



OIL AND GAS

IN THE FALL OF 1973 one piece of the industrial infrastructure emerged from the unnoticed background of American life and became a public obsession. The gas pumps had run dry. Any filling station with fuel to sell had a long queue of cars snaking around the block. Prices soared. Getting the tank filled—a chore that had been so routine it slipped beneath the level of conscious attention—was suddenly a challenge that called for strategy and guile, not to mention getting up before dawn.

The gasoline shortage lasted only a few months, but it made a strong impression on those who lived through it. For a while, American automobiles became smaller, lighter, and less thirsty. A pipeline from the North Slope oil fields across Alaska was quickly approved and built. A 55-mile-per-hour national speed limit was enacted as a fuel-saving measure and ruled the roads for more than 20 years.

The causes of the 1973 oil crisis were more political and economic than technological, but the event nonetheless prompted much sober thought about life in a world of finite resources. The oil was not running out in 1973, and 30 years later we are still pumping it out of the ground at a furious pace, but the idea that it will not last forever is now taken seriously—even by some of the oil companies. (Though evidently not by the owners of sport utility vehicles.) This chapter looks at the infrastructure of the oil industry, from the well through the refinery to the filling station. Finally, there is a section on the natural gas industry, which has a rather different culture.

THE NATURAL HISTORY OF AN OIL WELL

The drilling rig is the universal emblem of the oil industry: a tapered steel derrick, usually depicted with a gusher shooting up through the middle of it. Your chances of ever seeing a gusher are nil. Even to see a drilling rig, you need to be in the right

A spherical tank in Port Arthur, Texas (*opposite page*), forms part of a network of refineries, tank farms and pipelines that delivers the petroleum products to keep the United States fueled and lubricated. Spherical tanks are used to hold fluids under pressure; commonly they contain liquefied petroleum gas (LPG), which is mainly propane. The stacks atop the tank are vents for emergency pressure-relief valves. The bright red pipes carry water or foam for firefighting.

place at the right time. In the life cycle of an oil well, drilling is a relatively brief phase. It lasts a few months, and then the derrick is broken down and hauled away. You're much more likely to happen upon a well in the production stage, which can last for decades. Nevertheless, drilling is where it all begins, and it's certainly the adventurous side of the oil business, the world of wildcatters and roughnecks.

The Drilling Rig. When drilling is under way—when the rig is “turning right and making hole”—several things have to be going on at once. The bit, or tool, which does the actual drilling, has to be simultaneously turned and pressed downward so it will cut into the rock. At the same time, a lubricant called drilling mud has to be pumped down to the bit to carry away the cuttings.

The bit is connected to the surface by the drill string, which is not a string at all but a rigid, hollow steel pipe assembled from sections called joints, which are typically 30 feet long. A well 15,000 feet deep would need 500 such joints. The drill string transmits both vertical forces (pressing the bit down into the rock and later hauling it out again) and the twisting force that turns the bit. It also carries the mud through its hollow bore.

The twisting force comes from the rotary table, a motor-driven, spinning platter set into the main deck of the drilling rig. The rotary table grips a special length of drill pipe called the kelly, which has a square or hexagonal cross section matching a bushing in the center of the rotary table. The table turns the kelly in the same way that a wrench turns the head of a bolt; the kelly then turns the rest of the drill string. As the kelly is being spun by the rotary table, it is also free to slide vertically through the table so downward force can be maintained.

The drill string is not pressed down from above; pushing on a pipe three miles long would accomplish nothing but buckling the tubing. Instead, the downward force is provided by a series of extra-heavy joints of drill pipe, called drill collars, installed immediately above the bit; it is the weight of the collars that drives the bit

GETTING A LOOK

Much depends on where you live. If you come from the “oil patch” in Texas and Oklahoma, then the sight of a drilling rig is not a novelty, and a sucker-rod pump nodding over an oil well probably seems as commonplace as a traffic light. In other parts of the country, including most of the populous Northeast, oil wells are exotic rarities. The distribution of refineries is also patchy. There are clumps of them along the Gulf Coast and near some major ports, but fewer inland.

Virtually all elements of the oil-and-gas infrastructure are privately owned and receive visi-



tors only by special arrangement. On the other hand, petroleum is largely an “outdoor” industry: the machinery is not hidden away behind closed doors, and there is much to see from outside the fence.

Something to keep in mind, however: Refineries have often been cited by the FBI and the Department of Homeland Security as potential terrorist targets. If you stop by the roadside to take pictures, you may attract the attention of plant security forces or the local police.



A gleaming white rig in Marlow, Oklahoma, drills for natural gas on behalf of Chesapeake Energy. The rig is a portable, temporary structure. The derrick folds up and breaks into pieces for transport, and the rest of the rig consists of modules the size of truck trailers, which can be hauled to a new site and reassembled in a matter of days. The two stairways lead to the main drilling deck. The house-trailer-size shelter on the near side of the deck is the “doghouse,” where instruments and controls are kept out of the weather. The ramp on the right side is used to lift 30-foot “joints” of drill pipe up to the main deck. The rig is capable of reaching depths of about 15,000 feet; the well in Marlow was expected to tap a natural-gas reservoir at roughly 7,000 feet.



into the bottom of the hole. Ordinary joints of drill pipe weigh a few hundred pounds; drill collars weigh 2,500 pounds or more. An entire 15,000-foot drill string could weigh well over 200,000 pounds. The weight is supported by a block and tackle suspended from the top of the derrick, so that the string remains under tension at all times. Every few minutes a brake on the main hoist emits a loud squawk or honk as it automatically adjusts the tension—one of the characteristic rhythmic noises of the drilling site.

The routine of drilling is to keep “making hole” until you have “drilled down the kelly.” When the top of the kelly shaft has almost reached the rotary table, the crew stops the rotary, hoists the drill string far enough to expose the top of the uppermost joint of drill pipe, and inserts a clamping device called the slips, so that the drill string cannot fall down the hole. Then the kelly is disconnected from the string, another joint of pipe is inserted, and the kelly is reattached to the top of this joint. The joints fit together by means of threaded couplings, female at the upper end and male at the lower end. The threads are tightened with “tongs” that work much like a plumber’s pipe wrench but weigh several hundred pounds.

At intervals, the drill bit needs to be changed. This requires hauling the entire drill string back up to the surface, a process known as tripping out. The converse process, naturally, is tripping in. As a rule, the drill pipe is not broken down into individual joints during such a round trip; instead, sections two or three or four joints long (doubles, thribles, or fourbles) are stood on end in a rack inside the derrick.

Top Drives and Downhole Motors. The rotary-table-and-kelly mechanism for spinning a drill bit was a marvel of ingenuity—and it takes great ingenuity to keep it running. In recent years simpler schemes have come into use.

The drilling crew on the rig in Marlow adds another 30-foot joint of pipe to the drill string. Above, the threaded end of the new joint is wrestled into position; at right, the threads are tightened with enormous wrenches known as tongs. During this operation the entire length of drill string is supported by a two-handled device called the slips clamped to the drill pipe where it goes through the main deck.



A drilling rig with top drive has an electric motor mounted on vertical rails. The motor's drive shaft connects directly to the drill string, eliminating the need for both the rotary table and the kelly. Furthermore, the motor's torque can be used to tighten the threaded connections between joints, thereby dispensing with tongs as well. With top drive, smaller crews can drill faster and with less fuss.

So why don't all rigs have a top-drive motor? It turns out that mounting the main drive motor aloft requires a much stronger and heavier derrick, which takes longer to set up, break down, and transport. Most rigs move to a new hole every few months, so the loss of portability is a serious drawback.

Another drilling innovation puts the motive force at the opposite end of the drill string. Instead of turning the entire string from the surface, a motor is installed just above the bit, near the bottom of the hole. Thus, only the bit turns; all the rest of the drill string is locked in place. This down-hole motor is not an electric one. It runs on mud! It is a turbine that extracts energy from the stream of drilling mud pumped through it. (A dentist's drill works the same way, except that the fluid driving the turbine is air—thankfully—instead of drilling mud.)

The big advantage of a down-hole motor is that it facilitates steering the drill. Various geological strata can deflect the drill, and a correction is needed to bring it back to vertical alignment. Moreover, the driller may *want* to deviate from the vertical. Sometimes multiple wells fan out like spider legs from a single drill site to reach widespread areas of a petroleum reservoir. (This practice is particularly common in offshore drilling.) Sometimes a well is drilled vertically into an oil-bearing stratum and then deflected horizontally, so that it can collect oil from a larger region. Such wells can be drilled even with a conventional rig, but the technique is easier and more precise with a down-hole motor.

Mud. Drilling mud is the least obvious and most ingenious aspect of rotary drilling technology. It is also a major expense—a considerable fraction of the capital investment needed to drill a well goes into buying mud.

Drilling mud cools and lubricates the cutting edges of the drill bit, but that could be done as well by an ordinary fluid such as water or oil. Mud has other important properties. First, because it is a thick, viscous fluid, it holds the cuttings in suspension; rock chips carried away from the bit do not settle back to the bottom of the hole even when circulation stops. Second, because it is dense (sometimes twice the weight of water), it resists the pressure of fluids deep underground. And it leaves behind a coating of clay that helps to seal the bore of the well.

Legend has it that the first drilling mud was made by driving cattle through a bog, but the modern stuff is a high-tech product. The main ingredient is bentonite, a kind of clay. The density can be increased by adding barite, a mineral rich in the heavy metal barium.

