies are baghouses and electrostatic precipitators. The baghouse is just a filter. The bags are long and thin—maybe six inches in diameter and 20 feet long. Each bag is closed at the top but open at the bottom, where the dirty gas flows in to keep the bags inflated. Gases pass through the fabric, leaving the ash behind as a dust cake on the inner surface. As with a vacuum cleaner, the bags need to be emptied from time to time. In most plants this is done by briefly reversing the flow of air, driving the dust out of the bag and down into a hopper; some units also have a shaker that thrashes the bag back and forth.

Electrostatic precipitators rely on subtler physics than a vacuum-cleaner bag: they work on static cling, the force that makes a toy balloon stick to the wall after you rub it on your clothes, and that sometimes makes your clothes stick to you. Inside the precipitator are many parallel rows of vertical metal plates, with flue gas flowing horizontally through the lanes between them. Hanging down into the spaces between plates are fine wires energized with several thousand volts of electricity. Electrons are repelled from the wires and flee to the metal plates; along the way they attach themselves to passing particles of fly ash, which then stick to the plates. A "rapper" mechanism shakes the collected dust loose, and it falls into a hopper below.

The precipitator may sound like an exotic piece of machinery, but electrostatic air cleaners for the home work the same way. Laser printers and photocopiers are also based on the same principle of lending an electric charge to fine particles.

Neither baghouses nor precipitators can capture the sulfur dioxide in the flue gases. That's a job for a scrubber, which relies on chemistry rather than physics. The scrubber sprays the flue gases with a slurry of lime; sulfur dioxide combines with the lime to form calcium sulfate, or gypsum. The scrubber is a set of large cylindrical vessels (typically 60 feet high and 20 feet in diameter) where the flue gases enter at the bottom and travel upward through the descending mist of lime slurry. The scrubber comes last in line in the processing of the flue gas, after the precipitators or baghouses. A telltale sign of a scrubber in operation is pure white steam pouring out of the stack in the summer. (In cold weather the flue gases may form a visible vapor trail even without the moisture added by scrubbing.)

Not all plants have scrubbers. Some utilities have been able to meet air-quality standards without a scrubber by burning low-sulfur coal (much of it from the Powder River Basin in Wyoming). More controversially, a number of older plants are "grandfathered," or exempt from regulations enacted after the plants were built.

In recent years yet another pollution-abatement technology has begun to appear at coal-fired power plants. Selective catalytic reduction treats the flue gases with ammonia to deal with oxides of nitrogen (usually denoted NO_x). The NO_x is created when nitrogen and oxygen in the air combine in the high-temperature flame of the furnace. Ammonia reacts with the NO_x to form nitrogen and water.

The Stack. Even from miles away, the stack of a power plant is an impressive structure. Up close, what seems most remarkable is not the height but the girth and the bulk.

"Rappers" on the roof of the electrostatic precipitator knock the accumulated dust free, letting it fall into the storage hopper. Each rapper is the size and shape of a baseball bat. Inside is an electromagnet that pulls a steel plunger upward, then allows it to fall again, producing a sharp knock. The rappers are energized at seemingly random intervals, producing a haunting, syncopated music. (The rhythm seemed more modern jazz than rap.)

Electrostatic precipitators at Mayo remove fine particles of ash from the flue gases. The gases enter through the vertical duct at right, pass through the precipitators from right to left and exit downward through the vertical duct at left. Inside the four-story-tall precipitators are electrically charged plates and wires, which trap the particles and deposit them in hoppers below.

The function of the stack has changed over the years. Early industrial smokestacks worked much like a fireplace chimney: buoyant, warm air rose through the stack, creating a natural draft to draw fresh air into the furnace. The taller the stack, the more powerful the natural draft. Today, with a long twisting pathway from the firebox to the stack, natural draft is not nearly enough to keep the air moving, and large fans force the flue gases into the base of the stack. The height of the stack is determined not by the need for natural draft but by the requirement to release effluents high in the air, where they will be diluted and dispersed.

A plant that has three or four stacks has a corresponding number of firebox-boiler-turbine units. But a single stack doesn't necessarily mean that a plant has only one unit. Sometimes the stack is divided internally into multiple flues. In addition to the gigantic main stacks, most plants have a few smaller, stubbier stacks as well. They serve small auxiliary boilers. What looks like a small smokestack may also be a steam vent.

A scrubber in action at the Homer City power plant produces a characteristic white plume, which consists of water vapor from the scrubber condensing as it cools. Only one of the three units at Homer City is equipped with a scrubber, and so only one of the stacks emits a plume. Utility companies occasionally get complaints about these conspicuous emissions, but it is the stacks *without* a visible plume that are probably releasing greater quantities of sulfur.

The Boiler. The boiler of a toy steam engine is a little steel tank with a fire under it. But outside of toyland, a boiler is all about tubes, not tanks.

The power-plant boiler is one of those artifacts whose evolution is so complicated that you can't really understand how it works without also knowing where it came from. The distant ancestor of the modern boiler had a firebox of brick, with a bundle of water-filled steel tubes running through the middle of the combustion zone. Flames swirled among the tubes, boiling some of the water. As boilers got bigger and hotter, the firebrick lining became a factor limiting performance. No material could withstand the heat of the furnace for very long. As a strategy for cooling the brick, some of the boiler tubes were run vertically down the inner face of the furnace walls. Today, boiler tubes have been removed entirely from the interior volume of the combustion chamber; all the tubes are installed in "waterwalls" that line the furnace. The tubes are closely spaced and welded together with a webbing of steel to form solid, airtight panels. The brick that the waterwalls were once protecting has now been eliminated altogether.

Considering the inferno inside the firebox, the environment surrounding a big furnace is surprisingly benign. The waterwalls are swaddled in insulation. You can get up close to them, standing inches away from a 2,000-degree torch roaring 10 or 15 stories tall. The space is warm, but not uncomfortable.

The tubes in the waterwall are called risers, because heated water and steam rise through them. Another set of tubes, called downcomers, carry water flowing in the opposite direction. Risers and downcomers are joined at both the bottom and the top of the boiler to form a continuous loop. The entire assembly of boiler tubes, weighing hundreds of tons, hangs from the roof of the building, with no rigid supports underneath. This arrangement allows for expansion and contraction—the length may change by a foot or more—as the boiler heats up and cools down.

At the very bottom of the boiler, where the downcomers and risers meet, some plants have a mud drum, a cylindrical vessel that takes its name from what you find

A boiler hangs from the roof, allowing it to expand and contract with changes in temperature. The full weight of the boiler is supported by a forest of steel rods that hang from the uppermost girders of the frame and connect to the roof of the boiler structure. The cuplike objects at the very top of the frame are steam vents used for either routine or emergency releases. The cups are filled with baffles meant to reduce the noise level of high-pressure releases.

inside when the boiler is shut down for an overhaul. All the rust and scale accumulate there. The corresponding structure at the top of the boiler is the steam drum, which in fact is more than a drum; it's a complicated piece of apparatus, with lots of machinery inside. The main business of the steam drum is to separate steam from water so the vapor can be drawn off and piped to the turbine, while the liquid is recirculated through the downcomers. To those of us whose experience of steam comes mainly from teakettles, separating steam from water seems easy: the vapor just wafts off the top. But under the conditions inside a power-plant boiler—temperature 675 degrees Fahrenheit, pressure 2,600 pounds per square inch—steam and water are hard to tell apart. It takes a kind of centrifuge, or cyclone, to separate them.

A feature of all boilers, required by law, engineering codes, and insurance regulations, is a pressure-relief valve. Early in the age of steam, boiler explosions were a notorious technological hazard. Both railroad locomotives and the stationary boilers of factory steam plants were blowing up with enough regularity to inspire public dread comparable to modern worries about nuclear-power accidents. Those at greatest risk were the engineers who tended the boilers, and yet they resisted regulation of their work. Nevertheless, a safety measure was imposed and remains universal today: every boiler has a valve that automatically vents off steam at some preset pressure not too far above the normal working pressure. The valve is installed on the steam drum. It relies on the simplest kind of spring-loaded mechanism, which pops open if the internal pressure ever exceeds the strength of the spring. Modern boilers are equipped with other valves, tied into the central computer control system, that allow finer regulation of pressure. But the mechanical safety valve is there in case the computer ever crashes or someone falls asleep at the switch.

Downcomers form part of a loop of piping that allows water to circulate through the boiler. The photograph looks upward along one side of the boiler. Behind the metal sheathing and a layer of insulation is the water-wall, where pipes called risers are heated by radiant energy from the furnace. Some of the water in the risers turns to steam, but the rest is recirculated through the downcomers.

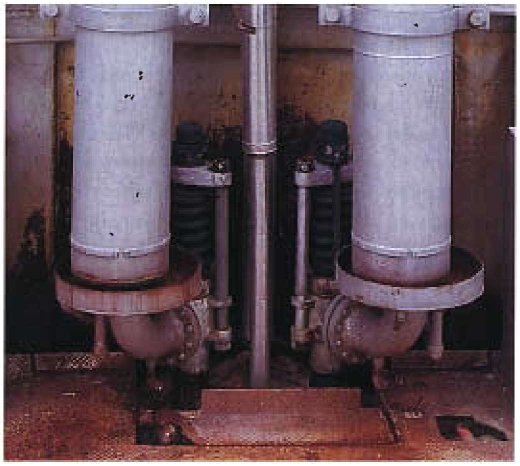

The steam drum *(below)* is where risers and downcomers join at the top of the boiler. Visible here is just one end of a long drum, filled with machinery for separating water and steam. Pressure-relief valves *(above)* are on the roof just above the steam drum. They are simple mechanical valves, which discharge steam whenever the pressure in the boiler is great enough to compress a spring that holds the valve closed.

The Turbine. The steam-driven turbine is a close cousin of a jet engine. In both machines, a hot, high-pressure gas spins a series of fanlike turbine wheels. In the process, the gas expands and cools. It's a simple idea, but power-plant turbines rated at a billion watts of mechanical power are not simple to build or operate.

The turbines and the generators they drive are the most expensive hardware in a power plant. Often, they are built atop their own special concrete-and-steel foundation, separate from the rest of the plant. This is done to control vibrations in the rotating machinery and to maintain precise alignment in the bearings that support the long, spinning steel shaft that runs through both the turbine and the generator.

A single stage of the turbine consists of a stator wheel (which is rigidly fixed to the frame of the turbine) and a rotor wheel (attached to the rotating shaft). Steam is steered through vanes in the stator and then passes through the blades of the rotor, turning it by the same principles that run a windmill or a waterwheel. The blades and vanes have graceful airfoil shapes, and they are carefully machined from fancy steel alloys that can withstand extremes of temperature, pressure, and mechanical stress, as well as a corrosive environment. A single blade breaking off would destroy the entire machine as the debris crunched through the downstream rotors and stators.

Typically a turbine has three units, all mounted on the same shaft. Steam straight from the boiler and superheater is fed into the high-pressure turbine, where it expands and cools somewhat. The steam then goes back to a reheater unit, where its temperature comes back up to about 1,000 degrees Fahrenheit, although the pressure is not restored to its original level. This warmed-over steam then goes through the intermediate-pressure turbine, where again it expands and cools. Finally, the steam passes into the low-pressure turbine. Note that the same quantity of steam—the same mass of water molecules—goes through all three turbines, but because the pressure drops in each unit, the volume of steam increases. As a result, the intermediate-pressure turbine has to be bigger than the high-pressure unit, and the low-pressure turbine is the largest of all. Judging from their relative sizes, you might guess that the big low-pressure turbine is doing most of the work, but the truth is just the opposite. The little high-pressure turbine puts out 60 percent of the total horsepower, and the massive low-pressure unit supplies only about 15 percent.

Turbines are so large and complex that their most mundane auxiliary equipment is more imposing than any of the machines most of us meet in everyday life. Pumps and motors larger than an automobile engine are needed just to keep the turbine supplied with lubricating oil. The bearings and seals along the main shaft also require large accessory pumps.

Another vital auxiliary is the governor that regulates turbine speed. The classic speed-control mechanism is the flyball governor, which became an icon of the industrial age and a textbook example of the concept of feedback control. The governor has two weights (the flyballs) attached to hinged arms that spin around a vertical shaft at the same speed as the turbine. As the shaft turns faster, the balls are flung outward, and the hinged arms are lifted up. A linkage attached to the flyball arms then closes the

The high-pressure turbine is the smallest but most powerful of three turbine units at the Mayo plant. The turbine itself is hidden under a thick blanket of insulation.

steam valve a little, slowing the turbine and allowing the flyballs to sink back toward their resting position. In this way the turbine is slowed every time it tries to speed up and is sped up every time it tries to slow down, so a steady speed is maintained. The flyball governor was the world's first version of cruise control. The feedback principle is still at the heart of turbine control, although now it's all done by computer.

