AVIATION

As the pilot maneuvers into position for final approach, the aircraft banks to the right and you get your first clear view of the airport below. Lights leading up to the runway threshold flash in sequence like a rapid-fire movie marquee. Crossbars of bright red and white lights draw the eye toward the touchdown zone. The runway itself is outlined in multicolored lamps, with another row of lights down the center. As the airplane descends, you notice the rotating beacon atop the control tower, flashing alternately white and green. Once the aircraft is on the ground there are still more lights—a forest of dim blue ones mapping out the taxiways. What are all these lights and signals? And what about the stripes, bars, chevrons, numbers, and other markings painted on the pavement? What do they all mean?

Aviation breaks free of the earth—that's what it's all about—and yet it has left large marks on the landscape. Major airports are highly conspicuous; some of them are as big as cities, with populations to match. Away from the airport are other telltale signs of the airplane's influence, such as the flashing beacon lights and the red-and-white stripes or checkerboard patterns on water towers and the tall masts of radio and television transmitters. Another important part of the aviation infrastructure is almost invisible: the network of air routes that most aircraft follow from city to city. Even though the air routes themselves are out of reach far overhead, you may come upon some of the navigational beacons that act as signposts along the way.

THE AIRPORT

Airports are not among our best-beloved public places. Railroad terminals get preserved as historic landmarks even as the passenger railroads themselves wither away, but airports are at best tolerated as a necessary evil to be hurried through on the way

Like the minaret of a mosque, the control tower soars above the vaulted roofline of Reagan National Airport, just outside Washington, D.C. The tower and terminal were designed by architect Cesar Pelli. Controllers in the windowed "cab" at the top of the tower direct aircraft on the ground and while taking off or about to land; other flight operations are handled from a regional center. The "hat" atop the tower covers a radar antenna used to monitor aircraft movements on the ground. Since this photograph was made, that radar has been moved to another part of the airport.

New terminals at Charles de Gaulle airport in Paris reflect the "inside out" structure favored in many recent airport designs. Passengers arriving by road or by rail are deposited in the middle of the complex and then fan out toward terminals and gates, which in turn are surrounded by taxiways and runways. Seen here in a photograph made in September 2003 are terminal 2F (in the foreground) and 2E (in the background). On a Sunday morning in May 2004 a section of the roof of terminal 2E collapsed, killing four travelers.

from here to there. Passengers complain of the inhuman scale, the sterile architecture, the unwelcoming environment. It will be interesting to see whether the disdain for airports persists in generations to come. Who knows—maybe 50 years from now there will be a popular movement and fund-raising campaign to save O'Hare International from the wrecking ball.

Airport terminals have undergone a surprisingly complicated evolution since commercial air travel began in the 1930s. The prototypical passenger terminal was just a building that served as the buffer between ground transport and air transport. You drove up to the ground side of the terminal building, bought your ticket and checked your bags, then you walked through to the air side, going out the back door, across the tarmac, and up the stairs into a waiting airliner. At the end of the flight you passed through a similar building in the opposite direction.

As the volume of passenger traffic increased, terminals had to grow more elaborate. On the land side, the building is often split into two levels, separating arriving and departing passengers. The upper floor usually has the ticket counters; the lower floor, the baggage-claim carousels. On the air side, the main problem is finding space to park more and more aircraft, which grow larger and larger. As a rule of thumb, a parking space for a wide-body four-engine jet takes up 150,000 square feet—which would make a fair-sized lot for a suburban house. In response to the demand for more space on the air side, many terminals grew various protuberances and excrescences—long finger piers, Y-shaped or T-shaped branching structures, bulbous loops on the ends of piers, and isolated satellite terminals reached by buses or underground passages. And passengers learned that to fly from New York to Chicago you have to walk as far as Cleveland.

The next major transformation of the airport was prompted by the spate of airplane hijackings and bombings in the 1960s and 1970s; it intensified after September 11, 2001. Security screening created a new set of zones in the passenger terminal. Now you pass not only from the land side to the air side, but also from the nonsecured to the secured zone. The secured zone is a lot like the sterile field of a surgical operating room: everything "outside" is officially classified as unsafe until it has passed through the metal detector or the x-ray machine at the security checkpoint; inside, everything is assumed to be safe. Note that the secured zone includes the interior of the aircraft, which means that the zone extends across whole continents. When you are cleared to board an airplane in Miami, you are still considered safe, and therefore you are admitted directly to the secured zone when you get off in Seattle. But if you stray a few steps back to the land side of the terminal, you are contaminated and must be inspected again for readmission.

Standard airport security practices actually create three zones. The outside world, where nothing can be trusted and where anyone may go, extends into the groundside areas of the terminal as far as the security checkpoint. The secured zone, where all persons and carry-on bags have been examined, includes the air side of the terminal and the interior spaces of the gates and jetways and the aircraft. The third zone is the part of the airport outside the passenger lounges—the aprons, taxiways, run-

GETTING A LOOK

Once upon a time, my father would take me to the airport on Saturday afternoon. We weren't flying anywhere; we just parked the car at the end of the runway and watched the planes roaring overhead as they landed and took off.

Times have changed. Several years ago our spot near the runway threshold was posted with No Parking signs, and eventually the entire road around the airport perimeter was closed to traffic. (Another change is that ten-year-olds these days probably don't consider a trip to the airport as proper Saturday-afternoon entertainment.)

More than any other realm of daily life, aviation has been transformed by worries over security and the threat of terrorism. The airport is definitely not the place for casual trespassing. Just opening the wrong door could land you in jail.

On the other hand, the public areas of airports remain very public indeed. Thousands of travelers pass through them every day. Much of the aviation infrastructure is readily seen from the departure lounge, or from a window seat on any flight, or even from the upper levels of the parking deck. A few airports still have observation areas open to the public.

Navigational aids installed away from the airport are always fenced off and bear the kind of stern warning sign reproduced at left. Unlike most such signs, it not only tells you to keep out but also explains why it's important that you do so. The reasons are certainly good ones, but if you stay outside the fence, you'll do no harm—and come to no harm. (Note, however, that some countries consider air-traffic-control hardware to be militarily sensitive, and they frown on picture taking.)

Atlanta's Hartsfield Airport is the prototypical design for the age of hub-and-spoke airline operations. Ground-side facilities are at the far left, surrounded by parking lots and decks. Most of the boarding gates are in the five concourses in the middle of the field, which passengers can reach only via an underground rail link. Eliminating automotive connections to these concourses allows aircraft free access to all sides of them, maximizing the number of gates in a given area. The layout is ideal for passengers making connections between flights; it is less than optimal for those whose trips start or end in Atlanta. The airport has four parallel runways, which can all be active at the same time. In this aerial photograph, made as part of the U.S. Geological Survey Urban Areas series, the runways and taxiways form an elaborately perforated gasket-like pattern.

ways, and other areas with access to the exterior of airplanes, as well as to luggage holds. These places are off limits to everyone but employees with the right kind of identifying badge. (International terminals with customs and immigration inspections have an even more elaborate division into zones.)

Airport architecture has also been altered by the shift to hub-and-spoke operations by most scheduled airlines. In decades past, an airline that served 10 cities might have flown most of the 90 possible routes connecting pairs of those cities. Designating one of the cities a hub and flying only between the hub and each of the spoke cities reduces the number of routes to 9. This strategy simplifies operations for the airline, but it also means that most trips require two flights, with a change of planes at the hub. Accordingly, new demands are made on the hub airport. At a spoke airport, most passengers are either beginning or ending their journey, and the main imperative is to minimize the distance from the gate to the parking lot or the taxi stand. At hubs, many passengers are merely changing planes, and they never use the land-side facilities. For them the highest priority is to minimize the distance between gates.

Designers have responded to this new traffic pattern with a wholly new style of airport, one that would have seemed quite improbable 20 years ago. The new hub airports divorce the land side from the air side, isolating the two functions in separate buildings as much as a mile apart. This strategy allows cars free access to the entire perimeter of the land-side building, while aircraft can taxi up to gates all around the air-side concourses. Moreover, putting the air-side building out in the middle of the landing field reduces taxi distances to and from the runway, allowing quicker turnarounds for connecting passengers. The price is paid by those passengers whose trip begins or ends at the hub airport: they have a long trek between the land side and the air side. In the United States the two main examples of this design are Hartsfield-Jackson International Airport in Atlanta and the new Denver International Airport. In both cases an automated underground rail line carries passengers from the land-side building to the isolated gate concourses out in the middle of the field.