A mud pump draws the mud from an open pit or tank and forces it down the bore of the drill string to emerge through nozzles in the bit. It returns to the surface



The main hoist on the Marlow rig (*above*) supports the drill string and regulates downward force. The black hose supplies drilling mud, which is pumped down the bore of the pipe. The rotary table (*below*) turns the drill string, transmitting twisting force to a special length of drill pipe called the kelly.



A pair of mud pumps (*right*) drive the rig's circulatory system. Each pump has three pistons that force the mud down the bore of the drill string and back up the annular space surrounding the pipe. The white sheds in the background house the prime movers—the diesel-driven generators that power all the machinery on site.

The blowout preventer (*below*) is the most important safety apparatus on the drilling site. The preventer, installed at ground level under the drilling deck, is a heavy casing with several hydraulic rams that can seal off the well if high-pressure fluids underground threaten to escape. The main controls for the preventer are on the drilling deck, but auxiliary controls (*bottom*) are at ground level for use in emergencies. The red steel cylinders hold hydraulic fluid under pressure in case the hydraulic pump should fail.



through the annular space between the drill pipe and the wall of the hole or the casing. The crew keeps a close watch on the returning mud; it's an important source of information on what's happening at the bottom of the hole.

Layout of the Drilling Site. The rig and derrick are at the center of the drilling operation, but there is much else on the site as well. On the ground near one side of the rig structure are the prime movers—the engines that supply power for all the machinery. The prime movers are usually diesel engines, three or four of them, with a total power output of several thousand horsepower. The mud pumps, as major consumers of power, are placed nearby.

The mud pits take up most of the space on another side the rig. Yet another large area is occupied by the racks where hundreds of joints of drilling pipe are laid out. Other racks hold the larger and heavier pipe called casing. From the racks, joints of pipe or casing are dragged up a ramp to the drilling floor.

Elsewhere on the site is a row of trucks and trailers that provide office space (and sometimes living space) for the numerous consultants and contractors who are an essential part of the drilling operation. For example, one specialist provides the drill bits (known in the trade as tools). Another offers instruments and expertise in well logging—recording the geological strata the drill passes through.

There is more office space up on the main deck of the rig, in a structure amiably named the doghouse. It serves as lunchroom and gathering place for the crew, as well as housing a full set of controls for the rig.

Gushers and Blowouts. In the Hollywood version of the oil business, every well is a gusher. As crude thunders out of the ground, the crew celebrates in a rain of black gold, like boys playing in a mud puddle. A real gusher would be a disaster: the rig might be destroyed, fire is likely, people are injured or killed, and at the very least there is a monumental mess to be cleaned up.

The main line of defense against such events is the blowout preventer, a series of valves mounted at the wellhead, under the deck of the drilling rig. The valves close off the well as soon as it starts to “kick,” or show evidence of back pressure. A common sign of kicking is mud that continues to flow out of the well even when the pumps are shut off. The kick can usually be stopped by closing an annular preventer, which seals off the annular space around the drill pipe. If the annular preventer fails, there are ram preventers that crush the drill pipe or, if necessary, shear it off. Once the kick is controlled, the crew will mix a heavier mud and pump it into the hole.

The blowout preventers are operated by hydraulic pressure. The main control levers are on the drilling deck, but a second set of controls is placed some distance away for emergency use. In case the hydraulic pump fails, a battery of steel cylinders stores hydraulic fluid under high pressure.

Casing and Completion. After a well is drilled but before it goes into production, several more steps need to be taken.

First, the well must be cased and cemented. Casing is a steel pipe that lines the hole, keeping the oil out of overlying formations. In a deep well there will be several strings of casing, nested one inside the other, though only the innermost goes all the way to the bottom.

If you are in a petroleum-producing area and you see pipe being unloaded at a rail siding or hauled by truck, it is more likely to be well casing than drill pipe. The reason is that much more casing is needed, since each string of casing remains permanently in the hole, whereas drill pipe is reused many times. For the same reason, casing is a major item of expense in completing a well.

Casing is bonded to the well with cement. It must have taken a fair amount of courage the first time this procedure was tried. Having gone to the trouble and expense of drilling a well and setting casing, the crew then fills the casing with cement. If this cement hardened in place, it would totally plug the well. The trick is to pump in another fluid (such as mud) under enough pressure to force the cement out the bottom of the casing and up the annular space to the surface.

Even when the final casing is set and cemented, the well is not yet ready to produce. There is work yet to be done deep underground. Diagrams of wells often make it appear that the drill string pokes through the roof of a hollow cavity filled with liquid oil, but there are no such cavities deep in the earth. Oil and gas are actually dispersed in microscopic pores within a matrix of solid rock. Fluids move so slowly through this material that a newly drilled well could recover oil and gas only from within a few feet of the well bore. For drilling to make economic sense, this range has to be extended, allowing a larger volume of oil and gas to flow into the well. The process is known formally as well stimulation or well development, but workers in the oil field call it fracking (short for *fracturing*). Passages through the rock are opened up by explosions, by acids, or by surges of high-pressure water; then the cracks are held open by pumping in sand or a synthetic “proppant” of crush-resistant spherules.



Another Chesapeake Energy well near Marlow had been drilled but was not yet in production when the photograph above was made. The wellhead, at left, still has a blowout preventer in case any high-pressure fluids are encountered. The apparatus at right is a sand separator, used to remove excess sand or “proppant” from the product stream. The proppant (*below*) consists of crush-resistant synthetic spherules that are forced down the well to prop open cracks in the oil- or gas-bearing rock layers.



After fracking, a well might tap the hydrocarbon resources within a few hundred feet of the well bore. This is still a small area—perhaps 10 acres. It is sobering to reflect that the oil and gas recovered from under such a small patch of ground can be



The offshore drilling rig Rowan Gorilla III is seen in an outfitting yard along the Sabine River between Texas and Louisiana, where many such rigs are built and maintained. It is a jack-up rig: the platform moves up and down the tall, girder-like legs. When this photograph was made, in 2001, new legs were being constructed. At one time, the Rowan Gorilla III was the world's largest jack-up rig. For a sense of scale, note that the basket dangling from a crane on the left side of the rig holds two workers, just visible in silhouette.

worth enough to repay the multimillion-dollar cost of drilling the well. But there is another side to the economics of oil. Most existing wells (as opposed to newly drilled ones) have a much smaller output. The national average is just 11 barrels per day, and thousands of wells produce less than a barrel. Even when the price of oil is up around \$50 a barrel, owning an oil well is not guaranteed to make you a millionaire.

Offshore Drilling. Boring a deep hole in the ground is hard enough on land; it's even harder when the ground is hundreds of feet beneath a hostile sea. Offshore drilling is therefore worth the cost and the risk only when the potential payoff is very large.

In the Gulf of Mexico and in several areas along the California coast, drilling operations begin not very far offshore: you can see drilling platforms from land or reach them with a small boat. For that matter, you can see a lot of the equipment onshore, where rigs are brought into port for maintenance. The machinery is so big you don't have to get very close to get a clear look.

Exploratory offshore drilling is done with a mobile drilling unit. These barges and ships and specially built semisubmersible rigs remain floating over the hole as drilling proceeds. The vessels are either anchored in place or held on station by a dynamic positioning system, which uses computer-controlled thrusters to correct any drift. In shallow water another kind of mobile unit is popular: the jack-up rig. It is a barge with three or four tall girder-like legs that slide vertically through openings in the hull. The rig is floated to the drill site with the legs raised; then they are lowered to the sea floor and the hull is jacked up above water level.

In proven oil and gas fields, drilling is done from permanent platforms, which also serve as production platforms once all the wells have been completed. These platforms are much larger structures than mobile drilling units. Some have a frame of hollow steel tubes, which is floated to the site on its side and then turned upright and sunk by flooding the tubes. The tallest structure in the world is a platform of this kind, called Bullwinkle, 150 miles southwest of New Orleans in the Gulf of Mexico. It towers 1,600 feet over the sea floor and almost 300 feet over the sea surface.

Wells in Production. After all the commotion of drilling is over, a producing well is a lonely place. A few pipes and valves poke out of the ground, and there may be a pump, but there is seldom anyone around to tend this equipment. An employee known as a lease pumper is supposed to stop by once a day to check the machinery and note the amount of oil produced, but automated instruments for remote monitoring are making even these brief visits unnecessary.

A free-flowing well—one with enough pressure underground to push the oil or gas to the surface—is fitted with a tall stack of valves and gauges called a Christmas tree. Why so many valves? They serve various functions—taking samples, regulating flow and pressure, providing an emergency shutoff—but there's also deliberate redundancy. If a valve fails and needs to be replaced, the job is a lot easier if there is another valve below it that can close off the flow while repairs are made.

A producing well near Marlow, Oklahoma, is fitted with a Christmas tree: a stack of valves for controlling the flow of oil and gas that come out of the earth under their own pressure. On this well the upper valves regulate flow through the well's central tubing; some of the lower valves give access to the annular space between the tubing and the well casing.



A sucker-rod pump extracts oil from a well when subterranean pressure is insufficient to bring it to the surface. The pump's raised beam (called a walking beam) rocks back and forth, driven by the motor and linkage at left. This nodding motion is transmitted to the vertical string of steel sucker rods, which extend all the way to the bottom of the well. A pump chamber at the bottom lifts oil on the upstroke and refills on the downstroke. The pump in this photograph is on farmland near Plainville, Kansas.