The Condenser and Feedwater System. To make a turbine spin, it's not enough to push steam into the inlet port; you also have to let it out at the exhaust port. Lowering the pressure and temperature at the outlet is the job of the condenser, where the steam gives up the last of its heat. As the steam condenses, its volume is greatly diminished, and so the pressure falls too. Indeed, the pressure in the condenser is less than atmospheric; there's a partial vacuum, which actually sucks steam out of the turbine.

The water that collects in the bottom of the condenser is distilled water, which is usually considered the ultimate standard of purity. But the water needs further treatment, called polishing, before it can be returned to the boiler. Any minerals deposited inside the tubes of the boiler would clog up the arteries and could cause a dramatic kind of heart attack: the deposits would act as an insulating blanket, allowing the metal wall of the tube to overheat. If a tube splits open, everyone within a few miles of the plant hears it.

To remove suspended solids, the water is filtered through sand or charcoal, and magnetic separators extract particles of rust. An ion-exchange column works just like a residential water softener to eliminate troublesome magnesium and calcium compounds. Other chemical treatments adjust the pH—the acidity or alkalinity—and remove dissolved oxygen, which can attack metals.

Governors and throttle valves mounted on the high-pressure turbine control the speed and the power output of the entire turbine-generator unit.

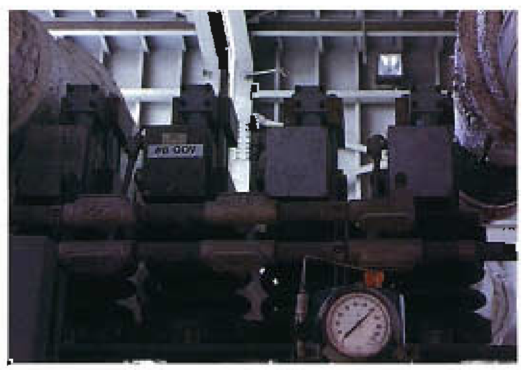

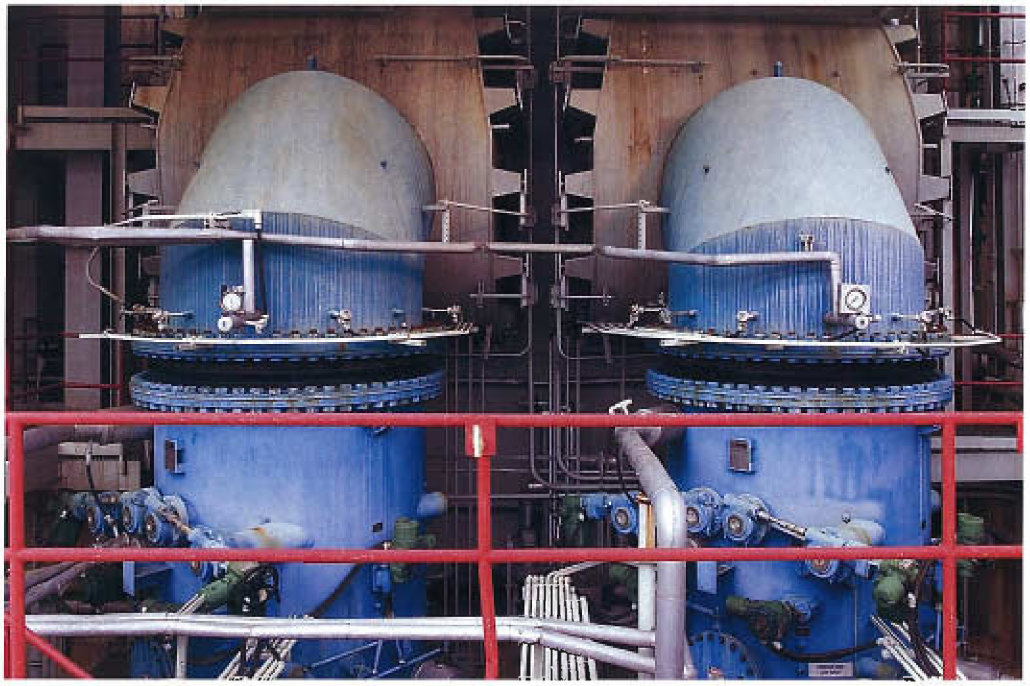

The condensers are among the largest devices in the plant, because they receive steam at its lowest pressure, when it has fully expanded. Here only the end caps of two condenser units are visible; they are the gray metal structures with flanges that look somewhat like gear teeth along their side walls. The giant blue ducts carry spent steam.

The treated condensate, as well as fresh "makeup water," goes back around the loop to be boiled into steam again. Pushing the feedwater into the boiler is not easy, however; the feedwater pump has to overcome the entire head of steam pressure. The feedwater pump is the biggest of the many pumps in a power plant. It is usually mounted on the floor below the turbine and generators and is driven either by a very large electric motor or by a steam turbine of its own. In either case the pump consumes 2 or 3 percent of the raw power output of the plant.

Generators. The ultimate purpose of everything in the power plant, from the coal pile through the furnace and the boiler to the turbine, is to spin the shaft of the generator and create an electric current. Generators rely on an effect called electromagnetic induction, discovered 150 years ago by Michael Faraday and Joseph Henry. Induction creates a voltage in a loop of wire whenever a magnetic field moves through the loop. In a generator, the magnet that creates the field is on the spinning rotor; the loop of wire is wound on the unmoving stator that surrounds the rotor.

In the type of generator used with a steam-driven turbine, each turn of the magnetic rotor produces one cycle of alternating current (AC) in the stator coils. The current flows first one way through the coil and then the other, like a tide sloshing in and out. A rotor turning at 60 revolutions per second generates alternating current at a frequency of 60 cycles per second, or 60 hertz. This speed—usually stated as 3,600 revolutions per minute, which amounts to the same thing—is the standard throughout North America. Every generator connected to the U.S. power grid is adjusted to this rate of rotation. In Europe similar generators turn at 3,000 revolutions per minute, producing power at a 50-hertz frequency.

The feedwater pump is mounted at the low point of the plant in terms of elevation, but it is the point of highest pressure in the entire boiler system. The pump has to attain this pressure in order to force water into the boiler against the head of steam.

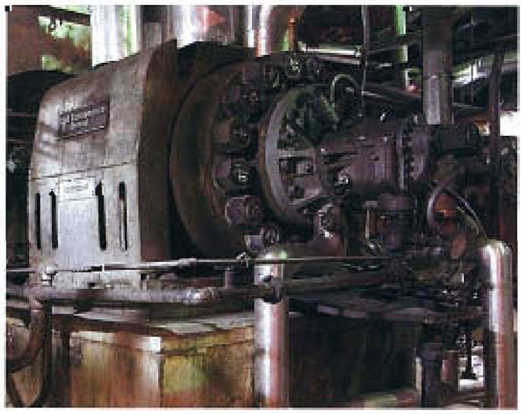

Here's something to puzzle over: to generate an electric current, you need a strong magnetic field, but to create a strong magnetic field, you need an electric current (because the magnet on the rotor is an *electro*magnet, a big coil of wire with a current flowing through it). Where does the current for the rotor magnet come from? The answer is that it comes from another, smaller generator called the exciter, mounted on the same shaft as the main generator. But the exciter also has a rotor magnet that needs a current—where does *that* come from? It comes from an even smaller generator, the pilot exciter. At this point it sounds like we're going to have an infinite regress of smaller and smaller generators, but in fact there's a stop to it. The pilot exciter is small enough to operate with a permanent-magnet rotor. In this way the many megawatts of the main generator are bootstrapped from the feeble stirrings of permanent magnets like the ones that hold notes on your refrigerator.

Big generators are remarkably efficient. Out of all the mechanical energy cranked into turning the shaft, the generator converts between 98 and 99 percent into electricity. But if a generator's output is a billion watts, internal losses of just 1 percent add up to 10 million watts of heat—the equivalent of 10,000 toaster ovens running at the same time. Getting rid of this heat is a major challenge.

The stator windings, where the heaviest currents flow, are water-cooled. The conductors in these coils are hollow copper tubes, and water is pumped through them at high speed. If running water through a high-voltage machine seems contrary to common sense, the cooling medium for the rest of the generator will strike you as even more unlikely. It is hydrogen gas. Hydrogen is chosen because among all gases it is the best possible coolant; the lightweight molecules carry off heat more effectively than those of heavier elements. The lightness also reduces "windage" losses, the energy spent moving the rotor through the atmosphere. But hydrogen has had a reputation for danger ever since the Hindenburg accident in 1937; infusing the flammable gas into a generator full of hot metal and high voltages seems to invite disaster. But hydro-

Pumps and valves below the turbine deck circulate gaseous hydrogen through the generator, cooling the rotor windings.

The single generator for the Mayo plant is rated at more than 750 megawatts. The generator is inside the cylindrical shroud in the foreground; the similarly shaped housing behind it is the low-pressure turbine.

The switchyard of the Gordon Evans Energy Center in Colwich, Kansas, is a thicket of transformers, switches, circuit breakers, lightning arresters, and other high-voltage devices. The main function of the switchyard is to raise the voltage to a level that can be transmitted long distances.

gen burns only in the presence of oxygen; the key to using it safely is to exclude all air. Before the generator is filled with hydrogen, it is purged with carbon dioxide.

Electricity is carried away from the generator on bus bars, heavy copper or aluminum conductors rated to carry as much as 40,000 amperes of current. (The heaviest wires you'll find in your home are limited to 100 or 200 amperes.) The bus bars are as thick as tree limbs, and they may be encased in protective tubes that make them look even thicker. They lead to a switchyard outside the plant.

The Switchyard. Although the generator puts out prodigious currents, the voltage level is only moderate by power-company standards—usually between 10,000 and 30,000 volts. Right outside the wall, a transformer boosts the voltage to a much higher level—often 230,000 or 345,000 volts, and in a few cases as high as 765,000 volts. The high voltage allows the power to be transmitted long distances with relatively little loss along the way.

The transformers and their related switches and circuit breakers are set up in a fenced-off area called the switchyard, which can be as large as the rest of the plant. The devices here are essentially the same as those in the substations at the other end of the transmission line, where the power is brought back down to lower voltages for distribution to neighborhoods. This machinery is discussed in the next chapter.

The switchyard brings power *into* the plant as well as providing a way out. A typical generating station absorbs 4 to 7 percent of its own electrical output for running machinery such as fans and pumps. When the plant is starting up, much of that equipment has to be running before the main turbines and generators are cut in. The start-up power is supplied by other stations on the power grid, and brought in over the same transmission lines that normally export the plant's own output.

What happens if all the power plants in a system are shut down at the same time? Until 1965 plant operators thought they would never have to answer this question because such an event seemed so unlikely. But on November 9, 1965, a blackout in the northeastern United States left some cities dark for more than 12 hours. One reason it took so long to restore power was that generating stations didn't have enough emergency power to restart without help from their neighbors—who were, of course, in the same predicament. The utilities companies promise it won't happen again.

COMBUSTION TURBINES

The demand for electricity fluctuates by the millisecond. When you turn on the coffee pot and the toaster in the morning, a power plant somewhere has to respond by opening the throttle a little. Demand also fluctuates on longer time scales. People use more electricity during the day than at night, and in most places they use more during the summer than the winter. It would be a great convenience to have generating units that could be run only at times of peak demand. Big coal-fired stations are

A power plant in a different architectural style, with all the machinery fully enclosed, the Ravenswood Generating Station is a well-known landmark for New Yorkers; it occupies a conspicuous site in Queens, just across the East River from Manhattan. It is also a landmark for power engineers, the home of a generator known as Big Allis (built by the Allis-Chalmers Corporation), the first generator capable of producing 1,000 megawatts. The plant burns natural gas, with oil as a backup fuel.