The largest metropolitan airports are hubs of economic activity as well as transportation. They employ tens of thousands, and at a busy hour their population can be that of a medium-size city. Some fraction of the purchase price of every airline ticket goes to support the airport's operation through landing fees and rents charged to the airlines. But this revenue stream is generously supplemented—and sometimes exceeded—by income from snack bars, newsstands, souvenir shops, and most of all from automobile parking. From an economic point of view, an airport is a device for persuading people to pay $50 to leave their car over the weekend or $5 for a warmed-over slice of pizza. It's no accident that some of the companies that build shopping malls are expanding into airport management, and vice versa.

THE TERMINAL APRON

Standing at the gate, you look out the window at the aircraft you're about to board. It stands on the terminal apron, the area where aircraft park while loading, unloading, and preparing for their next flight.

Parking a vehicle that's 200 feet long and almost as wide presents a considerable challenge. On the pavement you may see painted marks indicating the correct position for the nosewheel. Often there are several marks, labeled for different aircraft—Boeing 727, 737, Airbus 300, and so on. However, the pilot steering the airplane cannot see any of these marks when they are needed most because the nosewheel is under and behind the cockpit. Usually, the pilot follows the signals of a worker with flashlights or bright orange paddles. Some airports have a system of marker signs mounted on the wall of the terminal building that guide the pilot directly. Each sign is split into two pieces, with one piece mounted a few feet farther from the wall than the other piece. The two parts align correctly only when seen from the cockpit of a properly parked aircraft. Another system uses lights that change from red to green as the aircraft pulls into position.

At most airports the first item connected to an arriving aircraft is the telescoping jetway—the passage you walk through when you enter or leave the airplane. This strange articulated appendage is part building, part vehicle, and part bridge. It can nod up and down, extend and retract, sweep from side to side, and in many cases pivot at elbow-like and wrist-like joints. All of these motions are controlled with a joystick from a station next to where passengers board the plane.

Other connections to the parked aircraft are for utilities (fuel, water, electric power, air conditioning) and deliveries (luggage, meals).

Fuel. At most large airports today, fuel is delivered by a hydrant system. The driver of a small truck opens up an access panel in the pavement and connects a hose from the truck to a fitting in an underground vault; a second hose goes from the truck to a socket on the underside of the airplane wing. Actually, plugging in the hoses is not the first thing the driver does; before making any fuel connections, the driver clips on grounding wires that run from the airplane to the truck and to the underground vault to prevent sparks from static electricity. Once all the connections are double-checked, a valve is opened and fuel flows under pressure into the wing tanks.

The truck doesn't have to pump the fuel; the pressure is supplied by large, stationery pumps at the main fuel depot, which is generally on the periphery of the airport grounds. The machinery on the truck is mainly for filtering out air, water, or other contaminants in the fuel. The truck also measures the quantity of fuel delivered so that the airline can be billed. A few airports have filters and meters in each underground vault. This eliminates the need for the hydrant truck; the refueler just pulls a hose out of the vault and plugs it in.

Some airports have no underground fuel-distribution system at all. The fuel is delivered by tanker trucks—sometimes called bowsers—much like the trucks that bring gasoline to the corner gas station or fuel oil to homes. But the airport fuel trucks are larger, carrying as much as 8,000 gallons each. That makes them too big for the public roads, and yet they still can't carry nearly enough to quench the thirst of the biggest jetliners. A Boeing 747 holds roughly 50,000 gallons of fuel; in the long-range 747-400 model, the weight of the fuel is roughly equal to that of the airplane itself—more than 350,000 pounds. The refueling rate can reach 2,000 gallons per minute. If you had to fill up the tanks of a 747 from a gas-station pump, it would take the better part of a week.

Commercial jet aircraft burn a fuel called Jet-A, which is essentially kerosene and is similar to home-heating oil and diesel fuel. Piston-engined propeller planes burn high-octane gasoline (called avgas).

Electric Power. The heavy black cable that looks like a big extension cord is a big extension cord. It's usually draped along the underside of the jetway and plugs into a socket near the nose or on the underbelly of the airplane. If power isn't available from the terminal, the airline may roll up a generator cart that plugs into the same

The apron—the area just outside the terminal building where aircraft are serviced between flights—is the busiest part of an airport. At left, three aircraft nuzzle up to telescoping jetways at Miami International. On the opposite page, aircraft are guided and towed to parking places on the apron, and are connected to vehicles supplying air conditioning, fuel, and food. Below, deicing operations are seen from the point of view of one being deiced.

outlet. All larger aircraft have an on-board source of electricity—it's a small jet engine running an electric generator—but getting juice from the ground saves fuel and cuts down on noise.

Air, Water, Food, Baggage. A fat floppy hose about a foot in diameter, looking like a grown-up version of the hose that connects to the back of your clothes dryer, carries chilled air to the plane and thereby takes the load off on-board air-conditioning equipment. As with the electricity, the chilled air could come from a source in the terminal or from a mobile cart.

Water comes in a tanker truck, often prominently marked "Potable Water." Perhaps that's to distinguish it from the *other* tanker truck, which pumps out the sewage.

The catering truck is a van whose body is mounted on a hydraulic lift, like an oversize version of the jack you would use for fixing a flat tire. The truck rises so that carts of meal trays can be rolled directly onto and off of the passenger deck.

For wide-body aircraft, baggage trucks use a similar lift mechanism. The bags are preloaded into aluminum or fiberglass containers, which are then hoisted up into the belly of the airplane. On smaller aircraft, bags are loaded individually, using an inclined conveyor belt. The bags are hauled around in trains of rubber-tired carts.

What's most notable about all these trucks and other vehicles that buzz around tending the airplane is their specialization: they would have no use off the airport grounds.

Tugs and Tractors. Although jet airplanes have thrust reversers, they are used only to slow the airplane right after touchdown; they can't be used as a reverse gear for backing out of a parking place. Thus, airplanes usually have to be pushed away from the gate. The machine that does the pushing is a low-slung tug or tractor with a pole

that links onto the nosewheel strut. Keeping enough tugs operating is a critical requirement of airport management, as a shortage can delay flight after flight.

Deicing. Accumulations of snow and ice can crush buildings and topple trees, so it's no surprise that they can also keep aircraft on the ground. (The problem isn't just the weight; a layer of ice can also change the shape and hence the aerodynamics of a wing.) At many airports deicing is done at the gate by workers in bucket-lift trucks. They spray warm antifreeze to wash the ice off. The trouble with this system is that the wings can gather a new coat of ice by the time an airplane taxies out to the runway; then it has to come back to the gate to be deiced again, a process that obviously could go on indefinitely. Some airports where icing is a frequent problem have a permanent deicing pad, something like an automated car wash, out near the departure end of the runway. Doing all the deicing in one spot also makes it easier to recover and recycle the antifreeze.

THE AIRFIELD

Beyond the gate and the terminal apron is the realm of runways and taxiways. Out in these wide-open spaces, aircraft enter their own element.

Runways. The runway is where the rubber *really* meets the road; it is the exact point of interface between air and ground transport. At first glance a runway looks just like a road—but it's a road writ large. A fully loaded Boeing 747 can weigh as much as 750,000 pounds (more than 10 times the weight of a big tractor-trailer rig), and it lumbers down the runway at 130 miles per hour or more. A surface built to highway specifications would not hold up for long under such treatment. Accordingly, the concrete slab of a runway may be as much as two feet thick, laid atop a carefully prepared foundation. Sometimes a thin layer of asphalt is laid down over the concrete to improve traction. (People are always talking about "the tarmac" in connection with airport paving, but actually there is very little of the tar macadam that properly goes by that name.)

A runway is not only thicker than a highway but also wider—up to 200 feet, or the equivalent of about 15 lanes of automobile traffic. All in all, a heavy-duty runway two miles long consumes enough concrete to build 60 miles of two-lane roadway.

The length of a runway depends mainly on the kind of aircraft an airport serves; obviously, big jets need more room than small private planes. The critical factor is take-off rather than landing, because an airplane needs more room to gather speed than it does to stop. Other factors also enter the equation: high elevation and high temperature both call for a longer takeoff roll. The reason is that heat and altitude reduce the density of the air, which means they also reduce the lift generated by an airplane's wings. The difference can be dramatic. A Boeing 727 that needs 4,000 feet of runway at sea level on a cool day might require an 8,000-foot takeoff roll at an elevation of 5,000 feet in 100-degree weather. On an unusually hot day airlines may have to reduce load in order to take off safely, with the annoying result that your luggage may not get out of Denver with you.