The pressure that drives oil to the surface usually dissipates long before the oil itself is exhausted. Thereafter, the oil must be lifted artificially. The standard device for doing this is a marvellous contraption with a marvellous name: the sucker-rod pump. The business end of a sucker-rod pump is far below the ground, installed in the tubing at the level of the oil-bearing formations. The basic components are a cylindrical chamber and a pair of one-way valves, arranged so that the chamber fills on the downstroke and lifts the oil on the upstroke. The pump chamber is driven by the sucker rods—a string of rigid steel bars that run inside the production tubing. Like drill pipe, the rods come in joints about 30 feet long, but they are solid rather than tubular and only a half inch to an inch in diameter.

The topside part of the sucker-rod pump is a distinctive sight in any oil-producing region. The sucker rods are connected by a short length of wire rope to one end of a walking beam, which is held 10 or 12 feet off the ground and pivoted in the middle. The other end of the beam is driven by a crank arm, so the beam rocks about its pivot point, periodically raising and lowering the sucker rods. A large, curved guide at the sucker-rod end of the beam is called the horse's head, and the name couldn't be more





A battery of tanks near the well site holds oil pending shipment to the refinery. Not all the tanks are for oil. Storage is also needed for the salt water that comes out of the well.

apt, for it is horselike not only in form but also in its nodding motion. The face of the horse's head is an arc of a circle, and so it converts the rotary motion of the walking beam into linear motion of the sucker rods. It is important that the rods move straight up and down into the well, without any bending stresses.

The power source for a sucker-rod pump is usually an electric motor, but there are other possibilities, particularly for wells in remote areas. If a well produces significant amounts of natural gas, some of it can be used to fuel a small engine.

Sucker-rod pumps run at a stately, unhurried pace—typically 10 to 20 strokes per minute. Often you'll see a pump standing idle, but that doesn't necessarily mean the well has been shut down. Most pumps run intermittently—a few hours on, a few hours off—to avoid distorting the distribution of oil, water, and gas in the underground formations. Overpumping can permanently impair a well's production, and can even affect neighboring wells.

Up close, a sucker-rod pump makes a moaning or sighing noise with a slightly syncopated rhythm, since the upstroke puts more load on the machinery than the downstroke. In the evening the sound has a lullaby quality to it. The whole apparatus looks like a holdover from the age of steam. In fact it's not quite that old, but the design has changed little since it was introduced in the 1920s. Some individual pumps from that era are still running.

Field Processing of Oil and Gas. The fluids that come out of an oil well are not ready to be poured into your gas tank. They are not even ready to be pumped to the refinery. The stuff that comes out of the well doesn't qualify as crude oil until some preliminary field processing is done. The equipment for this processing is sometimes at the wellhead; more often a single processing station serves a cluster of nearby wells.

A heater-treater at one of the Marlow wells has the job of separating gas, water, and oil, which come out of the well mixed in a frothy emulsion. Under heat and pressure, the gas bubbles out of the mixture, while the denser water settles to the bottom.



Something to notice about a field-processing station is the bed of gravel that surrounds each piece of equipment. In the event of a small leak, the absorbent gravel reduces the fire hazard. It also makes leaks easier to spot.


The main task in field processing is to separate oil, gas, and the salt water that inevitably comes up out of the well along with the more valuable fluids. The simplest kind of separator is a settling tank with one inlet pipe and three outlets—at the top for gas, in the middle for oil, and at the bottom for water. A characteristic feature of these gravity separators is a sight glass used to gauge the water and oil levels. Gravity

SWEET-AND-SOUR HYDROCARBON SOUP

Crude oil comes in many varieties, some of which sound like the names of exotic coffees: Louisiana Sweet, Nigerian Bonny Light, Jobo Crude, Suez Blend, Escalante, Oriente, West Texas Sour. Every oil field produces something a little different.

The classification of crudes depends on two factors: heavy versus light and sweet versus sour. Heavy crudes are viscous, tarry, and dense; they have a higher concentration of the “bottoms” that give refiners fits. The sweet and sour crudes have nothing to do with soups served in Chinese restaurants. The difference between them is that sour crudes have more sulfur. I’m told that the terms came about because people really did taste the oil to detect sulfur.

Petroleum is not a single chemical substance but a complicated broth with many ingredients. Most of the molecules are hydrocarbons: little assemblies of carbon and hydrogen atoms (shown here as black balls and




white balls respectively). The carbon atoms are linked to one another in chains and necklaces, adorned with hydrogens. The rule for forming the molecules is that every carbon atom has four bonds, or places where it needs to be

attached to something else—either another carbon atom or a hydrogen atom. A hydrogen atom has just one bond.

The simplest hydrocarbon is methane, in which a solitary carbon is festooned with four hydrogens. Ethane has two carbons holding hands, with six hydrogens filling up the rest of the bonds. Propane and butane are chains of three and four carbon atoms, again with hydrogens stuck everywhere there’s a free bond.


Organic chemistry has a wonderful vocabulary with the unusual property that once you’ve



learned the full name of a molecule, you’ve also learned all about its structure. (What if people’s names were similarly informative?) This is not the place for a discourse on hydrocarbon nomenclature, but one basic principle is worth knowing. Beyond the four smallest molecules (methane, ethane, propane, and butane), the names are easy to remember if you know how to count in Greek: pentane has five carbon atoms; hexane, six; heptane, seven; octane, eight; nonane, nine; and decane, ten.

Not all the hydrocarbons are linear chains, with carbons lined up like beads on a string.

Among the four-carbon molecules, the straight-chain version is called normal butane, but there is also a branched version, identical in composition, called isobutane. Larger molecules have many possible isomers, or rearrangements. A



particularly important one in the petroleum industry is isooctane, which has eight carbon atoms in a lopsided cruciform shape. Isooctane is the gold standard for gasoline. The octane number posted on the gas pump is defined as 100 for pure isooctane, and other fuel constituents are ranked in comparison with it.

In addition to forming straight and branched chains, hydrocarbons can also fold up on themselves to form ring-shaped molecules.

Cyclopentane has five carbons in a closed ring. Ring molecules in one important family are called aromatics (because many of them have strong aromas). Benzene is the prototype of this group; it is usually represented as a six-carbon ring with alternating single and double bonds. Both the cyclic and the aromatic rings are present in gasoline, and benzene is a major constituent of some heavier products, such as jet fuel and diesel fuel.



A fully developed oil field seen from far overhead is a dense netting of wells, roads, and pipelines embroidered on the landscape. Each of the small white squares in this photograph is the gravel pad of a well and its associated equipment. As a rule, there can be no more than one well for each 40-acre parcel of land, which accounts for the regular spacing. The bright white lines threading through the array of wells are access roads; there is also a fainter and less orderly network of pipelines that gather the oil. The area shown is the Wasson oil field in west Texas, where drilling began in the 1930s. The gray patch near the bottom of the frame is Denver City, Texas. The photograph was made by an astronaut aboard the International Space Station from an altitude of about 200 miles.

separators are very slow, and thus, if you see one, it is probably safe to assume that the wells it serves are producing at a low rate.

Higher-capacity separators operate under pressure. Like other pressurized vessels, they can be recognized by their beefier pipe fittings, by their rounded form, and by the presence of pressure gauges and relief valves. One type is a long pressure vessel filled with hundreds of baffle plates where droplets of liquid collect and then drain to the bottom. The well stream enters at one end of the separator, and outlets for gas and liquids are near the opposite end. A variant has two horizontal barrels stacked one above the other; the lower barrel accumulates the liquids.

Separators work quite well for dividing gas from liquid, but they are less efficient in removing water from oil. Often, further treatment is needed to dry the hydrocarbons. Crude oil is not considered saleable if it has more than 1 percent water, and dehydration requirements for gas are even more stringent.

Crude oil and water are difficult to separate because they form a frothy emulsion, with the water dispersed in microscopic droplets. One way of breaking up the emulsion is to heat it, which causes the droplets to coalesce in the same way they do in an overheated béarnaise sauce. The usual device for this purpose is called a heater-treater. It can be either a vertical or a horizontal vessel with a burner underneath, typically fired by natural gas from the well. The stack for flue gases is a distinctive feature. As the heated emulsion breaks down, oil is drawn off from the upper part of the chamber and water from below.

The drying of gas is more of a safety issue than a commercial one. At high pressure, water vapor combines with some of the hydrocarbons in natural gas to form solids called hydrates, which can plug up valves and pipelines. Because hydrates are never seen in everyday life, they have an air of the spooky about them, and yet they



Pipeline bridges (*above*) are one of the few places where petroleum pipelines emerge from underground. Both of these bridges carry pipelines across the Colorado River near Needles, California. The one at right was built as a highway bridge (for U.S. Route 66); the much lighter suspension structure above was designed from the start for pipeline duty. Even where pipelines run underground, surface markers reveal their route. The marker below is along Black Bayou in southern Louisiana.



are more than a phantom menace to the gas industry, which expends much effort to suppress their formation.

The salt water separated from the oil and gas, called produced water, is quite a nuisance. It is corrosive, which makes it hard to store, and it is contaminated not only with salts but also with petroleum residues, which means you can't just dump it in the nearest ditch. A common solution is to put it back where you got it, by pumping it into another deep well. If you can arrange to pump it into the bottom of an oil-bearing formation, it will help to drive the oil up and out.