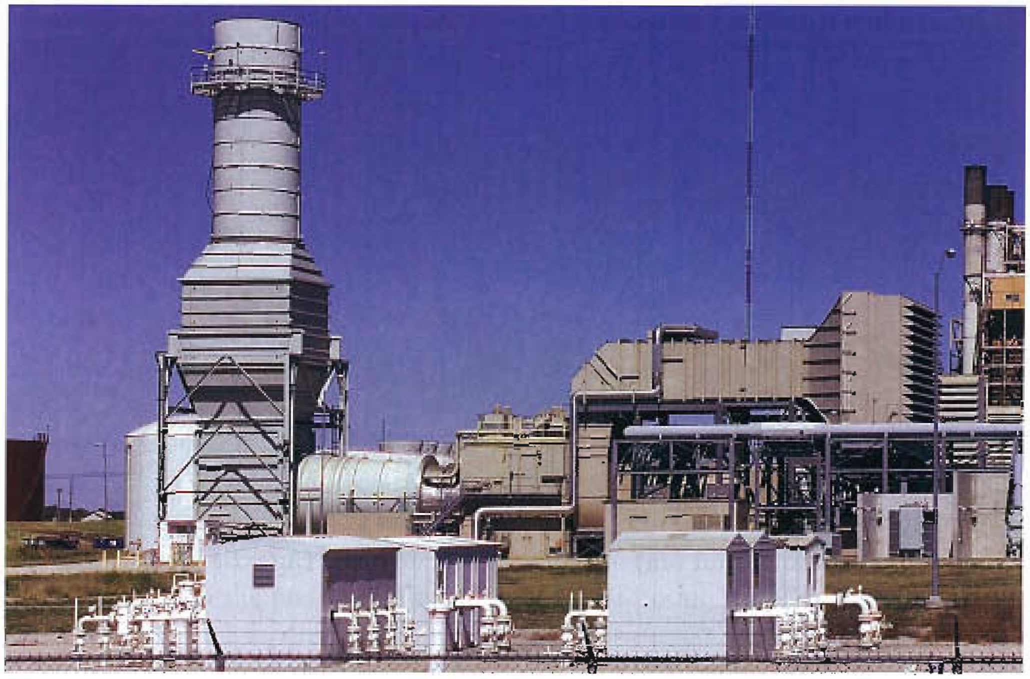

A combustion turbine at the Gordon Evans Energy Center supplements the capacity of the plant's main steam units. The tan duct extending horizontally to the right is the air intake. It arches over the generating unit. The turbine itself is in the square tan enclosure; it is followed by a flared horizontal exhaust duct, and then the large gray exhaust stack, which discharges vertically to reduce noise.

not well suited to this duty, because they take hours to start up and shut down. Hydroelectric plants are much more flexible in this respect, but, on the other hand, you can't build Hoover Dam just anywhere. The solution adopted by many utilities is a machine known as the combustion turbine or gas turbine. Power company employees call them jets, and for good reason: they evolved directly from the engines that power jet aircraft.

A combustion turbine relies on the same physical principle as the steam turbine in a coal-fired plant: hot, high-pressure gases expand against the vanes of a turbine wheel, exerting a force that causes the wheel to spin. But instead of steam, the hot gases are the products of combustion. Fuel and air are mixed, compressed, and ignited inside the turbine, where they expand and thereby turn the rotor vanes. The fuel is usually natural gas. A single combustion turbine has a power output of 10 to 100 megawatts, but it's easy to build clusters of them with larger aggregate power.

Combustion turbines are less efficient than the best steam turbines, but they have compensating advantages. First, of course, they can be started and shut down in a matter of minutes, sometimes just by pushing a button in a distant control room. Furthermore, because they don't require as much land as a full-scale power plant and because they burn cleaner fuel, they can be put closer to cities, which relieves congestion on electric transmission lines. And jets can supply start-up power for larger conventional plants. For this last reason, many steam plants have a few combustion turbines on the site. The standard mode of operation is to keep the big, efficient "base load" plant running all the time, and start the jets only at times of peak demand.

Combustion turbines vary in appearance, but a common feature is that the turbine itself is overshadowed by air intakes and exhaust stacks. One reason for the large intake and exhaust structures is that free flow through the system improves efficiency. But there's another reason: the structures are engineered to suppress noise. If the turbine plant has neighbors, noise is likely to be a major issue. Jet engines are no quieter on the ground than they are on airplanes.

Four new combustion turbines in Kearney, New Jersey, are seen under construction in 2002. The air intakes are atop the buildings that house the turbines.

NUCLEAR POWER PLANTS

It began as the technology of megadeath. Then in the 1950s and 1960s "the peaceful atom" promised a life of ease through boundless energy—electricity so cheap no one would bother metering it. By the 1970s the tide had turned again, and nuclear plants were regarded as a menace, at least in the United States. Who can say where this roller-coaster history will end. At the moment nuclear power looks like a zero-growth industry, but even if no one ever builds another nuclear plant, the existing ones will remain a major part of the energy infrastructure for decades. Almost 100 commercial power reactors supply about 14 percent of the nation's electricity.

A nuclear generating station has much in common with a coal-fired power plant. They are both closed-loop steam cycles. The generators and electrical switchgear are almost interchangeable, and the turbines are very similar. The only important difference lies in how the steam is produced.

The energy source in a nuclear plant is the disintegration of uranium, the heaviest of the naturally occurring chemical elements. In certain uranium atoms, the nucleus—the dense core of protons and neutrons—can spontaneously split in two. The splitting, or fission, is especially likely to happen after a nucleus absorbs an extra neu-

Eight combustion turbines are lined up along the East River north of the Ravenswood plant; there are eight more elsewhere on the site, as well as a new "combined cycle" plant that uses the hot exhaust from a combustion turbine to produce steam. In the units seen here, the brick-red stacks are air intakes, and the taller corrugated metal ones are exhausts. Visible in the background is the Queensborough Bridge.

tron. As the nucleus breaks apart, it gives off a small jolt of energy, which is what ultimately gets turned into electricity, and it also emits a few spare neutrons, which can go on to induce the splitting of other uranium nuclei. In this way a chain reaction gets started. It's just like one of those pyramid marketing schemes, except it works.

To keep the chain reaction going, all that's needed is a sufficient number of susceptible nuclei in a small enough space. It's also essential to control the reaction. Other substances come into play here. Water tends to enhance the reaction because it slows down neutrons, and slow neutrons are more likely to be absorbed by uranium nuclei. Carbon and boron tend to damp out the chain reaction by absorbing neutrons and making them unavailable. All three of these substances have roles in power reactors.

For use as reactor fuel, uranium oxide is molded into cylindrical pellets a third of an inch in diameter. The pellets are very heavy, and always warm with a glow from within. Or so I'm told. I've never held them in my hand. Outside of a few fuel-processing installations, the bare pellets are never seen. They are stacked up inside tubes made of a high-temperature zirconium alloy, and then the tubes are welded shut.

Nuclear power has certainly not made electricity too cheap to meter. The uranium fuel is not free, and the capital costs of building a plant have turned out to be

Two pressurized water reactors stand side by side at the Arkansas Nuclear One site on Lake Dardanelle near Russellville, Arkansas. The domed cylindrical structures are the containment buildings; behind them are the fuel-handling facilities At left in the foreground is part of a large cooling tower. The two units are siblings but not twins; they were built a few years apart from slightly different designs and components.

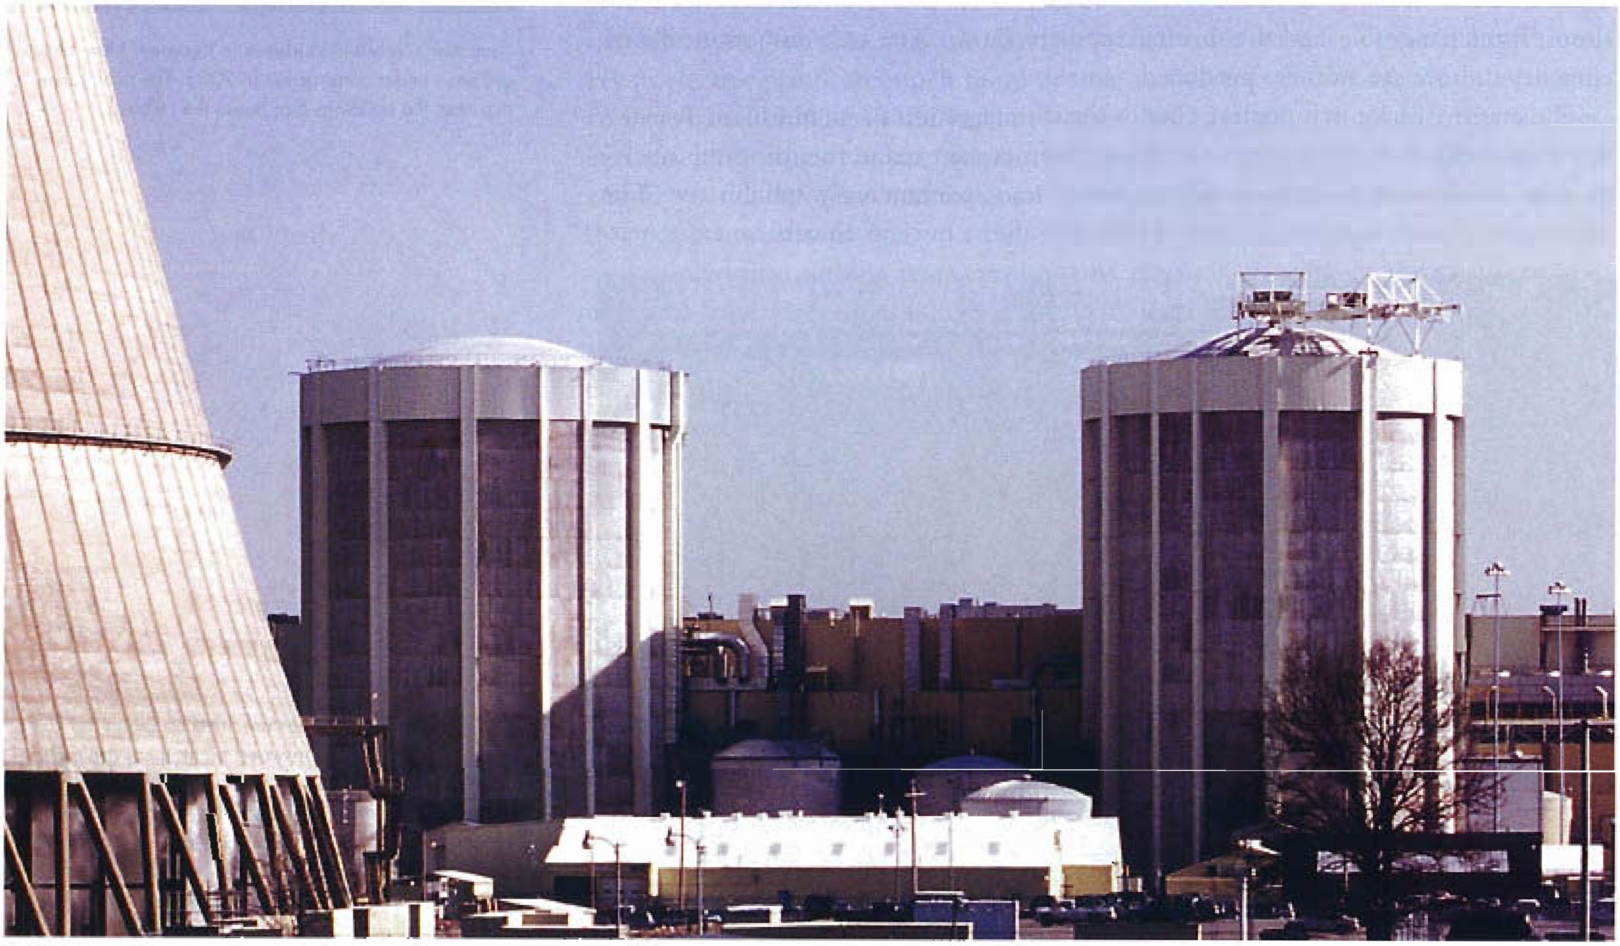

daunting. All nuclear construction is governed by special engineering codes, with elaborate schedules of inspection and maintenance. Every pipe and valve must bear a "Code N" stamp, which raises the price more than gold plating would. Someday, dismantling the plants may wind up costing even more than building them did. And there's also the cost of dealing with radioactive wastes and spent fuel.

Worldwide, there is considerable diversity in the design of nuclear plants, which may signify that engineers have not yet built enough of them to reach consensus on how best to do it. Just two designs dominate in the United States: the pressurized water reactor (PWR) and the boiling water reactor (BWR). Only those types are described here.

The Pressurized Water Reactor. The distant ancestor of the PWR is the U.S. Navy's program to develop nuclear power for ship propulsion. The defining feature is a reactor core fully immersed in liquid water, which is kept under so much pressure that it cannot boil even though the temperature reaches 600 degrees Fahrenheit. About two-thirds of the operating American reactors are PWR types.

The PWR relies on an indirect, two-stage process to drive the turbine and generator. Water heated in the reactor core is pumped to a steam generator, where it heats and boils water in an entirely separate circuit; it is the fluid in this secondary loop that drives the turbine. There is no exchange of fluids between the two loops; this is the safe-sex version of nuclear power. Because the steam that drives the turbine never enters the reactor, the chance of radioactive contamination should be slight.