Most airports that handle jet aircraft have at least one runway of 6,000 feet or more. For "heavy" jets, such as the Boeing 747 and the Airbus A-320, lengths of 8,000 to 10,000 feet are common. John F. Kennedy International Airport in New York has a runway 14,500 feet long, or nearly three miles. And there is an even longer runway at Doha International Airport in Qatar, where summer temperatures commonly reach 120 degrees Fahrenheit.

How can you determine the length of a runway? From the air, look for sets of white stripes, laid out in pairs on either side of the runway centerline. Typically there are single, double, and triple stripes, as well as broad white bars that mark the touchdown zone (the point that pilots are supposed to aim for when landing). Although there is some variation in the details of these markings, the distance from one set of stripes to the next is always 500 feet, and so you can use them as a ruler to estimate the total length.

Aircraft spend very little time on runways—takeoff and landing are the briefest phases of a flight—and yet those moments of transition between airborne and earthbound status are obviously crucial. On the opposite page, a U.S. Geological Survey aerial photo shows the markings on runways 34L and 34R at Fort Worth Alliance airport in Texas. The "34" designation means the runways are oriented with a magnetic bearing of approximately 340 degrees, or just west of due north. The stripe patterns are at intervals of 500 feet and thus can be used to estimate the length of the runway. At left, a runway and a taxiway are seen in the dawn light at Phoenix Sky Harbor Airport. A consistent marking scheme—white striping on runways, yellow on taxiways—helps avoid potential confusion.

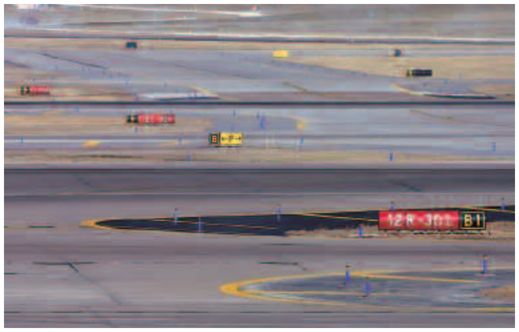

Red-and-white signs mark numbered runways; black-and-yellow signs point the way to lettered taxiways. At the bottom is a hold-short line, marking a point where aircraft must wait for clearance before entering or crossing a runway.

On the ground, during the takeoff roll, what matters most is not the total length of the runway but how much you have left ahead of you. For white-knuckle flyers who worry about such things, major airports have made the countdown easier. There are placards along the edge of the runway with the numbers 8, 7, 6, 5, and so on, marking off the remaining distance in thousands of feet. The placards have large white numbers on a black background. It's just like the Cape Canaveral countdown, but you want to reach liftoff *before* you get to 0.

Another runway marking is an identifying number, painted in numerals 20 feet high. The number represents the approximate magnetic compass direction of the runway, but there's a trick to reading it: the direction is rounded to the nearest 10 degrees and then the final zero is dropped. Thus, Runway 5 points toward 50 degrees, which is roughly northeast; Runway 18 has a bearing of 180 degrees, or south. Where an airport has two parallel runways, they are marked *L* and *R* for left and right. Runways can be used in either direction, so that Runway 9*R*, pointing east, would be marked at the opposite end as Runway 27*L*, heading west. Pilots learn to mentally calculate a "reciprocal bearing" to match up runways going in opposite directions. The arithmetic is simple: if the bearing is 18 or less, add 18; otherwise subtract 18. Because the bearings given are magnetic, they do not correspond exactly to geographic directions, although throughout the lower 48 states the difference between magnetic north and true north is no more than 20 degrees.

The centerline of a runway is marked with a dashed white line, and the edges of the paved area are outlined in solid white. These markings are similar to those of a highway, but they have a different meaning. They do *not* divide the runway into two lanes going in opposite directions; on the contrary, the pilot is expected to drive down the centerline, which would be unwise in an automobile.

Paved runway shoulders and paved areas beyond the runway threshold are painted in bold diagonal stripes or in a herringbone pattern. These are places where the pavement is not strong enough to bear the full weight of an aircraft, and the stripes warn pilots not to venture onto them. Why pave such places if you can't drive an airplane on them? The main reason is to prevent scouring of the soil by the blast from jet engines. There are various other strategies for dealing with the jet-blast problem, such as covering the area with stones or building a rampart of jagged concrete blocks. At the foot of a runway you may see a blast deflector, which looks like a giant immobilized snowplow, with a curved surface that directs the jet blast upward.

Taxiways. The taxiways that lead an aircraft to and from the runway are like frontage roads along a freeway. And as with freeways, some airports have high-speed exits that veer away from a runway at a shallow angle, allowing aircraft to clear the runway before they have slowed down to the usual taxiing speed of 20 or 30 miles an hour.

Taxiways are narrower than runways but are otherwise similar in construction and appearance. Indeed, one of the hazards of airport operation is that a confused pilot will try to take off from or land on a taxiway. To help dispel such confusion, all markings on taxiways are yellow, unlike the white markings on runways. Moreover, the taxiway centerline is not dashed, and the edges are marked by double yellow lines.

Whereas runways are numbered, taxiways are lettered. (Pilots and controllers pronounce the letters "alfa," "bravo," and so on, in the international phonetic alphabet.) Runway signs have white letters on a red background; taxiway signs are black and yellow.

Where a taxiway meets or crosses a runway, there is a "hold line": two solid stripes and two dashed stripes painted across the taxiway 100 feet short of the runway. Aircraft stop at the hold line until the control tower gives the pilot clearance. At various points around the airport you might also notice a hold line and a red-and-white sign with the legend "ILS." These mark positions where an aircraft would block the transmissions of an instrument landing system.

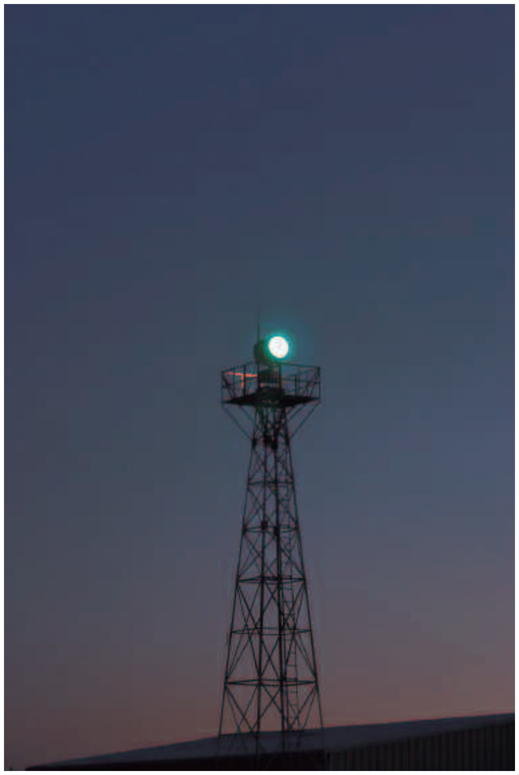

One of the longest-running traditions in aviation is the rotating airport beacon light, green on one side and white on the other. The one above beckons pilots in Hays, Kansas.

Layout of the Airfield. The mapmaker's symbol for an airport is three runways crossing in a triangular configuration. This design was once favored because it gives pilots six choices in selecting a runway aligned with the wind. (Both takeoff and landing are best done heading into the wind.) If the triangle is an equilateral one, there must always be a runway that deviates no more than 30 degrees from the wind direction.

The drawback of the triangular geometry is that only one runway can be active at a time. That limits the capacity of the airport during peak periods, as a single runway can accommodate only about 30 takeoffs or landings per hour even under the best of circumstances. To boost capacity, airport designers have turned to layouts with multiple parallel runways, spaced far enough apart that they can be in use simultaneously. One popular scheme puts the terminal in the middle, with parallel runways on either side.

A feature of parallel runways is that while coming in on final approach, you can sometimes look out the window and see another airplane performing the same maneuvers as your own. It's like having a mirror in the sky.

An airport's main parallel runways are lined up with the local prevailing winds. If space permits, an additional pair of cross runways can be built for those days when the wind shifts out of its usual quarter.

The classic triangular layout allows pilots plenty of options when choosing a runway headed into the wind, but only one of the runways can be active at a time. The landing field seen below is Flagler County Airport in Bunnell, Florida.