PIPELINES

Oil is sold by the barrel, and in the beginning that's how it was shipped too. Horse-drawn wagons hauled barrels to the railhead, where they were stacked on flatcars. Today the barrel remains a unit of measure (in the petroleum industry it's equal to 42 gallons), but oil, like cheap wine, never sees the inside of an actual barrel. Bulk shipments move over water by tanker and over land by pipeline.

The pipeline system is organized like a tree. Small collector pipelines in the oil field, called flow lines, are the fine roots of the system. They gather crude oil from many wells and bring it to the field processing station. Somewhat larger pipes carry the oil to the terminus of a main-line pipeline, which supplies refineries hundreds of miles away; this is the trunk of the tree. The products of the refinery are then distributed through another system of main-line pipes, which divide into smaller and smaller branches until they reach distribution depots—the leaves of the tree. The natural-gas pipeline system has a similar architecture, except there is no refinery, and the final twiglike branches go all the way to individual homes.

The pipe used for petroleum and natural-gas transport is made of high-strength steel. Joints 40 feet long are welded together in the field, and then the welds are x-rayed to check for defects. High-quality construction is critical because the pipeline



operates under high pressure, and a leak can be explosive. Because pipelines are buried, you won't see the steel pipe along most of its length, but the pathway is not hard to spot. There should be an identifying marker or warning sign wherever the pipeline crosses a road. One place where the pipeline may emerge from the concealment of its burrow is at a river crossing. Sometimes the pipeline is carried on a highway or railway bridge, but there are also bridges specially built for pipelines alone. Some of these structures are elegant and spare suspension bridges.

Corrosion is the great enemy of petroleum pipelines. It's not simply that the steel pipe rusts in the ground; chemical interactions with the soil set up electric currents that actively eat holes in the steel. To slow corrosion, the pipe is coated with various

HOT OIL FROM THE ARCTIC

The most spectacular pipeline—and the most controversial—is the one that carries crude oil from Prudhoe Bay, on the shores of the Arctic Ocean, 800 miles across Alaska to Valdez, on the southern coast. What's most unusual about the Trans-Alaska Pipeline is that you can see it. Over about half its length, the pipeline is not buried but instead is raised on stilts.

The reason for the elevated style of construction is that the arctic soil is permafrost: below a thin surface layer, the earth remains frozen year round. The oil running through the pipeline has the consistency of hot fudge, and the temperature of it too—up to 180 degrees Fahrenheit. If this scalding fluid were pumped through the permafrost, the result would be a boggy mess.

Even with the pipeline propped up off the ground, the builders had to take extraordinary measures to keep the foundations frozen. Ordinary steel or concrete posts would have conducted heat down into the soil. To keep the permafrost permanently frosty, the posts are fitted with "heat pipes" that actively transport heat upward. At the bottom of each post is a reservoir of liquid ammonia. If the base warms, the ammonia boils, extracting heat from the soil and carrying it upward; at the top of the post the heat is radiated away by fins; then the ammonia condenses again and dribbles back into the reservoir. There are 76,000 of these

supporting posts along the route, and most of them extend 50 feet into the subsoil. (More drilling was needed to set the posts than for all the oil wells at Prudhoe Bay.)

The pipeline is four feet in diameter, and the insulation wrapped around it adds another foot. At any given moment during normal operations there are nine million barrels (or 400

pipeline had to go under instead: at 24 "sag bends," a section of the pipe is buried in a specially prepared bed that's able to withstand cycles of freezing and melting. In a few cases a refrigeration plant pumps chilled fluid through the soil to keep it frozen.

In the first few weeks after the pipeline was filled, there were several leaks as well as a



million gallons) of petroleum in transit through the pipeline—enough to fill several tankers.

Over most of its route the pipeline is in wilderness terrain, and a big question was how herds of migrating animals would cross the pipeline route. The first proposal was to lift the pipe high enough to let animals pass under. This plan was fine for moose, but caribou are too wary to go under such a bridge. So the

major fire. The years since have brought further incidents but no major catastrophes. As is well known, the worst accident with Alaskan oil happened not on the pipeline but in Prince William Sound, where the *Exxon Valdez* was sailing through a deep, wide, and straight passage but nonetheless managed to run aground. (Photograph reproduced courtesy Bureau of Land Management.)



An anticorrosion device along the route of a petroleum pipeline in southern Arizona generates a small voltage opposing natural electric currents that would eat away the metal of the pipe. The circuitry for the protective device is inside the metal box; the black-and-white disk above the box spins when the current is flowing, as an indicator that the system is working.

kinds of gunk and wrapped with tape. A further measure is to plant “sacrificial anodes”; these are metal plates that protect the pipeline by allowing their own substance to be etched away. (You have one in your water heater, for the same purpose.) If electric power is available along the right-of-way, small transformers and rectifiers can feed direct current into the ground as another way of combating corrosion.

If you have ever turned on a faucet and been annoyed waiting for the water to run hot, consider how long it can take for a parcel of oil to make its way through a transcontinental pipeline. I was surprised to learn that oil does not gush through a pipeline in a high-velocity torrent, like water from a fire hose. The oil moves at a comfortable walking pace—between three and five miles an hour. Thus, a batch of oil pumped into one end of a 1,000-mile pipeline takes a week or two before it comes out at the far end. The quantity of oil in transit in the largest pipelines is greater than the storage capacity of the largest tank farms. The Trans-Alaska Pipeline, which is four feet in diameter and more than 800 miles long, holds nine million barrels.

Pipelines for refined petroleum products tend to be smaller than crude-oil lines, but their operations are more complicated. The products pipeline may run a few thousand barrels of gasoline, then some kerosene and some diesel fuel before more gasoline. Cutoff valves have to be turned at just the right moment to divert each batch to the proper destination; it’s like sorting railroad cars in a moving train, except that you can’t see the train.

Don’t the various products get all mixed up as they flow through the pipeline? In most cases the mixing is confined to a small portion of each batch, and it doesn’t cause much trouble. If two products are seriously incompatible, an inflatable sphere can be popped into the line to keep them separated.

Because their operations are strung out over such long distances, pipeline companies have always been on the leading edge of communications technology. Even the very first crude-oil pipeline, built in 1865 in western Pennsylvania and running a



A pumping station near Deming, New Mexico, boosts the pressure in a long-distance gas pipeline. The pumps (or compressors) are inside the metal-clad building; the tall cylindrical structure is a heater-treater for moisture removal.



A pig launcher stands out above other equipment in the pipe yard of a pumping station near Houma, Louisiana. A pig is a device sent through the pipeline to clean or inspect it. The spherical bulge in the elevated section of pipe is a valve that isolates the launcher tube. With the isolation valve closed, a cap is removed from the end of the launcher tube, and the pig is inserted; then the cap is replaced and the isolation valve is opened. Fluid diverted through the launcher tube drives the pig out into the pipeline through the downward-sloping tube.

total distance of six miles, had telegraph wires strung parallel to the pipeline itself. The telegraph was replaced successively by the telephone, the teletype, and microwave radio links. Now many pipelines have fiber-optic cables buried along the right-of-way. One pipeline company abandoned petroleum transport altogether and went into the telephone business—then came back again.

Pumping Stations. Pipeline transport is more energy efficient than shipping the oil by either rail or road, but it still takes hundreds of thousands of horsepower to push the oil through the pipe. The pushing can't all be done from the starting point; booster stations are needed every 60 to 100 miles.

The pumps themselves are usually roofed over to protect them from the weather, but sometimes the end walls of the pump building are left open to prevent explosive vapors from accumulating. Outside the pump building is a neatly arranged but still complicated-looking array of brightly painted pipes, mounted on stanchions that hold them a foot or two off the ground. There are valves everywhere to control the flow of oil through the system.

Another item to look for in the pipe yard is a pig launcher. Animal lovers can rest easy: as far as I can tell, no one has ever tried to run a live pig through an oil pipeline. A pig, in the pipeline trade, is a big brush or scraper or squeegee that gets pushed through the line to clean the inner surface. Most of them are shaped more like bullets than pigs; some look like oversize toilet brushes. I'm told the name comes from the squealing sound that a pig with metal scrapers makes as it passes through the pipe.

The pig launcher is a short section of pipe arranged as a branch from the main line, like the on-ramp of a highway. The several valves controlling the launcher are initial-

Four pumps driven by 6,000-horsepower motors occupy an open-ended shelter in Port Fourchon, Louisiana. The station pumps oil received at the Louisiana Offshore Oil Port (LOOP), a docking facility 20 miles out in the Gulf where large tankers unload. A submarine pipeline delivers the oil to the Port Fourchon pumping station, which then moves it farther inland.



A tank farm associated with a refinery in Philadelphia dwarfs the refinery itself (the small area at upper left with plumes of water vapor rising on a winter day). The farm includes both cone-roof tanks and floating-roof tanks. (Most of the latter appear to be empty.) The white tanks white are designed for volatile fluids; by reflecting solar energy, the white paint helps reduce evaporation. The black tanks are meant for more viscous materials, which warmth from solar energy helps to keep flowing.



ly set to seal and drain the launcher barrel, so the pig can be loaded through a door at the end of the barrel. Then the valve settings are changed so the flow of oil drives the pig into the pipeline and carries it along to the next pumping station. There it is retrieved in a pig trap, which looks the same but points in the opposite direction.