A PWR has a distinctive profile. The containment building, which houses the reactor, is a tall cylinder with a domed lid. Deep inside is a massive steel pressure vessel, and inside that is the reactor itself. Also in the containment building are the steam generators and pumps to drive the circulation through the primary loop. The pumps stand three stories tall and are powered by electric motors of 4,000 to 7,000 horsepower. Each pump has a flywheel that will keep it running for a few seconds after a power failure—long enough for other emergency cooling systems to kick in.

The reactor vessel is shaped like a medicine capsule standing on end, 40 feet high with steel walls nine inches thick. It weighs close to a million pounds, which means it can only be shipped by barge or rail. (There are no 500-ton highway trucks.) The inner surface is clad with half an inch of stainless steel as a defense against corrosion. And corrosion is a serious worry. In 2002 a work crew at the Davis-Besse nuclear plant near Toledo, Ohio, discovered a spot on the lid of the reactor vessel where acid had eaten away the entire thickness of the wall except for the stainless-steel cladding.

Within the reactor itself, several thousand fuel rods are packed into a volume about the size of a high-ceilinged bathroom. There are also control rods made of boron carbide—a compound of two neutron "poisons," or absorbers. With all the control rods in place, neutrons are blotted up quickly enough that a chain reaction can't sustain itself. The control rods are lifted out through the top of the pressure vessel to start the nuclear reaction. In the event of a power failure or some other malfunction, the rods fall back into place by gravity.

Adjacent to the containment building is the fuel-handling building—often larger in volume though less distinctive in shape. Here fuel-rod assemblies are stored, both fresh ones awaiting installation and depleted ones removed during refueling. The fuel rods can't simply be stacked on a shelf. They are kept immersed in water both as a radiation shield and as a coolant. The deep vat of water, known as the swimming

THREE MILE ISLAND

Sometimes an industrial accident seems to have the fatal momentum of a Greek tragedy. Terrible things keep happening, but nobody understands why until it's too late.

March 28, 1979, was a bad day on Three Mile Island, in the Susquehanna River south of Harrisburg, Pennsylvania. In the small hours of the morning, a shift foreman and two other workers were doing routine maintenance in one of the two nuclear power plants built side by side on the island. Both plants are pressurized water reactors. The maintenance work was in what would seem to be a noncritical section of the plant—the polishers that remove minerals from feedwater in the secondary cooling loop. But events in that obscure corner of the plant had consequences the whole country soon heard about.

The work crew was blowing compressed air into one of the polishers, and apparently the pressure drove water into an instrument air line, one of many small pneumatic tubes used for sensing and controlling conditions in the plant. The clogging of this particular air line had the effect of closing valves that controlled the flow of feedwater through the polishers. With the supply of water cut off, the main feedwater pumps shut down automatically (or, in power-plant argot, "tripped"). Three emergency feedwater pumps immediately started up, but they were unable to deliver any water because another pair of valves had mistakenly been left closed. The improper position of these valves was discovered and corrected eight minutes later, but by then a great deal else had happened.

Less than a second after the main feedwater pumps tripped, the turbine and generator tripped in turn. In the next three seconds, the pressure within the reactor and the primary coolant system rose to 2,255 pounds per square inch, at which point a relief valve opened up, draining steam and water from the reactor vessel into a tank at the bottom of the containment building. After another five seconds, the reactor itself tripped, and control rods were automatically inserted to halt the nuclear reaction.

Although this fast-paced cascade of emergencies sounds quite dire, there was as yet no reason for alarm. Turbine and reactor shutdowns are not routine events, but operators are trained to deal with them. In this case the two operators on duty in the control room immediately set out to perform what seemed to be the most urgent tasks—double-checking the status of the turbine and generator to be cer-

pool, is a beautiful sight. The water is still and clear; technicians test its purity and clarity by reading fine print at the bottom through binoculars. The fuel assemblies hang on racks 25 feet below the surface. The most active elements are enveloped in the haunting blue glow of Cerenkov light, which is emitted when electrons streak through the water faster than the speed of light in water.

tain that these expensive pieces of machinery would not be damaged. The all-knowing chorus in a Greek tragedy might have warned them that bigger worries were looming, but the operators at Three Mile Island did not have the benefit of such a warning.

Over the next two hours the scene in the control room grew more hectic; at one point 60 operators, supervisors, engineers, and others struggled to stabilize the system. The main focus of their attention was maintaining the right water level in the primary cooling system. Most of the time, the level seemed to be too high. A vessel called the pressurizer is supposed to be kept half full of water and half full of steam; the operators thought it was filling with liquid water, which would make it hard to control pressure in the system. Hence, they throttled back emergency systems that were pumping water into the reactor. Actually, the water level in the pressurizer was never too high; it was dangerously low. The operators had been misled by their instruments. The underlying source of the problem was yet another valve malfunction: the pressure-relief valve that had popped open three seconds after the start of the accident should have closed just 10 seconds later, but it remained open, allowing a massive leak. The stuck valve was not discovered until more than two hours later, by which time most of the primary coolant had boiled away.

By now it was too late to avoid serious damage to the reactor. Although inserting the control rods had halted the nuclear chain reaction, radioactive decay was still producing about 30 megawatts of heat, which could not be removed fast enough. Parts of the reactor core crumbled and melted. Also, the overheated zirconium-alloy cladding on the fuel rods reacted with steam to produce hydrogen gas, raising fears that a hydrogen explosion might rupture the containment building. The explosion never came; it turned out there was too little oxygen present to create an explosive mixture. Throughout the accident there were only small releases of radiation.

It took a month to coax the reactor into a safe state, and it took more than 10 years to clean up the mess. Several commissions investigated the accident—Greek choruses chanting of catastrophe after the fact. Factors cited as contributing causes included management policies that allowed the plant to run with emergency feedwater valves closed, operator training that put too much emphasis on one kind of accident and neglected other possible failures, and a reactor design that may have been too skittish for reliable control. Most of all, the investigators criticized the man-machine interface. The operators could easily have averted the damage if only they had known what was happening inside the containment building, but the hundreds of meters and gauges in the control room failed to communicate the information they needed. The indicator for the crucial relief valve showed that it had been ordered to close but did not register its true position.

What lessons should be learned from Three Mile Island? Opinions vary widely. Opponents of nuclear power interpret the accident as a demonstration of just how dangerous and uncontrollable the technology is. Proponents look at the same evidence and argue that the accident shows the inherent safety of nuclear reactors, since just about everything that could have gone wrong did go wrong, and yet there was no serious harm done to public health. Both sides would rather not see any further demonstrations of this kind.

Today the empty shell of the failed reactor still stands on Three Mile Island, next to its older sibling reactor, which was shut down after the accident but was restarted in 1986. (In the photograph on the opposite page, the active unit is on the right, the corpse on the left.) General Public Utilities, the operating company, built a visitor center and souvenir shop, where you could buy Three Mile Island tee-shirts and cookbooks. But at last report the visitor center was closed.

Every year or two, the reactor needs to be refueled. This is not like gassing up the car; it's more like a major engine overhaul, with the added complication that it's done underwater. The first step is to pop open the top of the reactor vessel. The containment building has a crane built in for lifting off this 150,000-pound item. Then a section of the building above the open reactor vessel is flooded to a depth of 15 feet; this pool of water connects via a tunnel with the swimming pool in the fuel-handling building. Old fuel assemblies are pulled out of the reactor and transferred through the water-filled tunnel to the fuel-handling building. New fuel elements come back in through the same tunnel, which you might think of as something like a pass-through between kitchen and dining room.

Another building houses the control room. It's airtight and pressurized (so that any leakage is outward). The ventilation system can close off all air intakes in seconds. The reason for these features is not hard to guess: in the aftermath of an accidental release of radiation, it's helpful if the operators can stay on the job, and survive.

One peculiarity of nuclear power can make the operator's job especially tense. When something goes wrong in a fossil-fuel plant, shutting off the fuel and air puts

A nuclear generating station built around boiling water reactors also has two independent units, but the entire plant fits under one roof. The reactors are in the wing of the building at the left; the fuel-handling facilities are in the lower wing at right. The multicolored forms in the foreground, which look like picnic tents, are actually the roofs of tanks for holding various water-treatment chemicals. The power plant, operated by the Tennessee Valley Authority, is in Browns Ferry, Alabama.

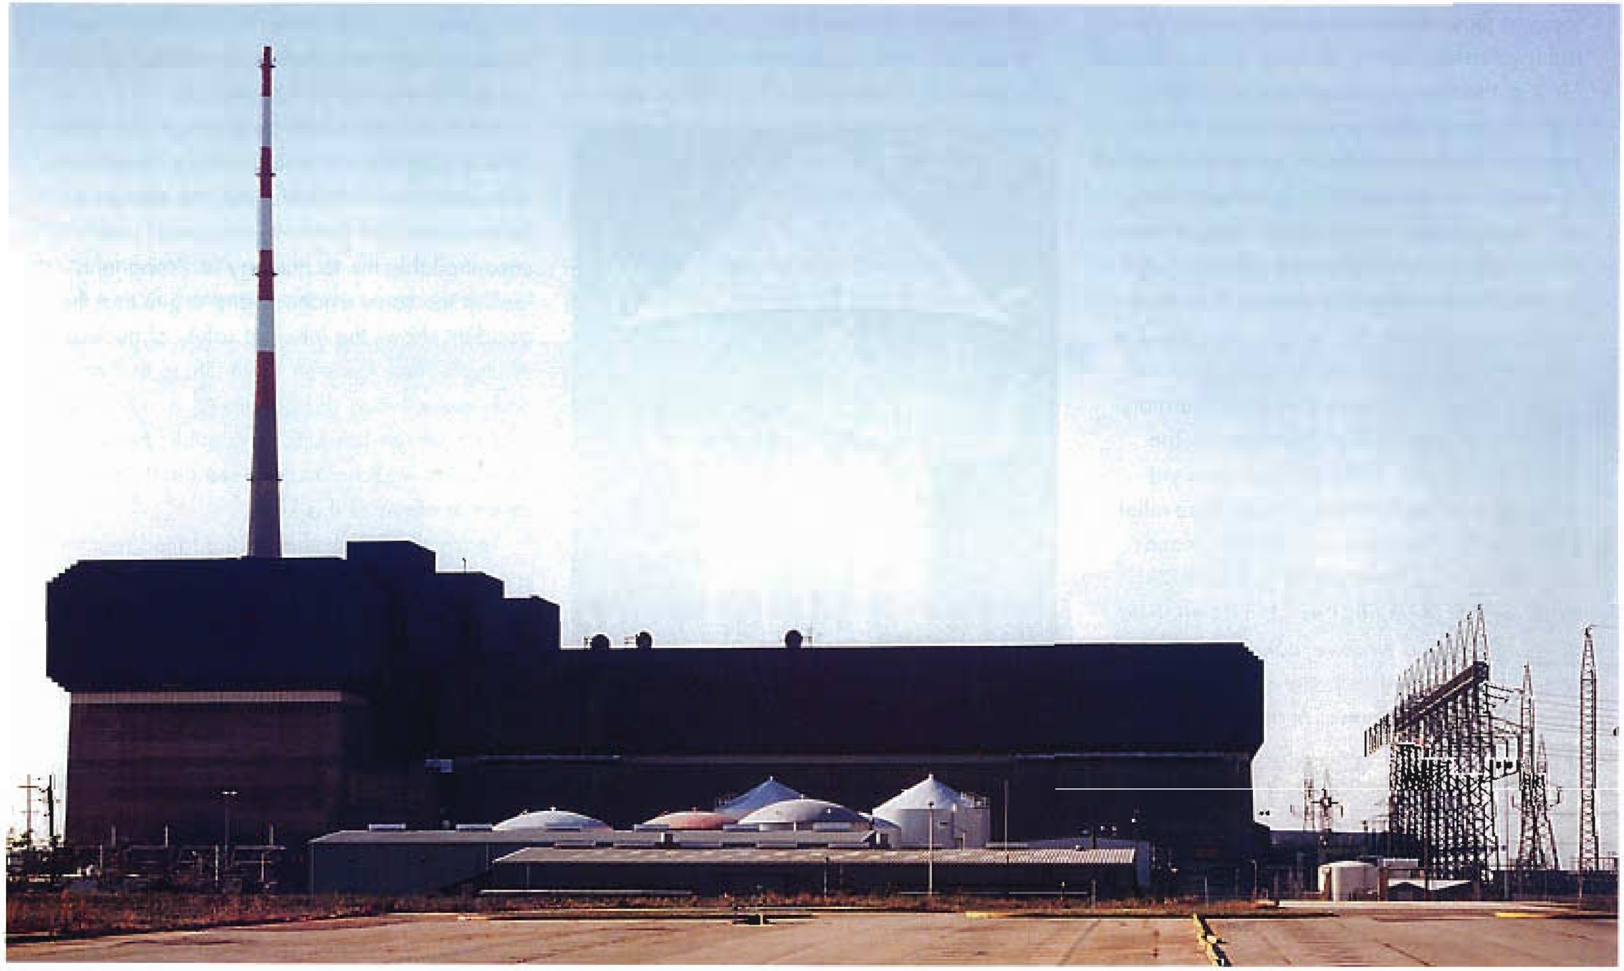

out the fire. With a nuclear reactor, dropping the control rods quenches the nuclear chain reaction, but that's not the end of the story. The fuel elements continue to produce megawatts of heat for hours afterward because of the ongoing radioactive decay of fission products. There's no switch or valve that turns this process off. As a result, the reactor needs a continuous supply of water for cooling even after a shutdown.