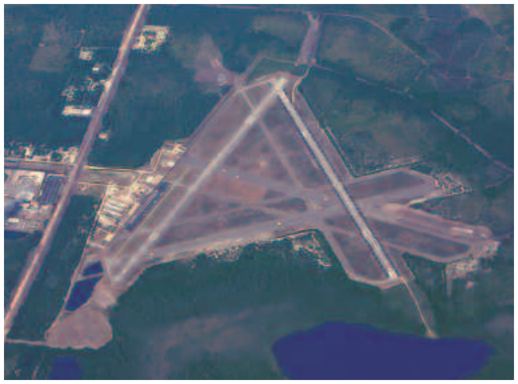

Airport Grading and Drainage. Airports tend to be *very* flat. The earthmoving needed to level the ground can be a major item in the construction budget, but there is not much choice about this investment. For each 1 percent increase in the slope of a runway, the length of the runway has to be increased by 10 percent for jet aircraft and by 20 percent for piston-engine aircraft. Extending the runway would cost even more than flattening the ground. This is one reason so many airports are built on marshes, meadows, and dry lake beds.

With vast expanses of flat, paved land, airports often have serious drainage problems. Runways and taxiways are crowned to help shed water, and often they are grooved to prevent hydroplaning, just as some highways are. Water drains into shallow ditches or collecting ponds and then is carried away either by open channels or by an underground storm-sewer system. The inlets to these sewers are not set at the curbline, as they would be on a city street, but are well back from the runway.

More Airfield Doodads. Here are a few more items you might spot as your plane taxies out to the departure runway.

The fuel depot. For a big airport, a month's supply of fuel is close to 100 million gallons, which means there will be a sizeable tank farm. Finding the best place to put the tanks is a delicate problem. On the one hand, they should be as close to the terminal as possible to minimize pipeline and pumping costs; on the other hand, storing 100 million gallons of flammable hydrocarbons under the path of arriving and departing flights is not ideal.

The weather station. Pilots, like sailors, keep a close eye on the weather, and often the official weather-reporting site for a city is at the airport. The instruments are installed in a grassy area (the National Weather Service requires it) somewhere well away from jet blast. The standard package includes a thermometer, a barometer, and a hygrometer (for measuring humidity), all hidden inside a louvered shelter, as well as a funnel-like rain gauge and a spinning anemometer for measuring wind speed. But for a quick check on the wind, pilots prefer a glance at the wind sock, which indicates both speed and direction. Every airport has one, somewhere out on the field near the runways.

WHIRLYBIRDS

Igor Sikorsky, the helicopter pioneer, flew his prototype machines dressed in a suit and a Homburg. No crash helmet for this test pilot. In old films and photographs he sits in the open cockpit looking like a lawyer or accountant on his way to the office. And, in fact, Sikorsky's dream was that the helicopter would become the everyday conveyance that ordinary folks would fly to work or the grocery store. The family helicopter would replace the family car.

If you live anywhere near a public heliport—or even a landing pad where hospital or TV-station helicopters touch down from time to time—you are probably grateful that Sikorsky's vision never came true. Helicopters have improved a great deal since the early days, but they still make a fearsome noise. If you look around a large city and imagine replacing every car on the streets with a helicopter whomp-whomp-whomping through the skies, this is not a trade that most of us would welcome.

In principle, helicopters can land just about anywhere—a backyard, a supermarket parking lot, a rooftop. But in practice they often land at a facility specially built for just that purpose. The typical heliport is a fenced-off concrete pad with a big letter *H* painted inside a square. The *H* signifies a public heliport; private ones often substitute a company logo. A hospital landing pad may show a white cross with a red H in the middle. Still another kind of marking is seen on the roof of a tall office building or a high-rise hotel: a circle with a number in the middle. The rooftop structure labeled in this way is not meant for routine helicopter operations, but it can be used for emergency evacuations; the number indicates the maximum weight the platform will support, in thousands of pounds.

The FAA suggests that heliports be paved with concrete rather than asphalt. The reason is that the skids or wheels of a helicopter might sink into a softer material, especially in hot weather. Apparently there have been a few accidents where a helicopter tipped over on takeoff because it was stuck in its own ruts.

Although a helicopter doesn't need a runway, it does require an approach and takeoff zone cleared of tall structures. The direction of the recommended final-approach pattern is indicated by the alignment of the *H* or other marker. Wind is also important on takeoff and landing, and the heliport will have a wind sock somewhere near the touchdown zone.

Visibility measurements. Flying blind is routine these days; airline pilots seldom need to look out the window while they're aloft. On the ground, however, it's a different story. Airports sometimes have to close because the fog is too thick for pilots to feel their way to the terminal. For a long time, visibility was "measured" by looking out the window of the tower at some distant landmark. Instruments now give more consistent readings of RVR, or runway visual range. The instrument consists of two pods alongside a runway. A light source in one pod is beamed toward a reflector in the other, which bounces the beam back to a photocell in the first pod. The photocell measures how much of the emitted light comes back, which is an indicator of atmospheric clarity. A related instrument measures the height of the "ceiling," or the bottom of the lowest bank of clouds overhead. Again, the instrument consists of two pods, but they point upward rather than at each other. One pod shines a beam of light at a slight angle from the vertical; the other pod tilts until it detects the reflection of the beam. The angle of tilt indicates the height of the reflecting cloud layer.

Bird control. Flying machines and flying animals don't mix. A collision with a seagull can crack a jet transport's windshield; "ingesting" a flock of smaller birds, such as starlings, can destroy an engine. And the grassy expanses of an airport make an attractive habitat for birds. A wide variety of solutions have been offered—which is a pretty sure indicator that none of them work. There are noisemakers and scarecrows, phony owls, mechanical hawks, and recordings of bird-distress calls. One of the common devices is the gas cannon, a stout pipe a few inches in diameter connected to a canister of propane just like the one that fires up the backyard barbecue. Gas trickles into the base of the pipe, then a spark from a battery-operated controller ignites it. The boom is of chest-thumping intensity. Nevertheless, gulls quickly learn to ignore it. If bloodless measures fail, the airport authorities may seek permission to use deadly force. They might set out poison-baited feeders; introduce foxes, skunks, or other predators on birds and their eggs; bring out trained falcons; or hire hunters with shotguns. (But the hunters have to be careful not to shoot the big silver birds.)

The twin pods staring at each other, above, measure runway visual range, or atmospheric clarity. Below is the most basic of all aviation weather instruments: the wind sock.

BEACONS AND BEAMS

The crucial difference between aviation and other modes of transport is that an airplane can't pull off to the side of the road or drop anchor when the pilot gets sleepy or the weather turns ugly. In an airplane, once the fuel runs out, you *will* come back to earth, ready or not. Hence, aviation puts a lot of emphasis on guiding the pilot to a safe landing even through darkness or cloud cover.

Runway Lighting. Landing lights do not shine down on a runway the way streetlights illuminate a road. Instead, the lights are embedded in the concrete and point upward, guiding the pilot to a safe touchdown. The lights are mounted flush with the surface of the runway, powered by buried cables. Airplane tires run over them all the time; the fixtures are built to withstand this abuse. (But tire rubber has to be cleaned off the lenses from time to time.)

The lights outline the runway, giving the nighttime pilot the same clues to position and orientation that painted markings offer in daylight. Along the centerline, 200-watt lights are installed every 50 feet. At the arrival end of the runway, the centerline lights are white. Starting at 3,000 feet before the departure end, alternate lights are red and white; in the last 1,000 feet they are all red.

The centerline lights are only the beginning of the runway lighting system. There are also edge lights, which are white over most of the length of the runway but yellow in the last 2,000 feet. Also, at each end of the runway is a row of lights installed across the full width of the pavement and spaced closely enough that from a distance they merge into a continuous bar of light; this threshold marker is green at the arrival end of the runway for a pilot about to land and red at the departure end for a pilot about to run out of room for takeoff.

Near the arrival end, still more rows of lights extend perpendicular to the runway centerline. There are six lights in each row, and there are typically 30 rows spaced 100 feet apart. The main purpose of these lights is "roll guidance"—helping the pilot to keep the wings level in the last few seconds before landing.

On the main runways of large airports the lighting system extends a further 3,000 feet ahead of the runway threshold. These approach lights consist of "barrettes" of five lamps erected at 100-foot intervals. The barrettes also hold flash tubes, which are fired in sequence to create the effect of a ball of light rolling toward the runway. This eye-catching display leaves no ambiguity about where to land. On a cross-country flight at night you can often spot these beacons far below and at great distances.

For lighting taxiways, the current preference is for small, flush-mounted green centerline lights, like a necklace of emeralds, but many airports still conform to an older standard, with blue edge lights mounted on foot-high stalks.