The rubber spheres that separate product batches are launched and retrieved with the same kind of equipment. And these days there are also “smart pigs” that carry instruments through the pipeline to check for corrosion and other flaws.

TANK FARMS

Rather than *tank farms*, it might be better to call them *tank orchards*, for the tanks are arranged in neat rows and columns like fruit trees. Tank farms are found at the terminals of pipelines, at refineries, and at tanker ports. There are smaller collections of tanks—perhaps we should call them tank gardens—at the depots where petroleum products are stored for retail distribution.

Seen up close, a tank farm seems immense; from another perspective, however, the tanks are amazingly small. Given the world’s thirst for oil, they hold only a few weeks’ supply. A typical refinery has enough tank capacity to store two weeks’ worth of crude stocks and four weeks’ worth of refined products. It seems petroleum molecules are like those insects that emerge from the earth, live and breed for a day, and then perish. Oil is burned up no more than a few weeks after it comes out of the ground.

Floating-roof tanks are the usual storage vessels for crude oil and for gasoline, both of which have volatile constituents that would evaporate into the empty space above the liquid surface in a conventional tank. In these tanks, at a marine terminal in Trieste, Italy, spiral stairs climb the outside of the tank, then telescoping stairs descend to the floating surface of the roof.





Berms surrounding oil tanks create a gently rolling landscape at a pipeline terminal in Greensboro, North Carolina. The berms offer both fire protection and pollution protection. They are sized to confine any possible leak from a tank.

Most petroleum tanks are cylindrical. Their proportions tell an interesting story. The smallest tanks are taller than they are wide—perhaps 20 feet in diameter and 30 feet high. Going to larger sizes, the diameter increases much more than the height, so the tanks become fatter and flatter. The very largest are 400 feet in diameter and 60 feet high (big enough to house a respectable professional sports arena). The reason for making big tanks shallower is that the internal pressure at the base of the tank depends on the height of the fluid inside, not the total volume. Higher pressures call for stronger and more expensive walls.

The walls are fabricated from curved steel panels welded together in a bricklike pattern. The most common size for the plates is 8 feet by 30 feet. (Knowing this, you may be able to estimate the size of a tank by counting panels.) Because the pressure is greatest at the bottom, the lower plates are thicker—an inch to an inch-and-a-half—tapering to about half an inch thick at the top.

Cone Roofs and Floating Roofs. The simplest cylindrical tanks have a conical roof, usually supported internally by a post in the center, like the pole of a circus tent. But tanks of this kind are suitable for only a few petroleum products, such as diesel fuel and home heating oil. These are liquids with a low vapor pressure: they have little tendency to evaporate. In contrast, crude oil and gasoline include volatile components whose fumes would fill the vacant space in a partially filled cone-roof tank. The vapors would then be driven out through the roof vent whenever the tank was filled. The loss of the volatiles would be costly, and would also be a source of air pollution.

The secret to storing gasoline and crude oil is the floating-roof tank, in which the roof is a pan or a buoyant platform floating on the surface of the stored liquid. Because there is no vacant space, very little of the liquid can evaporate. As the tank is filled and emptied, the floating roof slides up and down inside the shell. At the perimeter of the roof are seals or gaskets that wipe against the inner surface of the shell.

Firefighting provisions at the Trieste tank farm include permanently installed water cannons as well as berms. The red panels on the walkway around the perimeter of the tank are also firefighting stations.



Among all geometric forms, a sphere is the best adapted to resisting pressure, since outward force is exerted equally on all parts of the vessel. In the petroleum industry spherical tanks typically hold propane and similar light molecules that remain in liquid form only if they are kept under pressure of a few hundred pounds per square inch. Berms are not needed around such tanks, because any leaks would be released in gaseous form. A wind sock is more useful; it indicates which way to run in an emergency. These tanks are at the ChevronTexaco refinery in Pascagoula, Mississippi.



From above, floating-roof tanks are easy to recognize, since the roof platform is likely to be well down inside the hollow shell. Even from ground level, you can identify a floating-roof tank if you know what to look for. There is usually a flange, called a wind girder, that circles the tank near the top; on the largest tanks it is a walkway with a handrail. The wind girder adds stiffness to the shell; it is not needed in a tank with a conventional roof, since the roof itself stiffens the structure.

A problem with the floating-roof design is that it not only confines liquids under the roof but also allows rainwater to accumulate on top of the roof. How do you drain a roof that's lower than the surrounding walls? The usual solution is a flexible drainpipe that gathers rainwater from the middle of the roof, carries it through the petroleum compartment, and exits through the wall of the tank near its base. Another solution is to put a floating roof inside a tank that also has a conventional cone roof.

Pressure Spheres. The most distinctive tanks are the spherical ones, which hold liquids that have to be put under pressure to keep them from boiling away. (Propane and butane are the main petroleum products in this category.) The spherical form is chosen because it offers the best resistance to internal pressure, which can reach 250 pounds per square inch.

Spherical tanks are supported on piers that reach up to near the equator of the tank, like a baseball resting on fingertips. They cannot sit on the ground the way a flat-bottomed tank does, not only because they might roll away but also because all the weight would bear down on a single point. It so happens that a sphere is the most efficient of all shapes for a tank, in the sense that the largest volume of fluid is enclosed by the smallest quantity of steel. Nevertheless, spherical tanks are more expensive than cylindrical ones because the fabrication is so tricky. Many spherical tanks have a spiral staircase that follows a particularly elegant compound curve to the top.

Fire Berms. If a tank farm isn't an orchard, maybe it's a rice paddy. Notice that each tank sits in the middle of a large square plot, surrounded by an earthen berm three or four feet high. Pipelines, roads, or walkways that cross the berm go up and over rather than under or through. The area enclosed by the berm is calculated to hold the entire contents of the tank in the event of a major leak.

THE REFINERY

A petroleum refinery is the *ne plus ultra* of industrial landscapes. The maze of pipes, towers, vessels, stacks, and fuming vents defies comprehension. Looking at all that intricate plumbing, you cannot possibly trace the path of any given molecule. Still, it's possible to make some sense of the chemistry happening in a refinery. There are two important aids to understanding. First, some kinds of equipment are used over

An oil refinery suggests the image of a metropolis for hydrocarbons, the pipe manifolds like expressways, the distillation towers like skyscrapers. The refinery depicted here, in Rodeo, California, northeast of San Francisco, was one of the first on the West Coast of the United States. It has changed ownership several times; at last report it was operated by ConocoPhillips.



and over throughout the plant (although no two structures are exactly alike). Thus, if you recognize a feed heater or a fractionating tower in one place, you'll know it wherever you see it. Second, the refinery is organized into distinct units, each with a specific function. Although you may never be able to deduce the purpose of every pipe and pressure vessel, you might be able to identify some of the major units.

Petroleum refining is the prototypical process industry—a style of manufacturing where things are made in continuous streams rather than discrete batches. Oil could be refined one vat at a time, the way you make soup in a pot, but refinery engineers always try to avoid that mode of operation. Whenever possible, they design a process so that raw materials flow in steadily at one end and products come out the other.

Feed Heaters. Most of the processing steps in a refinery require high temperature—often a few hundred degrees Fahrenheit (enough to bake bread), sometimes a thousand degrees or more (enough to peel paint). For this reason, just about every unit has a furnace, or feed heater, where the incoming fluid is brought up to temperature.

The heaters have a distinctive shape, with sloping shoulders leading up to a tall metal smokestack, somewhat similar in form to an old stone fireplace with its chimney—but the hearth is the size of an entire house. At the bottom of this structure is a burner. The walls are lined with steel pipes, through which the feedstock is pumped.

A refinery might have a dozen of these feed heaters. The fuel burned in them is a mixture similar to natural gas, drawn from the refinery stream itself.

Fractionating Columns. Crude oil is a mixed-up stew of hundreds of chemical compounds. The refinery's first task is to sort them into groups according to molecular size. A fractionating column, or tower, is where the sorting gets done.

The tower, together with its feed heater, is a fancy kind of distillery. It separates substances according to differences in their boiling point. Inside the column are perforated baffles, called trays, stacked up one above the next at intervals of a few inches or a few feet. When the hot feed liquid is pumped into the bottom of the tower, much of it boils away. The vapor begins rising through the perforated trays, but meanwhile other fractions are condensing into a liquid and trickling back down through the same trays. As these counterflows continue, the most volatile molecules accumulate near the top of the tower while the heaviest sink to the bottom. Each intermediate component finds its own natural level.

Fractionating towers come in a range of proportions, from squat and tubby with a pronounced midriff bulge to extremely tall and svelte. A refinery could have dozens of columns. They are the objects that stand out most clearly against the horizon and give the refinery its characteristic skyline.

Reactor Vessels. Not every tall, metal, cylindrical structure with pipes around it is a fractionating column. There are also large vessels with other purposes, such as housing chemical reactions. Some of these reactors are upright and oblong, like a distilling

A feed heater at the ChevronTexaco refinery in Pascagoula, Mississippi, heats crude oil to about 700 degrees Fahrenheit. There are gas burners near the bottom of the unit; tubes lining the walls carry the feedstock to the heater. This furnace and another just like it at Pascagoula process 325,000 barrels of oil per day.



tower, but they are not as tall. Other reaction vessels are laid on their side. Many of them are built to withstand high pressure (up to 5,000 pounds per square inch, which is roughly the pressure at the bottom of the ocean). The pressure vessels have steel walls as much as six inches thick—not that you can tell from the outside. But there is a tell-tale sign of high pressure: hemispherical end caps give the vessel the shape of a medicine capsule. It's called a bullet tank. Vessels that don't have to hold high pressures have blunter ends, closer to the shape of a soup can.