The standard operating nightmare for nuclear plants is the dreaded LOCA—the loss-of-coolant accident. With a major leak in the primary loop, the water in the reactor vessel will quickly boil away; if the coolant is not replaced within seconds, the fuel rods will overheat and melt.

The Boiling Water Reactor. As the name suggests, the BWR system allows the primary coolant to boil on contact with the hot fuel in the core, so that liquid water and steam coexist in the reactor vessel. The steam is piped directly to the turbine and then condensed and pumped back to the reactor. Thus, the steam circuit has only a single loop, rather than the two-stage process of the PWR. The BWR design has the virtue of simplicity. On the other hand, steam running throughout the plant will routinely pick up low levels of radioactivity from its passage through the reactor. And if something goes wrong—such as the sudden failure of a fuel rod—large quantities of radioactive material would enter the steam circuit.

A BWR power plant looks very different from a PWR. Gone is the domed containment structure and the separate fuel-handling building. The reactor, the swimming pool, and all the rest of the nuclear machinery are in a single large building. Many BWR plants do have one visually distinctive feature: a very tall stack—as tall as one you might see at a fossil-fuel plant, though not as big around. The function of this unusual stack is explained later.

The BWR does have a containment structure; it's just smaller than the kind used in a PWR, so it fits inside a conventional building. The BWR containment is either cone-shaped (with the pointy end up) or lightbulb-shaped (with the screw-in end up). The reactor vessel is suspended near the top of this structure. Below it is a pool of water meant to absorb and condense steam released in the event of an accident.

The presence of steam in the reactor vessel requires some changes in the way the reactor core is designed. Because steam lines must connect to the top of the reactor vessel, the control rods can't enter the core from overhead; they have to come up from the bottom. But if the control rods are under the reactor, you can't count on gravity to insert them in an emergency. The emergency shutdown mechanism uses hydraulic pressure to drive in the rods, then latches them in place mechanically. The reliability of the hydraulic system is critical.

Because the water driving the turbine in a BWR plant passes through the intense radiation of the reactor, it becomes mildly radioactive even when nothing is leaking. Water quality is critical. Any minerals present will be transmuted to radioactive elements; if they are thereafter deposited in the turbine or elsewhere in the system, they can make the whole plant "hot." But even pure water is susceptible to irradiation; both

The elegant spire at Browns Ferry is designed to launch trace releases of radioactive gases high into the atmosphere where they should disperse. The form of the chimney—a slender hyperboloid—is calculated to give the gases maximum upward velocity.

Fan-driven cooling towers at the Browns Ferry nuclear station have a trapezoidal form, wider at the top, so that warm water cascading down the side will wash away any ice buildup. Each of the 16 shrouds atop the cooling unit houses a large fan blade that draws air in through the sides of the structure and discharges it upward.

the hydrogen and the oxygen atoms of the water molecule can be converted to radioactive forms. Worse still, radiation can break apart a water molecule, so that the hydrogen and oxygen are in gaseous form. The gases separate from the steam in the condenser, and a whole subsystem of the plant is needed to capture and dispose of them. That's what the tall stack is for. It is designed to launch the emissions well up into the atmosphere, where the gases disperse. The release of radioactivity is quite small. The Nuclear Regulatory Commission sets the maximum allowable amount at a level believed to be completely safe, and plants routinely stay below 1 percent of that standard. But that doesn't always set the neighbors' minds at ease.

COOLING TOWERS

Ever since the accident at Three Mile Island, the cooling tower has been the sinister symbol of nuclear power. Television reports on nuclear issues set the mood with a haunting image of the towers, often with a cloud of white vapor drifting above them, hinting at some toxic release. This choice of icon could not be less appropriate. In the first place, not all nuclear power stations have cooling towers, and not all cooling towers are installed at nuclear plants. Second, most cooling towers look nothing like the tall, tapered chimneys that have acquired such menacing associations. Finally, the cooling tower is not where the stinger is in nuclear technology. Nothing radioactive passes through it, and any release of radiation would have to come from elsewhere.

There is another irony in the evil reputation of the cooling tower. The reason for building the towers is not that utility companies earn money from them. On the contrary, they are a concession to environmental preservation. Their main function is the protection of aquatic life.

A perfect power plant would convert all the heat liberated by burning fuel or by a nuclear reaction into electricity. Real power plants fall short of that goal. For a coal-fired plant, only about 40 percent of the heat energy is captured in electric power; nuclear plants do even worse, with an efficiency of only about one-third. All the rest is waste. A nuclear plant with an electrical output of 1,000 megawatts must get rid of 2,000 megawatts of waste heat. The flow of water needed to carry off that heat can amount to 500 million gallons per day. This is more than enough warm water to provide a luxurious daily bath for the population of New York City. It's also enough to parboil the fish in a small river or lake. The cooling tower dissipates some of that heat to the atmosphere.

Cooling towers come in two basic types: the fan-driven tower, which is more common but less conspicuous, and the natural-draft tower, which is the one that has entered the public imagination. The choice between them is one of balancing operating costs against capital costs.

Fan-Driven Towers. The typical fan-driven cooling tower is a long, boxy structure, roughly 50 feet wide and 50 feet tall and as much as several hundred feet long. The end walls are of solid construction, but the long side walls consist of louvers to allow for the inflow of air. The warm water is pumped to the top of the tower and falls

A variation on the design of the forced-draft cooling tower wraps the structure into a circle and puts the fans in the center, drawing air through from the perimeter. This tower is at the Coal Creek power station in Underwood, North Dakota.

through a labyrinth of wood or plastic slats called fill. Meanwhile a fan pulls air through the fill into a central void, and then exhausts it upward. Thus, in the fill there is a cross-flow: water trickles downward while air is drawn inward.

Seen from the end, the cooling tower has the form of an upside-down trapezoid. It is wider at the top than at the base, and so the louvered walls slope inward. This shape is chosen to control icing in the winter. Because of the inward slope, the louvers are continually washed by warm water.

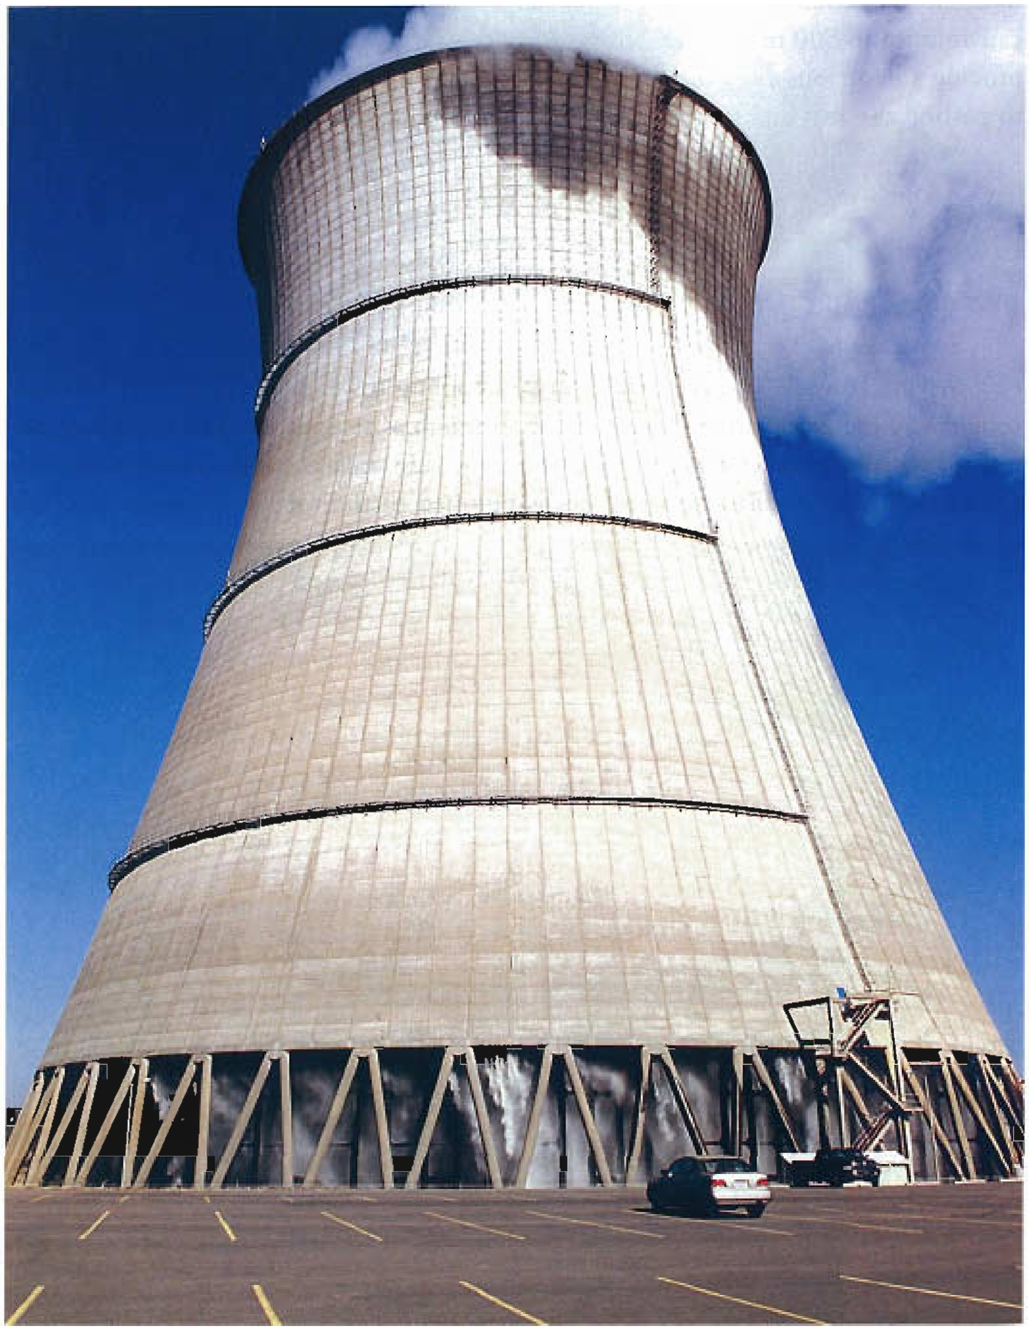

The natural-draft cooling tower at Arkansas Nuclear One has the classic hyperbolic form, tapering to a narrow throat and then opening to a slightly wider diameter. The shape is designed to produce optimum air flow for a given temperature difference between the water and the surrounding atmosphere. At the base of the tower *(detail on opposite page)*, the water is broken up into fine droplets to maximize evaporative cooling.

The fan in a power plant–size cooling tower can be 30 feet in diameter, driven by a motor of 200 horsepower or more. A flared shroud surrounding the blades increases the velocity of the air and reduces noise. Along the inner surface of the fill, lining the central void, are special louvers called drift eliminators. "Drift" is the mist of tiny droplets that get caught in the airflow. The drift eliminators force the airstream to make a sharp turn as it accelerates toward the fan, shedding entrained droplets.

Natural-Draft Towers. The 200-horsepower motor in a fan-driven cooling tower consumes about 150,000 watts, which is a lot of power even for the electric company. On a muggy day the fans can claim 3 percent of a generating station's electrical output. A cooling tower that needs no fan has an obvious advantage in daily operating cost; the disadvantage is that it costs more to build such a tower in the first place.

A natural-draft tower works just like the chimney above a fireplace. In both cases there is a source of heat at the base—burning logs in the fireplace, water from the condenser in the cooling tower. Air is warmed by the heat source and expands, thereby becoming less dense than the surrounding air. The buoyant air rises inside the chimney or the tower, and new air is drawn in at the bottom to replace it. The new air is heated in turn, so that a sustained draft is established.