Yet another system of landing lights is called VASIS, for *visual approach slope indicator system*. What VASIS tells the pilot is not where the airport is but rather where the aircraft is with respect to the runway. Specifically, it conveys information about the glide slope, the angle at which the aircraft is descending toward the ground. The customary glide slope is three degrees—which works out to a descent of 277 feet per mile of forward motion—and VASIS warns if the descent is either steeper than this angle or shallower. Horizontal bars of lights are mounted on each side of the runway at 500 feet and 700 feet from the arrival threshold. These specially rigged lights have lenses and colored filters arranged so that each light emits a split beam, white in the upper segment and red below. When the pilot is on the correct glide slope, the more distant bars (the "upwind bars") appear red, and the nearer ones (the "downwind bars") are white. If the aircraft is too high, both bars are white; if too low, both are red.

From a passenger seat in a jetliner I have never managed to glimpse the VASIS lights changing color as the pilot adjusted the landing trajectory, but you can watch them from a smaller aircraft by sneaking a peek over the pilot's shoulder. On the

Lights of the deepest cobalt blue mark taxiways.

Approach lights held aloft on breakaway stanchions extend 3,000 feet ahead of the runway threshold. The airport is Raleigh-Durham in North Carolina.

ground, the light fixtures are easy to spot as your flight begins its takeoff roll. They are low structures off to either side of the runway, installed on concrete pads and painted—like nearly everything else on the airport grounds—bright red and white.

Instrument Landing Systems. Multicolored lights and painted markings aren't much use if the pilot can't see them because of fog or rain. The instrument landing system, or ILS, is meant to work in any weather. The pilot is guided by indicators on the control panel and looks up from these instruments only at the last minute, when the runway should be straight ahead and just below. It is an act of faith.

The ILS works much like VASIS, but instead of colored lights, the signals that guide the airplane are radio waves. The basic idea is to create a radio beam pointing upward at a shallow angle from the runway; the incoming aircraft slides right down the beam to the touchdown zone. But there's a problem with this idea. The radio beam has to come from an array of antennas, which are large metal structures that can't be buried in concrete the way a centerline light is. So how do you arrange to have the airplane fly toward the antennas but not crash into them?

The ingenious solution is to divide the ILS beam into two parts, for horizontal and vertical guidance, and build two antenna systems. The one that aligns the approaching airplane horizontally—keeping it from straying off to the left or right—has to be installed along the runway centerline, but longitudinally it doesn't have to be in the touchdown zone. In practice it's put beyond the far end of the runway. The beam for vertical guidance—keeping the aircraft descending at the correct rate—has to come from just the right distance down the runway, but it doesn't have to be centered laterally; the antenna is built off to one side of the runway at the touchdown zone.

The positioning of these antennas makes them easy to recognize. The horizontal beam, called the localizer, comes from an array of antennas that often stretches the full width of the runway, perpendicular to the centerline, several hundred feet beyond the threshold. The antenna elements vary in shape; some look like a ladder, some like a picket fence; most of the newer ones resemble a grape arbor, with irregularly spaced round bars poking out from both sides of a central beam.

The glide-slope antenna is mounted on a mast about 30 feet high, offset to the right or left of the runway far enough to be outside the wingspan of any aircraft. When your flight pulls into position for takeoff, the glide-slope antenna is up ahead a little ways, and so you pass it early in the takeoff roll; when you land, it ought to be pretty nearly abeam the aircraft at the moment the wheels thump down. The antenna elements are multiple bars or paperclip-shaped loops or bow ties, mounted in two or three groups at various heights along the mast, with mesh reflectors behind them.

A glide-slope antenna provides vertical guidance to pilots making a landing on instruments. It is mounted near the touchdown zone of a runway, but off to one side. The three trough-shaped reflector antennas on the tower project radio beams that help the pilot maintain a three-degree descent toward the runway. The installation shown is at Raleigh-Durham.

To say that the ILS produces a beam for an airplane to follow is a bit of an oversimplification. Both the horizontal path and the vertical one are really defined by the intersection of two beams. The localizer antenna radiates two signals at the same radio frequency but with different audio tones impressed on them; the left side of the fan-shaped pattern is modulated at 90 hertz and the right side at 150 hertz. The aircraft is on course when the two signals are received at equal intensity. Similarly, the glide-slope antenna transmits waves modulated at 90 hertz above the correct slope and at 150 hertz below; again, the pilot follows the boundary line between the two regions. Of course the pilot doesn't have to do all this by ear; instruments automate the process.

The radio frequency of the localizer beam is in the band from 108 to 112 megahertz, which happens to lie immediately above the standard FM broadcast band. You might be able to hear the signals by tuning your car radio to the top of its range as you drive under the approach path. (Older radios work best for this. The new ones with digital tuning are just too good at rejecting off-frequency signals.)

The localizer and glide-slope signals keep an aircraft on track and headed toward the runway, but the pilot also needs another piece of information: an indication of how much distance remains to the touchdown point. This third dimension is added by another set of radio beacons, called the outer, middle, and inner markers.

Each of the three marker beacons comes from a transmitter directly below the flight path. They radiate upward in a fan-shaped pattern oriented perpendicular to the approach path. The outer marker is about four miles from the runway threshold, where the incoming aircraft is expected to cross at an altitude of 1,200 feet. The middle marker is 3,000 feet from the threshold, where the aircraft should be at an altitude of 300 feet. The inner marker—which many airports have not installed—is about 750 feet from the threshold, corresponding to an altitude of about 50 feet.

Many marker-beacon antennas have a distinctive Y shape; the upraised and outstretched arms create a flat, fanlike radiation pattern. Sometimes the design is simpler, just a few horizontal bars or loops on a vertical mast. Some beacon antennas could

A localizer antenna provides the horizontal component of the instrument landing system. The antenna elements, which might be taken for some strange sort of gym equipment, point down the runway and keep an arriving aircraft on the centerline. (The aircraft above is departing rather than arriving and hence is not using the landing system.) The localizer in the photograph is at Reagan National Airport in Washington.

be mistaken for the kind of television aerial that used to decorate almost every rooftop in America. But there's a clue to the identity of the aviation markers: the antenna will always be painted either solid red or red and white—the only colors the FAA buys.

HIGHWAYS IN THE SKY

The instrument landing system guides an aircraft only in the last few minutes of its flight, but other aids to navigation provide a frame of reference all the way from take-off to landing. These electromagnetic landmarks create an invisible network of aerial routes that spans the North American continent and much of the rest of the world.

In the early years of aviation, aids to navigation were all visual—town names painted on factory roofs, and a network of beacon lights marking recommended routes through mountain passes. A remnant of the beacon-light system survives today.

Marker beacons are another component of the instrument landing system, which tell the pilot how much distance remains to the runway threshold. Above is an outer marker, four miles away from the touchdown point, on one of the approach paths to Dulles airport in Virginia. Below is a middle marker, 3,000 feet out from the threshold.

Most airports have a rotating beacon, green on one side and white on the other, so what you see from a distance is a sequence of alternating green and white flashes. At military airfields the white beam is split, creating a double white flash. At a seaplane airport the beams are yellow and white.

But visual signals have mostly been supplanted by radio-frequency "navaids." The simplest of these devices is called a nondirectional beacon, or NDB. As the name suggests, the beacon emits a signal that's the same in all directions. An instrument in the aircraft homes in on the beacon by turning a small loop antenna until the signal is strongest. From this information, the pilot can determine the bearing of the beacon with respect to the heading of the aircraft. For example, if you are flying straight toward the beacon, its bearing is 0 degrees; at a bearing of 90 degrees, the beacon is off your right wingtip. If you know the bearings to two or more beacons, you can draw intersecting lines on a chart to fix your own position. There are more than 1,500 NDBs operating in the United States.

The typical ground installation for a nondirectional beacon is a vertical tower about 30 feet tall, sometimes with an umbrella-like arrangement of ribs and spiral wires at the top. High-power beacons may use a horizontal dipole antenna—a wire or a set of parallel wires stretched between two poles. Any of these arrangements could be mistaken for the transmitting antenna of a small radio station, or perhaps the communication facilities of a police department or taxi company—if it weren't for the FAA's do-not-touch sign.

Nondirectional beacons broadcast at low frequencies—between 200 and 500 kilohertz, well below the AM broadcast band. But as you get close to the transmitter, you may be able to hear the signals anyway, even on an ordinary car radio. Spin the dial across the AM band and listen for a few letters of Morse code repeated over and over. The signal you pick up will be at double or triple the beacon's assigned frequency. The Morse code is an identifier for the beacon. Every navaid in the world is assigned a unique label, usually consisting of three letters; for a beacon at an airport, the letters are the same as those you find on your luggage tags.