Heat Exchangers. At many places in a refinery, one stream of fluid needs to be heated and another needs to be cooled. A process engineer will never miss a chance to swap some heat in a situation like this. The two fluids are run through a heat exchanger. From the outside, a heat exchanger is just a long cylindrical tank, usually laid on its side. Inside is a big bundle of tubes. One fluid is pumped through the tubes; the other passes through the surrounding space, called the shell. Heat flows through the walls of the tubes from the warmer fluid to the cooler one.

Pipes, Pumps, and Valves. Pipes knit together the fabric of the refinery, carrying the various raw materials, products, and intermediate stocks, as well as steam, water, pressurized air, and exotic fluids such as hydrogen sulfide and ammonia. The racks of piping may look chaotic, but there is nothing ad hoc about their design; there are thick volumes of standards and specifications, and the piping engineer must have a lawyerly mastery of all this literature.

Many pipes are insulated. Some are heated with steam. A feature of these pipelines is an expansion loop, which looks like the hump in an inchworm. Expansion loops absorb the strain when the pipe expands or shrinks in response to temperature changes.

For every pipe, there's a valve—or a bunch of them. One engineering manual estimates that 8 percent of the equipment budget for a refinery goes into valves. There are two broad categories. Stop valves are meant to be either fully open or fully closed, like light switches. Throttling valves offer continuous control over the rate of flow. These days, most throttling valves are remotely operated from a central control room.

Pumps in a refinery have a prominent place in the minds of the operators, but they are hard to spot from a distance. They are usually mounted at ground level or below because they need to be lower than the vessel they draw fluid from. Pumps consume most of the refinery's electric power.

The Crude Unit. The crude unit is where the refining process begins. In the early years of the petroleum industry, it was also where the process ended. That is, the crude unit was the entire refinery.

The crude unit is a distillery with three main parts: a feed heater, an atmospheric-pressure distillation column, and a vacuum distillation column. This last item is also called a flasher. The atmospheric column is taller and slimmer; the flasher has a distinctive shape, with a bulge in the middle.



Fractional distillation column (*above*) is the most characteristic feature of a refinery. The catwalks at various levels provide access for maintenance. The column is at the Premcor refinery in Port Arthur, Texas. Insulated, bullet-tank reactors (*below*) are at the ExxonMobil refinery in Chalmette, Louisiana.



In the feed heater, the oil is brought to a temperature of 700 degrees Fahrenheit, then it is pumped to the atmospheric tower. In the molecular sorting process that goes on inside this column, the four lightest hydrocarbons—methane, ethane, propane, and butane—float to the top, where they are drawn off as overheads, or light ends. The fraction withdrawn at the next lower level is called straight-run gasoline. A century ago, straight-run gasoline was all the gasoline there was. Now quite a lot else gets blended into the product, and straight-run gasoline is a minority constituent.

Below the straight-run gasoline tap, several more fractions are drawn off: naphtha, kerosene, and something called gas oil. Then, at the bottom, there's the even heavier



One of the crude units at Pascagoula has two main components (apart from the furnace). The atmospheric distillation tower, in the background at right, operates at normal air pressure and separates the lighter components of the oil. The vacuum distillation tower, or flasher, in the foreground, boils heavier oil fractions at reduced air pressure.

stuff that just won't boil under these conditions; it's called resid (short for residual oil), or simply bottoms, and it is the problem child of the refinery. If you try to boil it away by raising the temperature further, the oil breaks down chemically, like sugar turning to caramel. But some of the resid will boil if you keep the temperature constant while lowering the air pressure (just as water boils a little easier in Denver than in Dallas). That's what the vacuum tower is for. The air pressure inside is about a third of normal.

The fractions boiled off in the vacuum tower are known as flasher tops. They are a raw material for the manufacture of motor oils and other lubricants. But if there are flasher tops, there must also be flasher bottoms! This is the residue of the residue—the gunk that won't boil even in a partial vacuum. Some of it can be made into asphalt for road paving and roofing.

The Gas Plant. The refining of the lightest fractions is the mirror image of what happens to the heavy resid. Where lower pressure helps the heavy stuff boil, higher pressure helps the light stuff condense. The gas plant has several tall and very slender towers. Height is needed because the light gases are hard to separate from each other, and so the columns must have a large number of trays. The towers can be narrow in cross section because the volume of material being handled is much smaller than it is in the crude unit. The pressure is about 200 pounds per square inch.

Four products come out of the gas plant. Methane (which is also the main ingredient of natural gas) fuels most of the refinery operations. Ethane can also serve as fuel, but it's worth more as a raw material for making other chemicals, including plastics such as polyethylene. Propane is marketed as liquefied petroleum gas. Butane is the clear liquid in all those transparent plastic cigarette lighters, but that's not where most of it goes. It's also a vital component of modern gasolines.

The Catalytic Cracking Unit. Cracking is just what the name suggests: Taking big molecules and snapping them apart into smaller pieces. For example, a chain of 16 carbon atoms might be cracked into a 10-carbon unit and a 6-carbon unit. The aim of this process is to convert some of the heavier petroleum fractions, for which there is little market, into more gasoline. The molecular bone-breaking could be done by heat alone, but a catalyst—a chemical facilitator—makes it happen faster and with better control over the outcome. Catalysts for this process are mainly zeolites, which are lacy molecular structures with many voids where hydrocarbons can lodge.

For some reason, "cat cracking" is better known to the public than other refinery operations. But I had a wrong impression of how it's done until a visit to a refinery clued me in. I had imagined feedstock trickling through a big vessel packed with beads of catalyst. That type of cat cracker, I learned, has been obsolete for decades.

The cracking process favored today is called fluid catalytic cracking. The catalyst isn't actually a fluid; it's a solid, but it's ground into a powder so fine that it flows. The catalytic powder and the feedstock are mixed at a temperature of about 900 degrees Fahrenheit and pumped upward into a reaction chamber. The chemical cracking

The fluid catalytic cracking unit at Pascagoula includes a reactor (*left*) where all the hydrocarbon chemistry takes place, and a much larger regenerator (*right*), where the catalyst is restored to its active state by a blast of hot air that burns away carbon deposits. (Photograph courtesy Stephen Renfroe, ChevronTexaco.)





Hydrocracking (*above*) is done in tall pressure vessels, where the feedstock and added hydrogen can percolate through a bed of granulated catalyst. A reformer (*below*) has pressure vessels filled with precious-metal catalysts. Both units are at Pascagoula.



takes place almost instantly. Then the catalyst and the hydrocarbons have to be separated so the catalyst can be reused. The separation is accomplished with a cyclone, a device that spins the mixture so the denser catalyst migrates to the periphery while the hydrocarbons remain near the center.

The fluid cracking unit has two more major parts. The fractionating column is similar to the atmospheric tower of the crude unit; it separates the products of the cracking reaction. The regenerator, which is usually the biggest part of the entire cracking unit, receives the spent catalyst and gets it ready for reuse. With a blast of hot air, it burns away carbon that plugs up the pores of the zeolite.

Where does the carbon come from? Answering this question calls for a detour into the chemistry of cracking. The feedstock molecules are chains of carbon atoms surrounded by hydrogens. You might think of such a molecule as a long banquet table. The carbon backbone is the table itself; the attached hydrogen atoms are the chairs lined up along both sides of the table and also (this is important) at the head and the foot. When the table is broken somewhere in the middle, two more chairs are needed to fill the newly created positions at the head and foot of the two fragments. The chairs cannot be created out of nothing—and neither can hydrogen atoms. Thus, the cracking of molecules leads to a hydrogen deficiency. The needed hydrogens are stripped away from other feedstock molecules, leaving behind a residue of carbon.

Hydrocracking. What if you ran a catalytic cracking operation but added hydrogen to make up for the deficiency? Roughly speaking, that's what happens in hydrocracking. It is another way of converting heavy oils into gasoline components. As it happens, hydrocracking is done in the kind of fixed-bed catalytic reactor that would be so old-fashioned in cat cracking. This arrangement works because no carbon clogs the catalyst bed, and there is no need for continual regeneration.

The main hardware components of a hydrocracker are hefty pressure vessels, usually of the bullet-tank variety. They are mounted vertically so the feedstock can trickle through the catalyst bed. Also part of the hydrocracking unit are feed heaters, separators to recover unreacted hydrogen gas, and—as everywhere—a fractionating column.

The Reformer. The reformer is not a crusading politician, but it *is* a do-gooder of sorts. It rearranges various hydrocarbon molecules, making them better citizens in the world of motor fuels. The basic idea is to take abundant straight-chain molecules and either give them branches or twist them into closed rings. These geometrically intricate molecules burn more smoothly in automobile engines.

The reformer is another unit based on high-pressure reaction vessels. They are filled with an unusual catalyst, rich in rare metals such as platinum and rhenium. The catalyst is worth a few million dollars, and it has to be replaced every few years.