The optimum size and shape for a chimney depend on the temperature of the heat source and the outside air. With a very hot fire (as in a modern home furnace) even a narrow, straight chimney can draw effectively. The lower temperature of a log fire demands a wider flue. In a cooling tower the temperature difference is only about 10 degrees, which means the chimney must be wide, tall, and carefully shaped. The largest towers are 300 feet in diameter at the base and 500 feet high.

The characteristic shape of a natural-draft tower, which maximizes the airflow, is a hyperboloid (based on the mathematical curve called a hyperbola). Rising air tends

A small hydroelectric installation at the base of the Monticello Dam west of Sacramento, California, exploits a high head (about 400 feet) to extract about 11 megawatts of energy from a modest flow of water. The turbine is a Francis wheel, a hybrid type that draws on both the impulse and the reaction principles. The discharge is above the level of water in the tailrace.

to cool because of the lower atmospheric pressure at increasing altitude; as a result the air loses its buoyancy and its ascent slows. The hyperbolic taper of the cooling tower compensates for this tendency by maintaining a nearly constant pressure up to the throat, the narrowest section of the tower. The slight outward flare above this point allows the air to expand and accelerate upward. The exhaust nozzles of rockets and jet engines have the same shape for the same reasons.

What is inside a natural-draft tower? Mainly nothing. The shell of the tower is raised up on concrete pillars or on a triangulated trusswork of steel, leaving channels on all sides for air to flow in. Just inside this perimeter is an assembly of louvers, fill, and drift eliminators little different from the one in a forced-draft tower. All this apparatus occupies a narrow ring at the bottom of the tower; the rest is a cathedral-like empty space, open to the sky.

HYDROELECTRIC POWER

Waterpower has a history going back to antiquity; it was a thriving and sophisticated technology long before electricity entered the scene. Waterwheels driving elaborate systems of belts and shafts ran the textile factories of New England, sawed timber in the West, and ground grain into flour everywhere. Some of these early mills have been preserved or restored, and a few of the waterwheels are still turning. Nevertheless, apart from sites of antiquarian interest, waterpower now means hydroelectric generation. The waterwheel has evolved into the turbine, much as the paddle wheel of early steamships has evolved into the propeller. The belts and shafts for power transmission have been replaced by electrical lines.

In the 1930s waterpower supplied about 40 percent of the electricity in the United States. Hydroelectric capacity has increased since then, and yet the proportion of all power coming from hydroelectric plants has fallen to only about 15 percent. The reason is that other power sources have grown much faster. The trend is likely to continue, simply because the best spots for hydropower are already occupied.

Two factors determine the power available from falling water: the height of the fall (called the head) and the quantity of water. A little water plunging off a high cliff can produce the same amount of power as a large mass of water falling over a low ledge.

The head and the quantity of water determine what kind of turbine a hydroelectric plant is likely to use. With a high head but only a modest volume of water, the turbine of choice is a Pelton wheel. It works on the impulse principle: nozzles direct high-speed streams of water against curved buckets on the rim of the turbine wheel. The wheel—or runner, as hydraulic engineers prefer to call it—is not immersed in water but turns in air. Pelton wheels spin very fast, and so they are used with the same kind of generator employed in steam power plants. The Pelton wheel and the generator are mounted on a horizontal shaft, which turns at 1,800 or 3,600 revolutions per minute in order to produce the North American standard 60-hertz power.

For lower heads and higher flows, turbines of another type work better. The runner, which has curved vanes rather than buckets, is immersed in a stream of water that flows through it. The runner is mounted on a vertical shaft and enclosed in a spiral scroll case, shaped like a snail shell. Water enters horizontally and flows inward to the runner, then makes a 90-degree turn as it is deflected by the vanes, and exits

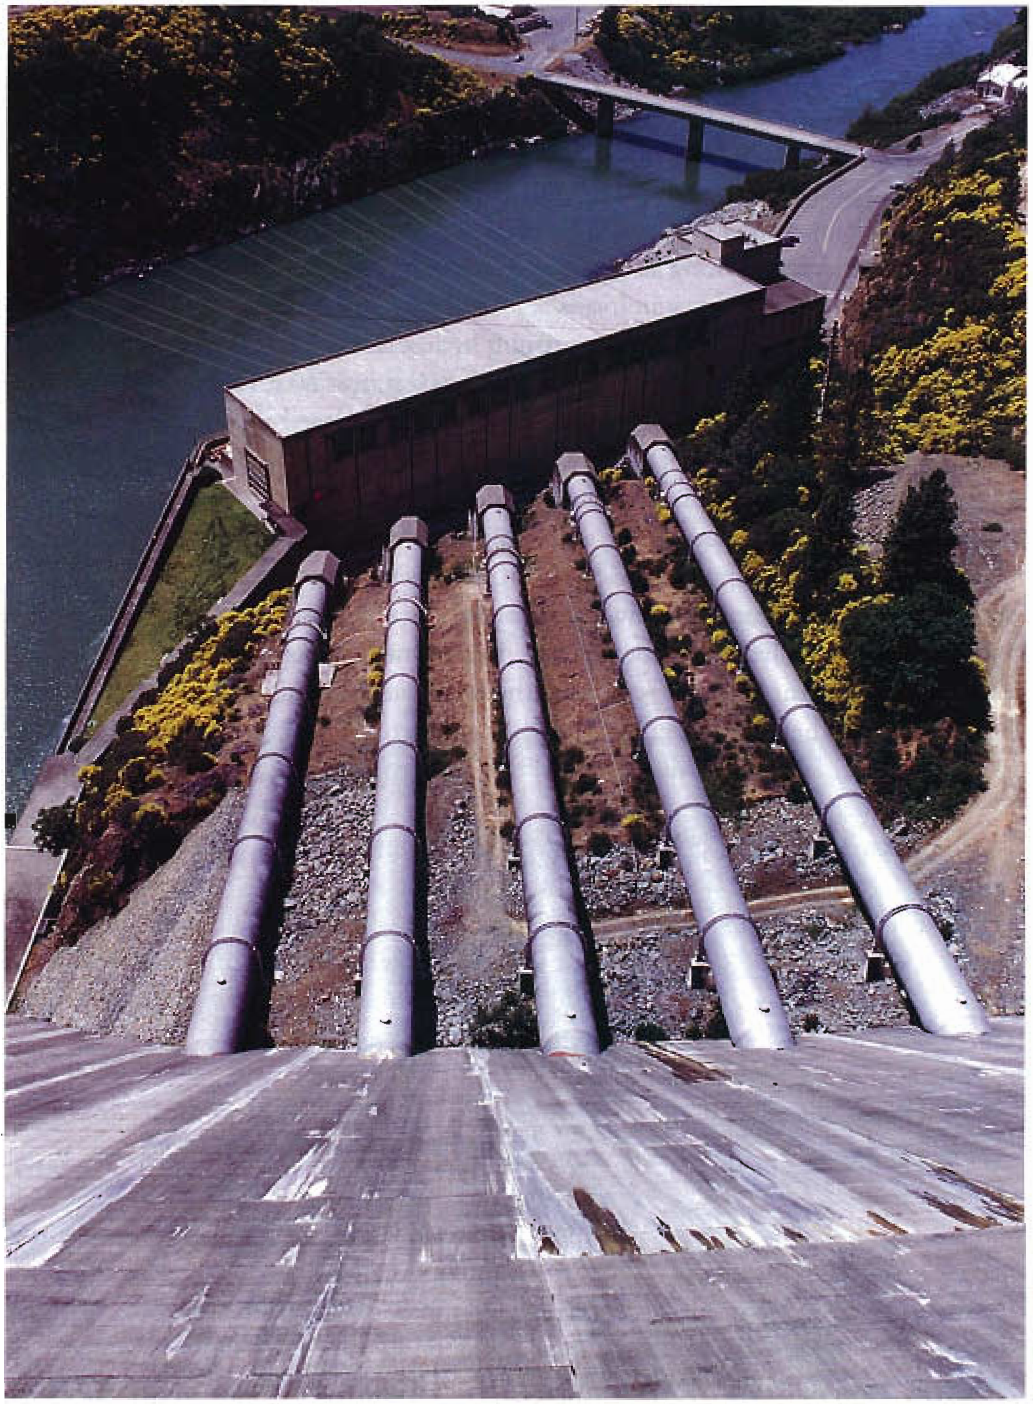

At the Shasta Dam in northern California, five large penstocks emerge from the face of the dam to drive turbines in the powerhouse below.

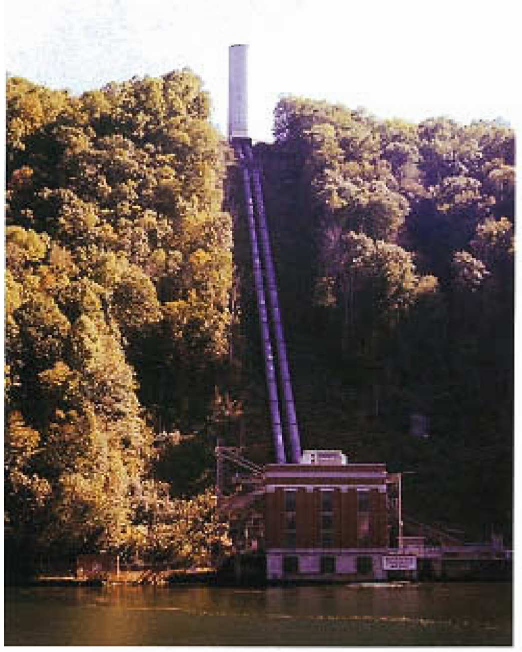

A surge tank at the crest of a hill helps to smooth the flow of water through penstocks feeding a small hydroelectric plant below the Fontana Dam in North Carolina.

downward, parallel to the axis of the turbine shaft. Turbines of this type—called a reaction-wheel turbine—turn much slower than the Pelton wheel.

The low rotation speed of the reaction-wheel turbine calls for a different kind of generator, one that can produce 60-hertz alternating current when turning at only a few hundred revolutions per minute. In the high-speed generators used with steam turbines, the rotor has a single pair of magnetic poles, much like an ordinary bar magnet. The generator produces one cycle of alternating current for each revolution of the rotor; thus, 60 cycles per second requires 60 revolutions per second, or 3,600 revolutions per minute. To generate the same output frequency with a machine that turns more slowly, you need a rotor with more pairs of poles. If the rotor is a cluster of 12 pairs of north and south poles, spaced equally around the perimeter, then on each revolution the output current will go through 12 alternating cycles. The generator produces 60-hertz power when turning at only 5 revolutions per second, or 300 revolutions per minute.

The generators employed in slow-turning hydroelectric plants have as many as 60 pairs of poles, yielding 60 hertz at a rotational speed of just 60 revolutions per minute. These generators are larger in diameter than the high-speed machines, in order to make room for the many rotor windings, but they can be shorter in the other dimension. Because the generator is mounted on a vertical shaft, it takes the form of a squat cylinder on the powerhouse floor, with a small turret at the center that houses the main thrust bearing supporting the shaft of both generator and turbine.

Often the powerhouse of a hydroelectric project is built into the structure of a concrete dam, usually at the foot. Water drops down through passages within the body of the dam, turns turbines installed near the baseline, and then rushes out into the tailrace. In other cases the powerhouse is a structure separate from the dam, possibly miles away. Water is conveyed from the reservoir to the powerhouse through a penstock, which is typically a welded steel pipeline 10 or 15 feet in diameter. Look for a surge tank above the penstock somewhere along the run. It is needed to smooth changes in the rate of flow as the load on the turbine varies. The tank is designed to be about half full during normal, steady-state operations. If the gates suddenly open wider, calling for more water, the surge tank is drawn down momentarily to help meet the demand. When the gates close suddenly, the surge tank is even more important: it gives the moving water somewhere to go as it decelerates, preventing the hard knock called water hammer.

The environment in the generator gallery of a hydroelectric plant is calmer than the turbine hall of a fossil-fuel plant. Gone is the shriek of steam. The noises are all low notes—hums, buzzes, groanings, rhythmic vibrations that you feel rather than hear. Workers—if there are any—can converse as quietly as in an office. The control of a hydroelectric plant is also less hair-raising than that of either a fossil-fuel or a nuclear plant. Power output is regulated by gates that control the flow of water through the penstock and into the turbine. An automatic governor system adjusts the gates to track variations in load and keep the generator turning at a constant speed.

One of the features of hydroelectric plants most welcome to power dispatchers is that they can be started up and shut down at a moment's notice. It takes as little as two minutes to get a unit up to speed and synchronized with the power grid. This makes hydroelectric power attractive as a means of satisfying short-term peak loads. When you come home in the evening and switch on the lights and the TV, somewhere a gate in a penstock has opened very slightly and sent a few gallons more down the penstock.