The alternative to a nondirectional beacon surely ought to be called a directional beacon, but instead it's called a VOR or a VORTAC. Here we enter the dizzy world of aviation acronyms, which fill whole dictionaries. *VOR* stands for *VHF omnidirectional range*, where *VHF* in turn stands for *very high frequency*. A VORTAC combines a VOR with the military navaid called a TACAN, which is an abbreviation of *tactical air navigation*. Note that *VOR* is pronounced as three separate letters, but *VORTAC* is said as a word.

Whatever you call it, a VOR supplies directional information by sending out two carefully timed signals on the same carrier wave. To understand how it works, it's helpful to think in terms of light beams rather than radio waves. Imagine a lighthouse with a beam sweeping clockwise around the landscape, and a separate strobe light that sends out brief flashes in all directions at once. The strobe fires once per minute, and the rotating beam takes exactly a minute to sweep a full circle; furthermore, the flashing and the rotating are coordinated so that the flash always comes at the instant when the rotating beam is pointing due north. With this system, determining your direction with respect to the lighthouse would require nothing more than a stop-watch. You start the watch ticking when you see the flash of the strobe light, and you stop the watch when the rotating beam sweeps past you. If the delay is 15 seconds, you are due east of the lighthouse; if 30 seconds, you are south; and so on.

In a VOR, the rotating lighthouse beacon is a radio beam that sweeps around the horizon 30 times a second, or 1,800 times per minute. The synchronizing signal—the conceptual equivalent of the strobe flash—is encoded in a warbling modulation of the same radio wave, with one cycle for each revolution. A special-purpose receiver in the aircraft disentangles the timed signals so that an instrument in the cockpit can simply point toward the VOR transmitter.

STRIPES AND CHECKS

Aviation leaves its mark on the land even miles from the nearest airport. Under regulations enforced by the Federal Aviation Administration, anything poking up into the sky far enough to snag an airplane gets painted red and white, or else it is decorated with blinking marker lights, or both. (Actually, the "red" paint has the official name aviation orange. It's the color of a ripe tomato.)

How high does a structure have to be before the FAA starts choosing your color scheme? That depends on where you are. On the airfield itself, an equipment shed the size of a doghouse or a knee-high lighting stanchion may get the red-paint treatment. At the airport boundary line, the height limit is usually 50 feet. Elsewhere, anything taller than 200 feet is a likely candidate for the paintbrush. This includes radio and television broadcast antennas, water towers, smokestacks, large oil and gas tanks, and high bridges. The standard treatment is to paint either the entire structure or the upper part of it in red-and-white stripes or checks. This pattern makes the object stand out in daylight. For night visibility, there are slowly winking red lights. In some cases only the lights are required; this is the common practice with tall buildings, which have enough bulk to be visible by day without the garish paint scheme. (The New York City skyline would certainly be a different place if every building had to be decked out in FAA plaid.)

In recent years the FAA has introduced an alternative to the red-paint and red-light

scheme. Many tall towers are now equipped with high-intensity strobe lights that emit brief but brilliant flashes of white light. The strobe flashes are easily visible by day as well as at night, and so there's no need for special paint. As a matter of fact, the strobes are so bright they have to be dimmed at night, or pilots might be dazzled by them.

A tall antenna tower will have three to five strobe units at roughly equal intervals along its height. They all flash simultaneously, at a rate of about 40 times a minute. A less common spectacle is a set of strobe lights timed to flash in rapid sequence from top to bottom so that a ball of light appears to be descending the tower. This distinctive marking is used on nearby pairs of towers to warn pilots that something is strung between them—usually a high-voltage power line—and so it wouldn't be a smart idea to fly through the middle.

At left is a water tower on the grounds of Andrews Air Force Base, just outside of Washington, decked out in gay FAA stripes and checks. Since this photograph was made, however, the tower has been repainted plain white. A flashing strobe light mounted at the top of the tank is all that remains to warn off low-flying pilots.

On the ground, a VOR installation looks like a broad-brimmed hat—usually a conical witch's hat, but sometimes a cylindrical, Abe Lincoln–style top hat. The brim is a circular reflector of solid metal or mesh that bounces radio signals away from the earth and toward the sky. The cone or cylinder in the middle—made of radio-transparent fiberglass—shields an antenna spinning at 1,800 revolutions per minute. If you can get close enough, you may be able to hear the motor whirring. At the edge of the hat brim there may be a dozen or so podlike attachments; they are antennas for monitoring and calibrating the VOR emissions. The whole apparatus usually sits atop a shed housing the transmitters and power supplies.

A newer style of VOR transmitter, called Doppler VOR, dispenses with the mechanical rotating parts (which are high-maintenance items). Instead of one spinning antenna, it has numerous stationary antennas arranged in a circle around a cen-

It's not a UFO that landed overnight in the Nevada desert. It's the Coaldale VORTAC, midway between Las Vegas and Reno, with the Silver Peak Mountains in the background. Inside the conical hat is an antenna twirling at 1,800 revolutions per minute, broadcasting a signal that aircraft instruments interpret as a directional beacon. The flat brim of the hat is a metal reflector, beaming the signal upward. Pods around the perimeter of the brim monitor and calibrate the signal.

tral emitter. The antenna in the middle radiates at all times, and the perimeter elements are activated in opposite pairs in a pattern that rotates at the standard 30-revolutions-per-second rate. For the airborne receiver, this electronic rotation is indistinguishable from a standard VOR signal. The circular array of antennas in a Doppler VOR would be hard to mistake for anything else. There are 52 antenna elements placed in a circle 44 feet in diameter, mounted on stalks above a metal reflector plane.

Much of the infrastructure for aerial navigation relies on machines that could qualify as antiques. The technology of the VOR dates back to the 1950s, and so does the ILS. Some of the radars and computers in air-traffic control centers are more than 30 years old. The FAA has been struggling to catch up, but technology has been moving faster than the agency can. For example, a replacement for the ILS, called the microwave landing system or MLS, was considered obsolete by the time the FAA finished testing it, and so the project was canceled. In the long run, all of the navaids described here will surely disappear. A long-range radio-navigation system called Omega has already been shut down. Another system called Loran (discussed in Chapter 12) is on the endangered-species list. According to one tentative plan, VOR and ILS services will be withdrawn sometime before 2010.

What will replace all these vintage technologies? The way of the future is satellite navigation. The reason becomes obvious when you consider that a passenger sitting in the back of an airplane, using a handheld GPS receiver that costs a few hundred dollars, can track the progress of a flight just as accurately as the pilots up front, who rely on multimillion-dollar avionics. All the thousands of VORs and NDBs worldwide can be replaced by a few dozen satellites in the Global Positioning System. Already the satellite signals are sometimes used for en route navigation. With just a little enhancement they'll be accurate enough to guide a plane right down to the runway threshold. (Some people worry about making transport so dependent on satellites, which are vulnerable to various kinds of attack and interference.)

Air Routes. The three-dimensional freedom of the skies is what's most alluring about aviation; you roam through the wild blue yonder, free as a bird. And yet, in practice, most flying is not so unfettered. Large tracts of airspace are regulated, and you may enter them only if you agree to abide by the instructions of air-traffic controllers. And most flights do not head out willy-nilly across the landscape. They follow well-worn tracks through the sky, invisible highways, where pilots are expected to stay in the correct lane, signal their turns, and obey traffic signs.

There are two systems of air routes over North America. The Victor airways are analogous to local blacktop roads. They get their name because they are designated V-21, V-44, and the like. *Victor* is pilot lingo for the letter *V*, as in VOR. The jet routes, with names such as J-14 or J-39, are the Interstate highways of the air, used by longer-range, higher-flying, and faster aircraft.

The Victor routes, at altitudes of 1,200 to 18,000 feet, run in straight lines from one VOR to another, so that you can lock onto a specific "radial" and then simply

Doppler VOR, above, emits essentially the same signals but without the spinning antenna. The telltale mark of the Doppler VOR is a ring of 52 podlike antennas mounted above a metal reflector plane. The tall stalk next to the VOR is the antenna of a distance measuring equipment (DME) station. Whereas signals from the VOR tell the pilot what direction to steer to reach the site, those from the DME indicate how far away it is. A nondirectional beacon, below, is an even simpler navigational device. It merely broadcasts an identifying signal uniformly in all directions; a direction-finding antenna in the aircraft is needed to estimate the bearing.