Alkylation. The cracking units break big molecules into small ones; the “alky” plant does the opposite, gluing together small molecules to make bigger ones. The catalyst

for this process is different from the various powders and granules used elsewhere in the refinery: it is liquid sulfuric acid. Another difference is that the alkylation unit doesn't have a feed heater. Instead, it has a big chiller, where the feed is cooled to 40 degrees Fahrenheit. This is the temperature where the acid does its best work.


The Bottom of the Barrel. The heaviest part of the crude oil has always been the unloved, good-for-nothing, poor relation of the petroleum industry. The stuff burns fine, but it is gooey, stinky, dirty, and almost worthless. Since the 1980s many refineries have built new facilities for dealing with the bottom of the barrel; they can't make the noxious elements go away, but they can concentrate them. The process is called coking, and it cooks the oil until there's nothing left but carbon. A coking unit is a bizarre—and unmistakable—sight. It has drilling derricks, as if the refinery has had the good fortune to strike oil on its own property! But a closer look reveals something odd about these drilling rigs. They are mounted a hundred feet in the air, atop tall steel vessels. The vessels are the coking drums, where the overheated residual oil breaks down, releasing all its volatile constituents but leaving behind a hard mass of solid carbon. The drilling rigs are needed to break the carbon loose. A coking drum is cooked for about



An alkylation plant at Pascagoula (*above*) also relies on bullet-tank pressure vessels, but in this case they are laid down horizontally; one large tank is visible here behind the smaller drums of a heat exchanger. The unit also includes several tall distillation towers.

Coking units at Pascagoula are an unmistakable feature of the skyline. The coke is formed in the tall drums in the base of the tower. Oil heated to more than 1,000 degrees Fahrenheit breaks down there, releasing all volatile components and leaving a residue of solid carbon. The drilling derricks are needed to remove this carbon after a drum has filled up. The drill creates a pilot hole through the carbon deposit, and then a high-pressure jet of water washes out the rest.

Flares flicker against the evening sky in Pascagoula. Most of the time, the flame atop each stack is just a pilot light. In emergencies, however, the stacks are required to dispose of very large quantities of flammable gases. One reason the stacks are so high is that the heat of an emergency flare could injure anyone closer to the flame than about 100 feet.



24 hours, until it fills up with carbon. Then the drum is allowed to cool, and the rig drills a pilot hole through the coke. A jet of water flushes out the rest of the carbon.

Petroleum coke is soft, spongy, and intensely black—like concentrated grime. It's too full of sulfur to be burned anywhere subject to air-quality standards. What becomes of it? Most of it is exported to less fastidious countries.

Product Blending. Refiners convert as much crude as possible into transportation fuels—diesel, jet fuel, and especially gasoline. Most of the molecules that go into gasoline have five to eight carbon atoms, but when I learned how gasoline is blended, one ingredient surprised me. It is butane, lighter than the rest, with just four carbons. I was surprised because butane boils at about 30 degrees Fahrenheit, and so I thought it would vaporize and escape.

Butane is added to gasoline for two reasons. First, it has an excellent octane number (the main rating of gasoline quality), and it's much cheaper than other high-octane additives. The second reason is the low boiling point, which turns out to be an advantage as well as a drawback. Vaporizing is just what gasoline needs to do in order to burn in an automobile engine. Having some highly volatile components is particularly important for starting the engine.

For economic reasons, refiners would put as much butane as possible in a gasoline blend, but they can't overdo it or much of the butane would evaporate. The evaporation rate varies with the weather. Hence, motor fuel is blended differently according to the season; gas for a cold climate has more butane. When you fill your tank, you may notice wavy fumes emerging from the filler pipe. That's butane escaping.

The flare stack. People who give refinery tours tell me the flare—that pulsating flame held aloft on a tall mast—is what everyone asks about first. At one time, refineries burned off large volumes of gas that was not worth selling, but that's not done anymore. In many refineries today the flare does not burn at all during routine operations. It is there for emergencies. If the gas plant has to be shut down suddenly because of a leak or malfunction, all the gases coming out of the other refinery units have to go somewhere. Burning is the safest and cleanest way to get rid of them.

The burner at the top of the flare is a lot like the one in a gas grill or broiler: it has many rows of small nozzles that spread the flame over a wide area and let air reach all parts of it. Bars called flameholders, mounted just above the nozzles, create turbulence to mix the gas and air. Some flares have additional nozzles to inject steam, which, paradoxically, makes the flame burn cleaner.

Sulfur recovery. Much of the sulfur in crude oil is released in the refinery as hydrogen sulfide, the rotten-egg gas. Until about 1970 the common way of getting rid of the hydrogen sulfide was to add it to the refinery fuel gas. There were two problems with this practice. First, it made the refinery a very stinky place. Second, it added a heavy load of sulfur oxides to the atmosphere. The burning of hydrogen sulfide has been stopped now; instead, a recovery plant converts it to elemental sulfur or sulfuric acid. People still hold their nose when driving by a refinery, but it's just a reflex. The ones I have visited lately don't smell at all.

THE GAS STATION

If the refinery is an exotic locale that most people never visit, the gas station is all too familiar. Still, even at the retail end of the oil industry, a few technological elements often go unnoticed.

The gas station has undergone a curious evolution. Early in the twentieth century, gasoline was sold as a sideline by general stores, with a pump at the curb. Then came the classic filling station, which grew into an emporium for all things automotive. The gas station of the 1950s, clad in gleaming white porcelain enamel, as if it were some giant kitchen appliance, would not only fill your tank but also change your oil, fix your flat tire, and replace your leaky muffler. Since the 1980s another

The metropolis for hydrocarbons also has a skyline like that of a large city. This is the ExxonMobil refinery in Chalmette, Louisiana, observed from a ferry crossing the Mississippi River.



The American gas station began as an adjunct to the general store and has returned to that role. A modern station (*right*) in Altoona, Pennsylvania, sells food, beer and cigarettes as much as fuel. An antique gas pump (*below*) is on display as a novelty item at a shop in Greensboro, North Carolina, that was once a gas station but has long since been converted to other uses. The glass vessel atop the pump allowed the buyer to confirm the quantity of gasoline before pouring it into the fuel tank. (The price on the placard is 17 cents per gallon.)



transformation has come over the gas station: it has fused with the convenience store. Now you can buy gas and pick up a six-pack of beer and a bag of chips all in one place, but you need to go elsewhere—to an array of more specialized franchises—for an oil change, a tune-up, a muffler, or tires.

In other ways too the gas station has come full cycle. The early curbside pump was often a self-service device. Later stations were staffed by squadrons of uniformed attendants who not only operated the pump but also swabbed the windshield, checked the oil, and issued Green Stamps. Now (except in a few states) we have returned to the tradition of do-it-yourself. Pumps equipped with a credit-card reader allow the entire gas-buying experience to be completed without having to speak with a human being.

Meanwhile, the gas pump itself—the heart of the gasoline retailing business—has also been evolving. Early pumps were cranked by hand, and the gasoline flowed first into an elevated glass vessel graduated in gallons (or liters). That way the purchaser could verify the quantity of fuel before it was allowed to flow by gravity into the car's tank. This principle continued to be honored in vestigial form for decades thereafter; as late as the 1950s some pumps still had a glass globe where you could watch the gasoline flowing through and turning a plastic turbine.

Early retailers stored gasoline in drums or barrels. Then, as larger volumes were sold, the storage tanks were moved underground, both to save space and for safety. But burying the tank didn't necessarily bury the problem. Steel tanks would eventually corrode and begin leaking. For a long time, minor leaks were not considered a serious issue. The only cost of a leak was the lost gasoline. Now, an undetected underground leak is the gas station owner's worst nightmare. Hundreds of tons of contaminated soil have to be dug out and carted off to be treated as hazardous waste.

To deal with the leak hazard, steel tanks have been replaced by corrosion-free fiberglass ones. Monitoring wells are drilled all around the tanks for early detection of any hydrocarbons entering the soil or groundwater. Sometimes electronic leak-

detection systems are installed. And in some parts of the United States, above-ground storage tanks have come back into favor, especially for kerosene and diesel fuel, which don't present quite as much risk of fire and explosion as gasoline does.

A typical tank at a large gas station holds 12,000 gallons, and there are three or four such tanks to accommodate the various grades of fuel. The tanks are filled through fittings recessed into the concrete apron that surrounds the pumps. A standard color code is developing for the various hatches and access caps. White marks regular unleaded gasoline, blue is the "plus" grade, and red is premium gas. Yellow and green are for diesel fuel—the yellow for low sulfur and the green for high sulfur. Finally, brown is for kerosene. Monitoring wells are danger orange: you wouldn't want the delivery truck to pump 10,000 gallons of high test into the well.

In heavy-smog areas, gas stations have had to make another change in the past few years: installation of vapor-recovery systems. As noted earlier, gasoline is rich in butane, a very restless molecule that takes any opportunity to escape confinement. The gas-pump nozzle is fitted with a plastic hood, and an interlock switch allows the pump to run only when the hood is seated to the filler pipe. Suction draws the vapor through the hood and into a separate hose, eventually to be returned to the storage tank.

NATURAL GAS

The petroleum and natural gas industries are twins separated at birth. They have a great deal in common; indeed, the raw materials for both industries often come out of the same hole in the ground. And yet they have grown apart, developing different cultures, practices, and technologies.