OTHER ENERGY SOURCES

Fossil-fuel plants, nuclear reactors, and hydroelectric plants account for 99 percent of the electric power generated by utility companies in the United States. Everything else—all the "alternative" energy technologies—amount to just 1 percent, and so they are pretty marginal in economic terms. But the alternative energy sources have a conspicuous place in the landscape and in public consciousness, even if they don't yet make much of a dent in the energy budget. And their contributions are growing. Three of these technologies are described here: wind power, solar power, and geothermal power.

Blowin' in the Wind. Wind power, like waterpower, has a long history. The Old World windmill, with its broad cloth-covered blades, or sails, goes back at least 800 years. But wind technology has been evolving rapidly in recent decades, and modern windmills look nothing like their ancient prototypes. They are tall and spindly, with narrow blades like those of an airplane propeller but on a vastly larger scale.

Generators at Hoover Dam are mounted with the shaft vertical. Each generator has 60 pairs of magnetic poles and turns at 60 revolutions per minute to produce 60-hertz alternating current. The eight generators seen here are on the California side of the dam; there are nine more on the Nevada side.

They aren't even called windmills anymore; the preferred term is *wind turbine*. Moreover, they are usually brought together in large wind farms, with hundreds of turbines lined up along ridges or scattered across a broad plain.

In the United States, wind farming got its start as a California thing; for a time, that one state produced half of the world's wind energy. The three biggest California wind-energy areas are at Tehachapi Pass, 100 miles northeast of Los Angeles, where a range of hills separates the Central Valley from the Mojave Desert; Altamont Pass, near the town of Livermore east of San Francisco Bay, where another range of low hills divides the coastal plain from the Central Valley; and San Gorgonio Pass, in the southern California desert near Palm Springs, where once again the wind has to rise over hillsides to reach an interior valley. All three areas have major highways running through them, so you can easily get a look at the machinery. (Some European wind-energy developments are even more tourist-friendly, with visitor centers and picnic areas out among the fields of turbines.)

The great boom in California wind power was launched by tax laws in the 1970s that encouraged experiments with alternatives to fossil-fuel and nuclear power plants. But the first generation of wind turbines proved to be expensive and unreliable, and they had a hard time competing against the more mature smokestack technologies. The result was a slump in the wind-power industry during the 1980s. Today the wind is rising again, however, even though some of the tax incentives have expired. The new generation of wind turbines—many of them built in Europe or inspired by European designs—are more efficient and cheaper to maintain, and they also work in a wider variety of wind conditions. One result is that wind farms are no longer just a California crop; you'll find them in Texas and Minnesota and Iowa and Vermont, and in years to come they may well sprout on hillsides everywhere. The state with the richest potential for wind power is North Dakota; if the farmers of North Dakota farmed wind instead of wheat, in principle they could supply a third of the electricity consumed in the United States.

As of 2003, the total capacity of wind turbines in the United States was about 6,000 megawatts—the equivalent of five or six nuclear plants. Europe has far more wind-energy capacity: well over 14,000 megawatts in Germany alone, and another 6,000 in Spain.

When you look at the spindly propeller-like rotor of a modern wind turbine, and compare it with an old multiblade farm windmill, you might conclude that the new technology is letting most of the wind slip by without getting any benefit from it. But that's an illusion; the new wind turbines rely on different physical principles. The older windmills are drag devices: they arrange the blades so that the wind pushes against a broad surface. The modern ones are lift devices: the air passes over an airfoil, like an airplane's wing, and pulls the blade through the air. Drag devices produce higher torque (turning force), but in most conditions they extract less energy from the air.

Just as engineers have varied the thickness of the blades over the years, so too have they disagreed over the ideal number of blades for a wind turbine. The blades are

Wind turbines are planted in strict military rows along a northern California hillside, but their twirling blades are not to be disciplined. The wind farm is near the town of Tracy.

expensive, and so a design with fewer of them might be expected to reduce the cost of the machine. The minimum number, obviously, is one, and one-bladed rotors have actually been tried. They look funny, to say the least; even though a small counterweight keeps the machinery in balance, the visual impression is of something dramatically out of kilter. But that's not the big problem with one-bladed designs; more serious is that the one-bladed rotor has to turn faster to produce the same energy output as a turbine with more blades, and higher speed brings more strain and noise.

Two-bladed rotors have a subtler problem. The balance of the blades is perfect, but trouble comes whenever the wind shifts direction and the turbine has to swivel—or yaw—to stay pointed into the wind. When both blades are vertical, there's no resistance to yawing, but as the blades turn toward the horizontal, the inertia increases. This cyclic change in resistance to yawing—going from maximum to minimum twice in every revolution—creates vibration and stress, shortening the life of the blades.

The vast majority of modern wind turbines have exactly three blades. Apparently three is just enough to solve the problems of speed, balance, and vibration; any more than three would be a needless expense.

A wind farm near Tehachapi Pass in California had more than 600 turbines spinning when this photograph was made in 1999. The machines are planted like orchard trees in a rectangular array; many other wind farms line the machines up along ridges.

Still another contentious issue in wind-turbine design is whether to mount the blades on the upwind or the downwind side of the machine. Letting the blades trail behind the rest of the turbine has one big advantage: the blades can act like a weather vane, automatically turning the machine to face into the wind. When the blades are mounted in front, some kind of steering mechanism is needed to sense the wind and forcefully pivot the turbine whenever the direction shifts. Nevertheless, the prevailing design has the rotor in the front, with a complex power-steering unit to keep it properly pointed. The reason is one I never would have guessed: with a rear-mounted rotor, the turbine blades pass through the "wind shadow" of the tower structure on every revolution. The result is a cyclic variation in wind force that can set the blades vibrating, thus creating yet another source of fatigue and premature failure.

All this leads to a portrait of the typical wind turbine. It has three blades, each about 50 feet long, made of carbon fiber or some other lightweight ultrastrong material. The blades are attached to a hub, which in turn pokes out the front of a streamlined housing called a nacelle. Inside the nacelle, which is the size of a moving van, are the generator, a gearbox, and other machinery needed to control the turbine. The nacelle is mounted atop a hollow steel pylon, 100 feet high, 20 feet in diameter at the base, and tapering gradually toward the pinnacle.

On top of the nacelle you might notice a small airplane-shaped weather vane—just like the ones you see on suburban lawns. This is the sensor for the mechanism that keeps the turbine facing into the wind. Elsewhere on the wind farm, scattered among the massive turbines, are tiny, spinning cups of anemometers on tall masts. These instruments are there to keep records of wind speed for use in analyzing turbine performance and also to shut the turbines down if winds approach dangerous levels.

Every wind turbine is designed for a limited range of wind speeds. Too little wind, and it's not worth starting up. Too much, and the machine could destroy itself. On some turbines the blades can be "feathered," or twisted so that the wind won't spin the rotor, when speeds get into the danger range. Others have aerodynamic "spoilers" with the same purpose. The final line of defense is a mechanical brake that binds the main shaft—but the operators at a Tehachapi wind farm told me they're not eager to use that one. Climbing the tower in a gale to tighten down the brake is more excitement than they're looking for.

In normal operation most wind turbines spin at a fixed rate. You might think they would speed up and slow down as the wind varies, but instead they are designed to adjust the pitch of the blades so that the speed stays constant even as the energy output changes. Running at a constant speed makes it easier to maintain the steady frequency of the alternating current that the turbine supplies to the power grid. The speeds are slow enough that you can count the revolutions. At one big wind farm in northern California I found that the turbines were making 40 turns per minute; at another farm down the road the speed was 72 revolutions per minute.

Most wind turbines turn clockwise, as seen from the hub side of the rotor. But there's no fundamental reason for this, and a few machines spin the other way.

The three-bladed propeller of the wind turbine drives a generator inside the nacelle, atop the mounting mast. Within the range of operating wind velocities, the turbine rotates at constant speed, adjusting the pitch of the propeller blades to regulate power output.

The most unusual of all wind-turbine designs is the vertical-axis machine developed by D. G. M. Darrieus. Instead of an airplane propeller, it's an eggbeater: a vertical shaft with two thin blades bent into bow shapes so that they can be attached at the top and the bottom. The big advantage of the Darrieus design is that it responds equally well to wind from any direction, with no need to pivot when the wind shifts. Also, the generator can be mounted at ground level, which makes it more convenient for maintenance and allows a lighter structure. Nevertheless, the design seems to have gone out of fashion. In the northern California wind farms, the few Darrieus machines still running were looking pretty tired and careworn when I last saw them.

A large wind farm, with hundreds of turbines, makes a powerful visual impression. From a great distance, they look like cheerful daisies or sunflowers planted in neat

LONG BEFORE THE WIND FARM, THE FARM WINDMILL

The multibladed, pinwheel-like farm windmill was an American invention in the middle of the nineteenth century that became an icon of American rural life. By the 1890s the windmills were a standard item in the Sears catalogue, and traveling salesmen peddled dozens of brands to farmers throughout the Midwest and Southwest. An estimated 100,000 of them are still at work in the United States, mostly pumping water on ranches in the western states. In aggregate they may put out 250 megawatts.

The most famous brand of farm windmill is the Aermotor, designed by Thomas O. Perry. At the peak of production in the 1890s some 20,000 per year were being made. The company is still in business, in San Angelo, Texas, where they manufacture about 500 windmills annually.

Most farm windmills are erected directly over a well shaft. A crank arm connected to the fan wheel operates a piston down in the well tube. Some later models have gearing to reduce the speed of the pump and increase the force available. The windmills come in many sizes, but the most common ones have a rotor eight feet in diameter and can pump up to 10 gallons per minute.

As with other styles of wind machines, the big challenge in building a farm windmill is making sure it doesn't fly to pieces when the wind blows too strongly. Over the years, designs were equipped with spring-loaded vanes or centrifugal weights or other contraptions to furl the blades or turn the fan wheel parallel to the wind when the speed reaches dangerous levels. Another engineering issue is the need to grease the bearings of a windmill mounted atop a tower 20 or 30 feet tall. Nobody ever wanted to climb up there in the middle of the winter. A number of tricks were tried (including hinged towers that fold in half to bring the works down to ground level), but nevertheless there are a lot of squeaky old windmills out there.

The Fairbanks-Morse New Eclipse model pictured below was still twirling cheerfully, despite a peppering of bullet holes, when I photographed it in northern California in 1999.

Darrieus wind machines take a different approach to wind energy: by mounting the rotor on a vertical axis, they eliminate the need for a turbine to swivel to follow changes in wind direction. These four Darrieus turbines were photographed in the Altamont Pass of northern California.

rows. As you get closer, their gargantuan scale becomes apparent. And when you finally stroll among them on an afternoon with a fresh breeze, there is something almost comical about their appearance: they seem to be waving their arms frantically, signaling to unseen friends on the next hill, or else they are turning cartwheels along the ridgelines like overexcited children. The sounds of the turbines are equally extraordinary: the swish of the blades slicing through the air, the whir of the gearbox, the hum of the generator, an occasional groaning or booming as turbines yaw with the shifting wind. Some wind machines produce a deep-bass whomp-whomp-whomp as the blades pass through the wind shadow of the support pylon. And I've even heard a few squeaky wheels.

As it happens, these very sights and sounds have become an impediment to further development of wind energy: people don't want to see wind turbines on the skyline, or hear them. There's an irony in this. Modern wind power began as a "clean" alternative to nuclear and fossil-fuel plants, promoted by environmental activists and resisted by utility companies that were skeptical of the economics. Today the opposition to wind power comes from environmental groups that see the turbines as despoiling the landscape. There is also concern about the turbines as a hazard to birds, which sometimes wander into the blades at night or in fog. Meanwhile, the utilities have begun to warm up to wind energy, as costs have come down.

Wind has proved itself as a supplement to conventional power sources. But if we want to rely on it for a large fraction of the base load, there's a problem: you can't tell the wind when to blow. In the argot of the power engineer, wind is not a "dispatchable" energy source. This puts a limit on wind's total contribution to the energy budget, but we are still far from reaching that limit.

Mirrored-trough solar collectors cover a square kilometer of the desert floor at Kramer Junction in southern California. The peak output of the array is about 150 megawatts. The mirrors appear deep blue from this point of view because they are reflecting the sky.

Let the Sun Shine In. The friendly star whose neighborhood we inhabit sends planet Earth a steady energy flux of 175 billion megawatts, which is equivalent to the output of a few hundred million nuclear power plants. Now, admittedly, a third of that energy is reflected back into space before it ever reaches the ground, but there's still plenty left over if we could only figure out how to collect it. But, as of 2000, U.S. utility companies were gathering only about 5,000 megawatts of solar energy.

There are two quite different types of solar power plant. In a solar-thermal system, sunlight is simply a source of heat. Photovoltaic plants generate electricity directly from light, with no moving parts.