The pagoda control tower at Dulles airport, designed by Eero Saarinen, is one of the most admired towers in aviation. Nevertheless, it has apparently outlived its usefulness. With ongoing expansion of the airport, Dulles has undertaken to build a new, higher tower. On the opposite page are some other distinctive towers: the bipedal design at Logan airport in Boston and a round glass tower at Charles de Gaulle in Paris, topped by a very French beret.

follow that course until the plane passes directly over the navaid. The jet routes are at altitudes of 18,000 feet and up. They also run from one VOR to another, but generally they make longer hops.

Air routes are multilane highways, but the lanes are not side by side as they are on a terrestrial freeway; they are stacked up one above the other. Furthermore, lanes going in opposite directions are interlaced. For example, a flight heading east on a certain air route might be assigned an altitude of 9,000 or 11,000 or 13,000 feet, whereas a westbound flight would have to choose 10,000 or 12,000 or 14,000 feet.

Although air routes are invisible, if you're patient you can sometimes notice aircraft tracing them out across the sky. Years ago I used to bicycle out to Farmington, Minnesota, south of the Twin Cities, where several air routes converge over a VOR at a major air-traffic-control facility. When the weather was right, high-altitude jets left condensation trails chalked against the dry midwestern sky. As the aircraft switched from an inbound to an outbound radial, I would see dozens of white trails pivoting on the same invisible pole directly overhead.

But aerial navigation is not the rigid, point-to-point discipline it used to be. For a long time pilots had to hop from one navaid to the next because it was too hard to calculate a direct course from departure to destination. With the simple cockpit instruments of earlier aircraft, you could steer straight toward a VOR or directly away from it, but more complicated routings required too much figuring. All that has changed with the advent of computers in the cockpit. Now a computerized instrument called a flight director or flight-management system continually monitors signals from several VORs and plots the aircraft's course with respect to all of them. The system is called RNAV, which is the aviation industry's strange way of abbreviating *area navigation* (or maybe it means *random navigation*—sources differ). When you hear your pilot announce, "We've been cleared direct from Baltimore to Cleveland," that word *direct* probably means you'll be following an RNAV route instead of skipping from one VOR to the next.

Air-Traffic Control. The controller, with a headset clamped over his ears, peers into the radar scope and calmly reassures the 12-year-old who has taken the controls of the airplane after the entire flight crew became stricken with ptomaine poisoning. How many times has this scene been played in the movies? Surely more than it has in real life. And yet it is not outside the job description of an air-traffic controller.

The airport landmark associated with air-traffic control is the tower, although only a small percentage of all controllers work in the glass-walled room at the top of the tower. This room is called the cab. It is where controllers handle flights that are just about to land or are ready to take off, as well as aircraft on the taxiways—basically all moving aircraft that are within sight of the tower. It's for this reason, of course, that the tower is a tower: the controllers have to be able to see everywhere on the airport grounds, which necessarily means that the tower itself can be seen from just about everywhere. The bigger the airport, the taller the tower. The one at the new Denver International Airport is 327 feet high.

Usually the windows of the cab are tilted inward at the bottom to avoid distracting reflections, or the glass may have a curved profile. (The same trick is sometimes employed with department-store display windows.) Sunscreens of various colors and degrees of transparency can be pulled down like window shades when needed.

All towers are tall, but apart from that they are anything but standardized. Airport architects have not been shy about making these structures distinctive as well as conspicuous. There is the two-legged Colossus of Logan in Boston, and a bright new minaret at National Airport in Washington; nearby Dulles International's tower is a kind of elevated pagoda; the new tower at Kennedy International in New York looks like an axe poised to fall; Chicago O'Hare has a golf tee covered with bright ceramic tiles. Some other designs look a bit too much like the guard towers of a prison.

The roof of the tower often has a small forest of antennas. These are mostly simple whip antennas used for radio communication with aircraft and with various ground vehicles (fuel trucks, snowplows, fire trucks). There may also be a rotating

radar antenna on the roof, perhaps enclosed in a protective radome. And the spinning green-and-white beacon light may be up there too. The tallest structure on the airport grounds is an obvious spot for installing all these items.

Most controllers track aircraft with radar rather than binoculars, and their surroundings are just the opposite of those in the tower cab; they live in a dim, windowless room, the better to see the blips on the screen. In some cases the radar room is on a tower floor below the cab, but it could also be in a nondescript building elsewhere on the airport grounds, or even some miles away. The controllers here work with aircraft at a somewhat greater range than their colleagues in the tower cab. They handle approach and departure patterns, making sure that all the flights converging on the airport and radiating from it don't get in each other's way. It is the approach controller who puts you in a holding pattern when traffic backs up.

There is still another cadre of controllers whose place of business is generally nowhere near the airport. They work in en route centers, directing traffic in the wide-open spaces at cruising altitude. There are 20 en route centers in the United States. They are mainly inconspicuous government buildings—there's no need for a tall tower—although some of them are located near VORs or other major navigational facilities. One distinctive feature you may spot near an en route center is a mast with a ladder-like array of long dipole antennas. Typically there are about 20 antenna elements, and the biggest of them is roughly 50 feet long. The whole contraption looks like a monstrous rooftop TV aerial. It is used for voice radio communications.

Controllers sometimes help with certain aspects of navigation—on request they will issue "vectors" telling a pilot which way to steer—but their main function is to

Still more airport control towers: at Venice *(top)*, at O'Hare in Chicago *(above)*, and at St. Louis *(right)*.

SMALL-AIRPORT ETIQUETTE

Instrument landing systems and stroboscopic runway lights are standard fixtures at airports serving commercial airlines, but there are also hundreds of smaller airports that have no need for such high-tech facilities. At some of these airports, the most important maintenance task is mowing the runway.

Protocol at a small airport is very different from the rules that govern metropolitan flight operations. Here there is no control tower, and there are no air-traffic controllers keeping watch on radar scopes to prevent collisions. With no central authority to adjudicate the right-of-way, pilots must negotiate among themselves, much as drivers do at an intersection without traffic signals. If you watch the skies for an hour on a sunny Saturday afternoon, you may be able to figure out the etiquette of your local airstrip.

Most often, landing aircraft follow a three-legged, left-handed pattern. An approaching plane enters the pattern on the downwind leg, parallel to the active runway but in the direction opposite to the landing (and takeoff) direction. The aircraft makes a left turn, and then a second left turn lines it up with the runway centerline for the final approach and landing.

How does the pilot know which runway is active and which way to land on it? Even airports without a control tower may have a radio service offering such information, but there is also a visual indicator. It is called the segmented circle (segmented because it is drawn with a dashed line). Look for it at a conspicuous spot on the landing field, such as near the intersection of two main runways. It's at least 100 feet in diameter, and the individual segments making up the circle are at least three feet wide and six feet long. Most often the segments are painted markings on a paved surface, but sometimes they are beds of white stone laid in a grassy field. (It can look a little like Stonehenge, and no doubt future archaeologists will speculate about its ritual significance.)

In the middle of the circle, mounted on a mast, is the wind sock. This device looks just like what you would expect from its name: a long cloth cone held open at one end by a hoop, and allowed to swivel so that the opening always faces into the wind. A sock is preferred over a simple weathervane in that it indicates the force as well as the direction of the wind. As the wind grows stronger, air fills more of the sock and extends it to a greater length. A fully horizontal wind sock indicates a wind of about 30 miles an hour.

In the absence of any more definitive information, an arriving pilot will prefer to land into the wind and will look at the wind sock to choose a runway accordingly. But some airports have a device within the segmented circle to visually indicate the current active runway. It is a T-shaped or arrow-shaped sign made of wood or metal, big enough to be seen clearly from a few thousand feet up and mounted under the windsock. The arrow or the long arm of the T is turned so that it points along the active runway. If the wind shifts, someone has to go out on the field and spin the indicator around to establish a new traffic pattern.

Around the perimeter of the segmented circle, you may see a few L-shaped markers oriented in various ways. These tell the arriving pilot what kind of an approach pattern is customary for each runway—whether to make all right turns or all left turns when approaching the runway.

make sure that aircraft don't run into each other. Aircraft flying at the same altitude are kept at least three miles apart. When flight paths cross each other, the minimum vertical separation is 500 feet.

Radar. The revolving radar antenna is one of the most readily recognized features you're likely to spot at the airport. It's also one of those objects with the pleasing property that how it looks tells you what it does. You can learn a lot about a radar's purpose and operation just from the appearance of the antenna.