The word *gas* is a variation on *chaos*. The name was suggested by Jan Baptista van Helmont, a seventeenth-century Flemish chemist, who thought of gases as unruly spirits given off when solids or liquids are heated. (Or maybe he was just being whimsical.) The stuff we burn today is called natural gas to distinguish it from manufactured gas, which came first historically. All through the nineteenth century and up until about 1950, gas was made by roasting coal or heavy oil residue to drive off a mixture of combustible gases—mostly carbon monoxide but also some hydrogen and methane and traces of much else. The gas was a poor fuel, not to mention a deadly poison. The residue left behind at the gas works, called coal tar, was even nastier. Efforts to find a use for the stuff or somehow get rid of it were so assiduous and enduring that they pretty much created the discipline of organic chemistry, and then the dyestuff and pharmaceutical industries.

Manufactured gas hasn't been manufactured for decades, but many of the companies that once made it live on as distributors of natural gas. Companies such as Brooklyn Union Gas and Public Service of New Jersey converted their old gas works to storage facilities for the new fuel flowing in from Texas and Oklahoma. The most difficult part of the conversion was readjusting millions of stove burners and furnaces.

The storage of gasoline and other fuels at retail outlets has also undergone significant change. Originally stored in elevated tanks, gasoline was later put underground for safety reasons. But in many jurisdictions, tanks are now being installed above ground again because of the risk of underground leaks. The tanks below are at a gas station in western Missouri.



Natural gas is a simpler fuel than manufactured gas. It is mostly methane, the smallest of the hydrocarbon molecules, with few impurities except water. The only major products of combustion are carbon dioxide and water. One impurity is deliberately added to the gas before it enters the distribution pipelines: methyl mercaptan, the mercury compound that gives gas its characteristic smell. (The methane itself is odorless.)

Natural gas became a practical fuel only with the construction of long-distance pipelines. Other forms of transport are simply too expensive, which means that gas fields not served by a pipeline can't be developed. The gas in the Alaskan North Slope, for example, has no way of getting to market, and so it is reinjected into the wells, where it helps push oil to the surface. A few oceangoing tankers carry natural gas as a cryogenic liquid (temperature -259 degrees Fahrenheit), but commerce in liquefied natural gas (LNG) has not taken off the way the world petroleum market has.

Gas received through the pipeline system is stored in the equivalent of tank farms. The traditional storage facility is a gigantic, collapsible canister, or holder, that rises



Rigid-walled gasholders loom over a residential neighborhood in Genoa, Italy. Most tanks of this kind were built in the years before 1950 and originally held manufactured gas. Often they are the tallest structures in the area and therefore serve as important landmarks. Inside the tanks, a floating piston separates a layer of natural gas below from air above.



Telescoping gasholders in Elizabeth, New Jersey, rise and fall within their skeletal framework as gas is stored and withdrawn. In this photograph the gasholders are fully deflated. They had probably been emptied in preparation for demolition; since the photograph was made in 1998, the gasholders have disappeared, like many others of their kind.

and falls with changes in the supply and demand for gas. In New York City, generations of commuters came to know two such gas tanks near the Long Island Expressway in Elmhurst, Queens. The Elmhurst tanks have been demolished, like many others around the country.

The holders are essentially inflatable structures, but made of telescoping sections instead of elastic fabric. The walls are concentric rings called cups, linked together by flanges. As the empty tank starts to fill with gas, the innermost cup (which supports the slightly domed roof) starts to rise. A flange at the base of the innermost cup engages a matching flange at the top of the next cup, which is therefore pulled up as the filling continues. The joint between sections is sealed with water, which is steam-heated in winter to keep it from freezing. At full extension a typical holder stands more than 200 feet high and holds roughly 10 million cubic feet of gas. The sections are held in alignment and braced against wind loads by a cylindrical exoskeleton of steel trusses.

Because of the immense volume, the pressure needed to inflate the holder and lift the cups is very slight—only about one-half pound per square inch above atmospheric pressure. You could blow up the holder with your breath, like inflating a giant beach toy, but it would take a while—perhaps 300 million lungs full.

There are also rigid gas holders, which look from the outside like plain steel tanks without moving parts, but inside they have a floating piston that separates gas in the lower part of the holder from air above. (Keeping air and gas apart is the point of all the storage arrangements. Gas is reasonably safe as long as there's no oxygen present.)

The big low-pressure gas holders are being replaced by high-pressure tanks and by insulated tanks for liquefied gas. The typical facility for high-pressure storage is a bank

High-pressure cylinders store natural gas in a smaller volume than the low-pressure rigid and telescoping gasholders. The cylinders are in Ljubljana, Slovenia.



Liquefied natural gas (LNG) is stored at even greater density than the high-pressure gas. The LNG depot at right, shrouded in foggy vapors on a winter afternoon, is in Everett, Massachusetts.

A vent stack on a city sidewalk serves as a pressure-relief valve for the gas distribution system. Ordinarily the valve is closed and nothing escapes through the vent; if a malfunction causes high-than-normal pressure in the mains, it is better to release the gas through a vent like this one than through leaks in underground pipes or inside buildings.



of cylindrical tanks with hemispherical end caps. Each tank is about the size of a house trailer. A dozen of these tanks have as much capacity as one of the giant old telescoping holders. That's about a day's supply for a major city.

Larger stocks are stored in liquid form. Six hundred cubic feet of natural gas condenses into a single cubic foot of LNG. Tanks for LNG don't have to withstand pressure, but they have to be insulated to keep the fuel cold and be built of materials that retain their strength in the deep freeze. In addition to the storage tanks, an LNG depot will have a refrigeration plant, and an evaporator, which accomplishes the opposite task—making the gas a gas again.

Although storage technology has changed, the gas-distribution system remains much as it has been for the past century. The gas moves through underground pipes at very low pressures—roughly 0.2 pound per square inch over atmospheric pressure. This is less than the change in barometric pressure when the weather turns from fair to stormy. Thus, the gas is very gently wafted through the pipes. The low pressure entails much larger pipes than a high-pressure system would need. Some of the gas mains in New York are six feet in diameter. Low pressure is a safety feature; it reduces the leakage rate when a pipe fails or a pilot light goes out.

The traditional material for gas mains is iron pipe, but distribution lines laid in the past few years are mostly plastic, colored bright yellow for easy identification.

FUTURE FUELS

Petroleum is routinely called the life's blood of the industrial economy. The metaphor is not far-fetched. Societies rely on oil and gas to meet a share of almost all energy demands, and transportation in particular would go nowhere without petroleum. The

least hiccup in supply sets off international alarms. Nations do not hesitate to go to war over oil.

But given the central and essential place of petroleum, it's startling to reflect that oil is a newcomer to the world economy. The industry sprouted up only a little more than a century ago. What's more, in another century or so it will have disappeared. The petroleum age is a brief episode in human history.

No one knows how much oil remains in the ground, or how much of it can be recovered at acceptable cost. There is surely more to be found and more to be pumped than the oil companies know about—or tell about. Nevertheless, it is beyond dispute that the resource is a finite one. We are burning petroleum much, much faster than the earth is making it. And there is reason to suspect that we are near the high-tide mark in oil production and consumption. In the United States, petroleum production peaked in 1970. Production elsewhere has continued increasing since then, but Kenneth S. Deffeyes, a geologist with close ties to the oil industry, has argued that the upward trend cannot last more than a few years; indeed, he has gone on record predicting that the peak will come on Thanksgiving Day in 2005. After that, he says, it's all downhill—although the decline might take as long as the build-up did.

The prospect of running out of oil a few decades from now should not be cause for panic or despair. Given that the whole infrastructure of the petroleum industry was built in less than a hundred years, there should be plenty of time to create its replacement. It will be interesting to see what that replacement is, and what new features it adds to the technological landscape. And if a world without gasoline seems unimaginable, look back to the 1850s, when a world without whale oil and a whaling industry must have seemed equally unlikely and forbidding.

THE HYDROGEN ECONOMY?

In one popular vision of the post-petroleum future, automobiles powered by fuel cells will convert hydrogen and oxygen into electricity, producing only water as exhaust. President George W. Bush endorsed this idea in his 2003 State of the Union address, and proposed a \$1.2 billion program of research.

The two main elements of this plan—the use of fuel cells and the use of hydrogen—are almost independent. If fuel cells turn out to be the best way of powering automobiles, they can probably be made to work on fuels other than hydrogen. At the same time, if hydrogen is readily available as a fuel, it could be put to

work in many ways, including burning it in engines much like those of cars today.

The big question is where the hydrogen comes from. Some envision splitting water molecules either electrically or by applying intense heat. Thus, the hydrogen economy would be a closed cycle: water would be broken down into its constituent hydrogen and oxygen molecules, which would later recombine to make water again. It's an attractive scheme, but an energy source is needed to make the cycle run. The electricity or the heat for prying apart water molecules would have to come from a power plant of some kind. In other words, hydrogen is

not an energy source in this plan; it's merely an intermediate or carrier, a way of turning coal or nuclear power into automobile fuel.

There is another source of hydrogen. Almost all hydrogen used in industry today is made from methane, by stripping the four hydrogen atoms away from the single carbon atom in a methane molecule. There's a well-developed technology for doing this, and it could be scaled up to supply larger quantities of hydrogen as fuel. Nevertheless, this process will not solve the problem of what to do after the oil and the gas run out. The methane from which hydrogen is made comes from natural gas.