The simplest solar-thermal technology uses flat-plate collectors, which work a lot like greenhouses. Pipes carry water or some other fluid through a glass-covered box tilted toward the sun; the pipes and the inside surfaces of the box are painted black to absorb as much heat as possible, and the glass cover retains the heat. It's simple and reliable, but the maximum temperature is well below the boiling point of water, so you can't produce steam to turn a turbine and generator. Most flat-plate collectors are rooftop installations used for water heating and space heating.

To reach higher temperatures, you have to concentrate the sunlight, collecting over a wide area, and focusing it on a smaller patch. In principle, lenses might be used to do the focusing, but in practice it's always done with mirrors. The ideal shape for a reflective solar collector is a parabola, because this curve has the property that parallel rays of sunlight striking the mirrored surface are all reflected to the same focal point.

One style of collector is a long trough with a parabolic cross section, mirrored on the upward-facing surface so as to focus sunlight on a tube that runs parallel to the trough at just the right position to receive all the concentrated light. The biggest installation of parabolic troughs in the United States is at Kramer Junction, a crossroads in the southern California desert near the city of Barstow. Each trough is about 15 feet across and 150 feet long, assembled from 300 curved glass panes. Altogether there are 546,000 panes, with a total area of one square kilometer (about 250 acres). At the focus of each trough is a receiver pipe housed in a glass vacuum tube to retain heat. The receiver pipe is painted black for best absorption, but when the plant is operating and the sun is shining, the pipe glows bright white, like a fluorescent lamp.

The receiver pipes are filled with oil, which is heated under pressure to more than 750 degrees Fahrenheit. The hot oil is pumped to a heat exchanger, where it generates steam at 670 degrees; the steam then runs a fairly conventional turbine and generator. The Kramer Junction solar array is divided into five independent plants, each capable of producing 30 megawatts of electricity in full summer sun. It would be more efficient to run one big unit rather than five small ones, but at the time the plant was built, tax incentives for solar power were limited to plants of 30 megawatts or less.

Each trough collector at Kramer Junction has a parabolic cross section, which concentrates the sun's rays on a tube installed at the focus of the parabola. The tube is painted black, but it glows white when the collector is operating. Oil pumped through the collector tubes gathers the solar heat and generates steam to run a turbine.

The troughs at Kramer Junction are aligned on a north-south axis, and during the course of the day they tilt, facing east in the morning, then directly overhead at local solar noon, and finally turning toward the west at sunset. The tracking is done automatically with a sensor that tries to keep the focused image of the sun centered on the receiver pipe. Standing near a trough, you can hear it adjust itself every few seconds.

With 250 acres of glass in the middle of a dusty desert, washing the mirrors is a full-time job. It's done at night, with high-pressure hoses. It takes about two weeks to wash off the entire field of collectors; then the washing starts over again.

Even with a large parabolic trough, the temperatures are not as high as power-plant engineers would like to see for maximum efficiency. To reach still higher temperatures, the trick is to focus the sun's light not on a line of pipes but on a single point. One way to do this is with a mirror in the shape of a paraboloid, like a satellite-dish antenna or a radio telescope, but building a really big paraboloidal mirror that can tilt to track the sun is an engineering challenge. A better idea is to set out lots of small mirrors, called heliostats, which can be adjusted individually so they all reflect sunlight onto the same point. According to legend, the principle was invented by Archimedes in 212 BC, when he had a troop of Greek soldiers at Syracuse use their bronze shields as heliostats to burn the ships of an invading Roman fleet.

Heliostats were the basis of the Solar One project at Daggett, another town near Barstow in southern California. Some 1,800 flat mirrors, with a total area of 17 acres, were continually adjusted so that they all reflected the sun's image onto a black receiver at the top of a 300-foot tower. Water pumped through the receiver boiled to produce steam at about 900 degrees Fahrenheit, which then drove a turbine and generator at the base of the tower. Solar One operated as a pilot project in the 1980s. Later, a new receiver was fitted to the tower and the plant was recommissioned as Solar Two. Instead of boiling water directly, the solar energy was now absorbed into

The Solar One and Solar Two projects at Daggett, California, achieved higher temperatures by focusing the sun's light on a point rather than a line. The last of the California experiments was shut down in 1999, but the program continues in Spain.

molten salt, which could be stored for a few hours so the plant could continue generating electricity even after sunset. Solar Two was shut down in 1999. A similar plant called Solar Tres is under construction in Cordoba, Spain.

Photovoltaic technology has almost nothing in common with solar-thermal power beyond the basic fact that both rely on sunlight as the ultimate source of energy. A photovoltaic device dispenses entirely with boilers, turbines, and generators; it converts light directly into electricity in one step. The transformation is accomplished with no moving machinery. All the complexity is hidden in the microscopic structure of the photocells, which are high-tech products of the semiconductor industry.

The photoelectric effect was first noticed more than 150 years ago, and the first good explanation came from Einstein in 1905 (that's what he won his Nobel Prize for—not for relativity theory). The key idea is that light comes in packets, or particles, called photons, each of which carries some definite energy. If a photon's energy is great enough, it can kick an electron out of its stable orbit inside an atom, making the electron available to carry an electric current. These forced evictions are happening all the time; a coin sitting in the sunshine is seething with liberated electrons. But in most cases all the activity comes to naught because the electrons just wander around for a while and then fall back into the atomic orbits they came from. A photovoltaic cell is designed to capture the ejected electrons and put them to use.

Most photovoltaic cells are made of silicon, and if you can get a closeup look, you may find them to be quite beautiful objects, with crystal facets like frost on a windowpane, in various shades of blue. Stripes or grids of metal electrodes are laced across the surface to collect the electric current.

The output of a photovoltaic collector is direct current (DC) rather than the alternating current (AC) of the national power grid. Also, the voltage produced by an individual cell is closer to that of a flashlight battery than that of a power-plant generator. Thus, connecting a panel of cells to the utility grid calls for special electronics to boost the voltage and to convert from direct to alternating current.

Then there's the matter of cost. Even though the fuel is free, photovoltaic power remains substantially more expensive than electricity from coal-fired power plants. As a result, you are most likely to see arrays of photocells in places where utility lines haven't reached—powering emergency telephones along highways, powering the lights on marine buoys, powering remote homesteads. And most remote of all are the many spacecraft that have relied on photovoltaic power.

The cost of photocells has been coming down steadily for two or three decades, and interest is finally growing in utility-scale projects. The pioneer in this field in the United States is the Sacramento Municipal Utility District in California, which operates more than eight megawatts of photovoltaic collectors. One big array of photocells is next to a decommissioned nuclear plant, but most of the collectors are distributed around the utility's territory on residential rooftops and in parking lots.

It's sometimes said that to run the country on solar power we'd have to pave the whole landscape with collectors. It's not nearly that bad. According to one estimate,

Photovoltaic cells, which generate electricity directly from sunlight, are made of polycrystalline silicon. The array of cells above is installed at Montgomery College in Germantown, Maryland. A detail, below, shows the grain structure of the silicon material and the grid of metallic conductors laid down over it to collect current.

photovoltaic plants that could meet the electricity needs of the United States would occupy a little less than 12,000 square miles. That's a lot of land, but it's only about one-third of 1 percent of the total land area of the nation. So there's no need to pave over the whole country, just the state of Maryland.

Warmth from the Earth. Why bother burning coal to make steam when you can just drill a hole in the ground, and let the steam come whistling out? This is the idea behind a geothermal power plant. It's a great idea. It dispenses with the whole fuel supply and the furnace and the boiler. The trouble is, it only works at a few places in the world—rare hot spots where the heat of the deep earth bubbles up unusually close to the surface.

The simplest geothermal plants take steam straight from the well and, after minimal processing to remove a few impurities, pipe it to the input port of a turbine to generate electricity. In the earliest plants, the spent steam leaving the turbine was just vented to the atmosphere. That's no longer done, for two reasons. First, the steam carries some obnoxious contaminants, chiefly hydrogen sulfide, that could make a geothermal plant a worse polluter than a coal-fired one. Second, it turns out that the supply of underground steam is far from inexhaustible. To keep it flowing, you have to recycle it, pumping water down recharge wells to prevent the reservoir from drying out. Recapturing the spent steam requires a lot of additional equipment: a condenser, pumps and valves, and cooling towers to get rid of excess heat. The cooling

The A. W. Hoch geothermal power plant *(right)*, near Calipatria in the Imperial Valley of southern California, generates electricity from steam and hot brine brought up from several thousand feet below the Earth's surface. The fluids come out of the ground at a temperature of 450 to 500 degrees Fahrenheit. One of the insulated pipelines that carry well fluids can be seen in the foreground. The remaining equipment is needed to separate steam from water and to remove salts and other impurities. At one of the nearby wells *(above)* the pressure gauge reads 110 pounds per square inch.

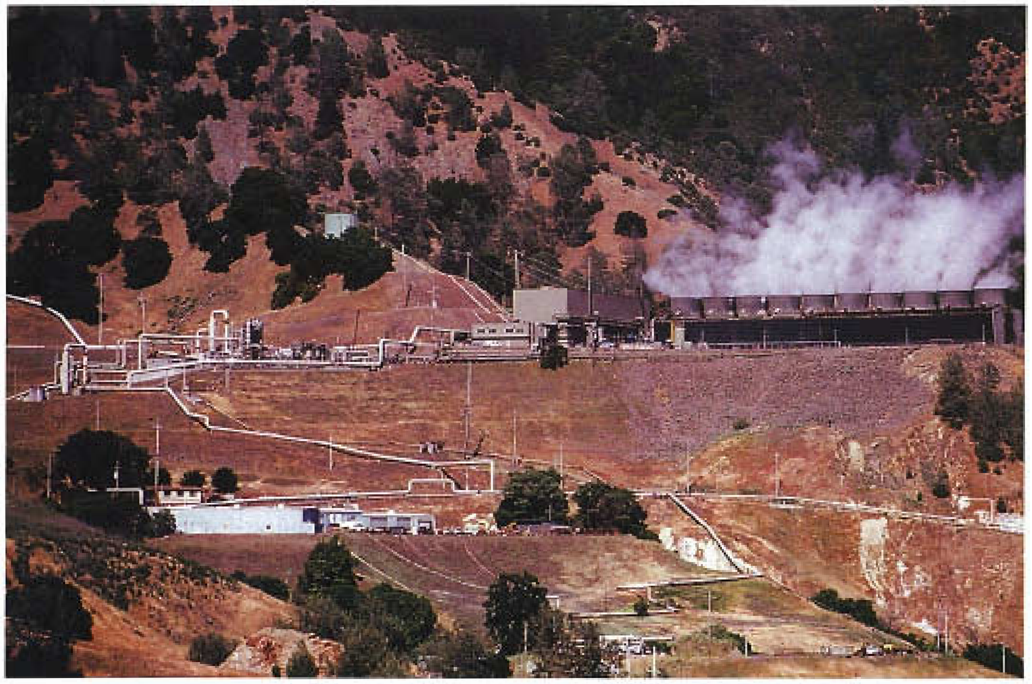

Geothermal steam at higher temperatures allows a simpler energy cycle at the Geysers in northern California. This "dry steam" can be piped to a turbine with little preprocessing; then it is condensed and reinjected into the earth.

towers tend to be the largest component of the entire plant, and their plumes of rising vapor make them even more conspicuous.

Piping steam from the ground directly into a turbine is a plan that works only if the steam is superheated to a few hundred degrees above the boiling point. It's called dry steam, since it's too hot for water to condense in pipelines or holding tanks. Unfortunately, few geothermal sites produce enough dry steam for commercial power production. One of these areas is Larderello, near Pisa in northern Italy, where power production began almost 100 years ago. The one spot in the United States where dry steam comes out of the ground is the Geysers, a landscape of hot springs and hissing fumaroles tucked among the wine-growing valleys north of San Francisco. At the Geysers geothermal power plants have been running, off and on, since the 1960s.

Although dry steam is a scarce resource, many areas of the world have enough subterranean heat to produce large quantities of mixed hot water and steam. But getting useful power out of these lower-temperature fluids calls for more elaborate machinery. Steam and water have to be separated, and then some of the hot water can be persuaded to vaporize by lowering the pressure in a device called a flasher. In another type of plant, heat from the geothermal fluids is used to boil a more volatile liquid such as butane or ammonia that circulates through a turbine in a closed cycle.

In any geothermal area, one thing you're sure to notice is a network of pipelines that feed steam or water from widely scattered wells to the central power plant. The pipes are often of large diameter, and they look even fatter because of a thick blanket of insulation. At intervals there are big inchworm-like loops to allow for expansion and contraction as the temperature of the pipe changes.