The idea of radar is simple, although getting it to work reliably was an engineering challenge that took the better part of a century. The twirling radar antenna sweeps a narrow beam of radio waves around the horizon, then listens for echoes reflected from aircraft or other "targets." From the echoes, the radar unit determines both the direction and the range of the target. (The term *radar* began as an acronym for *radio direction and range*.) The direction part is easy: it is the direction the antenna

Voice communication between pilots and controllers on the ground is the function of much radio hardware you might spot near airports. At top, a cluster of five towers in North Carolina provides a link between aircraft en route and a regional air traffic control center. The structure in the middle photo looks much like a rooftop television antenna, but it is built on a larger scale; the longest of the crosspieces is about 50 feet long. At the bottom is a direction-finding antenna. Both of the latter devices are installed near MacArthur Airport in Islip, Long Island.

is pointing at the moment the echo arrives. To calculate the distance, the radar must measure the time needed for a radio pulse to make the round trip from antenna to target and back.

Two factors make the job of the radar difficult. First, only a minuscule fraction of the transmitted signal is reflected back to the antenna. Thus, the system has to emit megawatts and detect microwatts. Second, because the signals move at the speed of light, the timing of the radio pulses has to be controlled and measured with an accuracy of microseconds. (Bouncing a pulse off a target one nautical mile away takes 12.36 millionths of a second.)

When you look at a radar antenna, most of what you're seeing is a curved reflector that forms the radio signal into a narrow beam in much the same way as the reflector in a flashlight does. The shape of the reflector can tell you something about the shape of the radar beam. The larger the reflector in any one dimension, the more narrowly it focuses the beam in that direction. In most radars for air-traffic control, the antenna reflector is a low rectangle, longer than it is high. As a result the beam is strongly focused horizontally but not vertically: it is a thin wedge as seen from overhead but a broad fan when viewed from the side. This beam geometry allows the radar to locate aircraft precisely in the horizontal plane while sweeping up targets at all altitudes.

For some short-range radars the antenna is more a bar than a rectangle; it looks something like a rotating log. At airports these radars are often used for ground surveillance—that is, for keeping track of aircraft on runways and taxiways. Similar radar units are installed on boats and ships.

More often than not, the reflective surface of a radar antenna is a lattice or mesh—like a window screen or chicken wire fence—rather than a solid panel. A reflector full of holes wouldn't work very well for a flashlight, so how does it manage for radar? Doesn't the beam just leak through the sieve? The explanation lies in the difference between light and radio signals. For a reflector, holes and other irregularities don't matter much as long as they are smaller than the waves being reflected. For visible light, that means the surface has to be smoothly polished down to a scale of well under a millionth of an inch. But most radars operate at wavelengths of a few inches or even a foot or more, and so the openings in the mesh reflector cause no trouble. Roughly speaking, the waves are too big to pass through them. By reducing weight and wind resistance, the mesh eases the burden on the mechanical system that twirls the antenna.

The parabolic reflecting surface is the largest and most conspicuous part of a radar antenna, but the real business end of the system is a much smaller device called the feed horn. It is mounted in front of the reflector (that is, on the concave side) and rotates with it. If you can get a clear look—this is easier if the radar is shut down and not spinning madly around—you'll see a metal pipe connected to the feed horn. The pipe is called a wave guide, and it carries the radio-frequency signals from the transmitter to the antenna and then from the antenna back to the receiver.

One of the tricky parts of building a radar is connecting the rotating antenna to the stationary transmitter and receiver. The crucial component is called the rotary coupler joint; like something in a lawn sprinkler, it has to be loose enough to allow rotation but tight enough to prevent leaks. It tends to be a high-maintenance item.

The radar unit you are most likely to see at an airport is a short-range surveillance radar, also known as a terminal control radar. "Short range" means 50 to 100 miles; it covers all the traffic arriving at and departing from the airport, or perhaps from a

Airport surveillance radar keeps track of aircraft within 50 or 60 miles of an airport. The unit shown here is on the grounds of Miami International Airport. The most important components have helpfully been painted aviation orange for ease of identification. Atop the tower is the rotating antenna, which turns at 12 revolutions per minute. The curved, concave reflector is for the primary radar, which bounces microwave pulses off the surface of airplanes and detects the faint echoes. Note the two "feed horns" pointing into the reflector; they emit and receive the signals. The flat "picket fence" above the primary antenna interrogates a transponder installed on all larger aircraft; the signals returned by this device allow the radar display to identify individual aircraft and give their altitude. The two orange tubes installed along the near side of the tower are wave guides that carry microwave signals between the antenna and an electronics shed at ground level.

Long-range surveillance radar, reaching out several hundred miles, requires larger antennas, which are almost always protected from the weather by a "radome." The site shown here, on a hilltop near Tonopah, Nevada, was built for military purposes, but at the time this photograph was made, in 2000, was being operated by the FAA.

cluster of nearby airports (as in New York, where there are three major airports within an area about 20 miles across). An important characteristic of a radar is the sweep rate, the speed at which the antenna rotates. A typical airport radar turns at roughly 12 revolutions per minute, or 5 seconds per sweep. (This means that the controller sees airplanes moving in jumps every 5 seconds.)

All radars, incidentally, rotate clockwise. As far as I can tell, there's no law of man or nature that requires them all to turn in the same direction, but that's the way it is. They go the same way in the Southern Hemisphere as in the Northern. So if you ever see one going counterclockwise, you'll know you have crossed over into a mirror universe.

Many airport radars have a second antenna—either a flat mesh screen or an elongated "hog trough"—mounted above the main reflector and rotating with it. This is part of a beacon-and-transponder system that allows aircraft to identify themselves individually on the controller's radar screen. A "passive" radar merely bounces signals off the skin of an aircraft and detects the faint echoes. With nothing more than this to go on, a hang glider and a 747 look much the same. Furthermore, the radar echoes carry no indication of altitude. The beacon and transponder solve both of these problems. The transponder is installed in the aircraft. When the radar beam sweeps past it, the transponder broadcasts a coded message that identifies the aircraft and gives the current reading of the altimeter. With this information available, the controller's display can label each blip of light with a flight number and an altitude.

The radars that keep an eye on ground movements of aircraft need a range of only a few miles, but they have to give controllers a detailed map, accurate to within a few feet. The antennas are small, rapidly spinning, and often mounted atop the control tower, since the radar, like the controllers themselves, needs to have an unobstructed line of sight to all important areas of the airport.

Aircraft at high altitudes are tracked by long-range surveillance radars, which can reach a radius of 250 miles or more. It takes about 20 such radar installations to cover the lower 48 states. The antennas are larger and rotate slower: about 5 revolutions per minute instead of 12. Unfortunately, you are unlikely to see any details of the antenna because it is usually enclosed in a radome, a housing that protects the hardware from wind and weather. (A rotating antenna the size of a large billboard is not something you'd want to have out in a gale.) The radome looks like a giant golf ball or soccer ball and is made of fiberglass or a plastic transparent to radio waves. The entire structure has to be assembled without metal components, which would produce strong reflections bouncing around inside the dome.

Radomes like this one in Pennsauken, New Jersey, also house antennas for weather radar.

In recent years, smaller radars have also been covered up by radomes, mostly to reduce maintenance costs. The latest thing is the spinning radome, often seen atop an airport control tower housing the ground-control radar. Surprisingly—and perhaps disappointingly—the spinning of the radome has nothing to do with the rotation of the antenna inside. The radome's rapid twirling is meant merely to shed rain, snow, and ice, which can distort the radar signals.

Some of the fanciest radars have no moving parts at all. They are called phase-array radars: the antenna is an array of hundreds of small antenna elements, and the beam is steered by electronically controlling the phase of the signals fed to these elements. A phased-array radar can flick its beam across the sky much faster than a mechanically steered antenna, but building big arrays is expensive. Some aircraft have a small phased-array antenna in the nose, hidden behind an aerodynamically shaped radome. On the ground, most phased-array radars are operated by the military. The big ones that stand on alert for missiles arcing over the North Pole look appropriately ominous and mysterious—pyramids in an arctic desert.

The wavelengths used by radar systems are the same as those of microwave ovens. (Indeed, an early brand of microwave oven was the Radar Range.) The power output of the radars is thousands of times greater than that of the kitchen appliance where you pop your popcorn. The hazard this suggests is quite real. If you stood directly in the beam emitted by a large radar feed horn, you too would pop. But the intensity falls off rapidly with distance. Stay outside fenced areas and you should be safe.