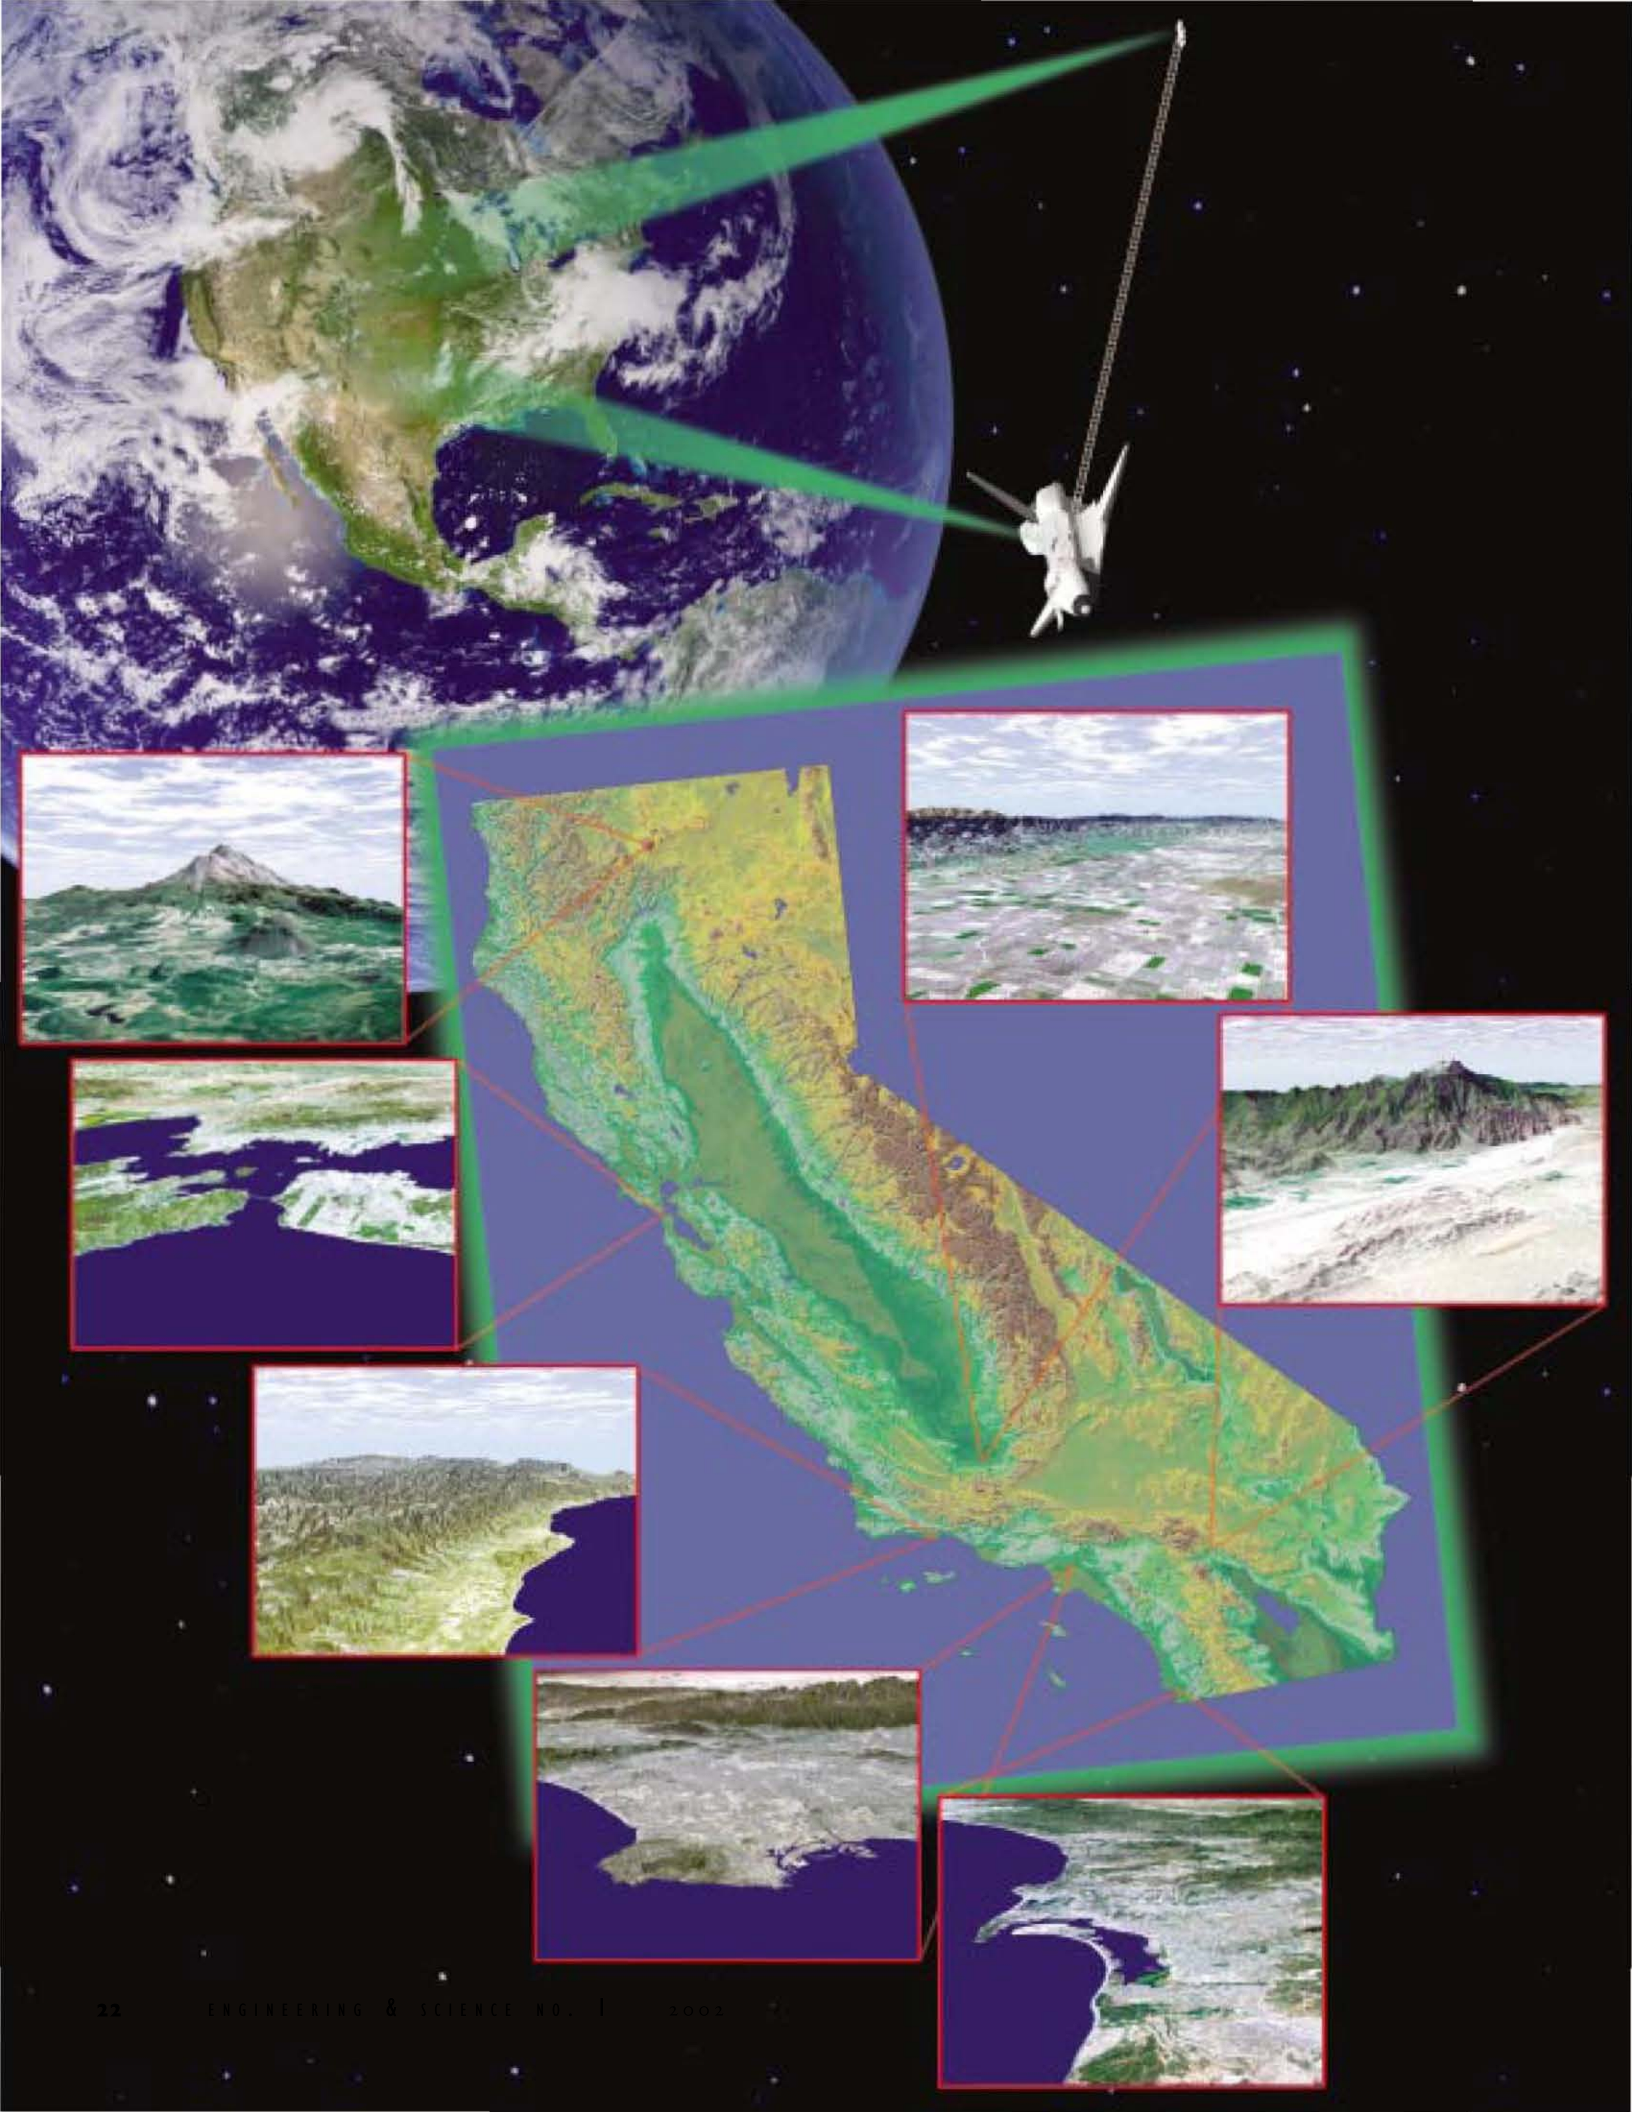

Planetary Phrenology: The Lumps and Bumps of the Earth

by Michael Kobrick

Phrenology is the study of the shape and protuberances of the skull, based on the belief that they reveal character and mental capacity. It was very popular in Victorian times, but has since been discredited as a scientific way of understanding the mind. On the other hand, studying the lumps and bumps of the earth to work out what's going on inside is solid science: the shape of the outside of the earth—its topography—really is the key to what's going on underneath. Topography controls the flow of surface and subsurface water, provides clues to the structure of the earth's crust, gives us places to put antennas, sometimes falls on us during earthquakes, and allowed the invention of skiing.

In February 2000, the Shuttle Radar Topography Mission (SRTM) aboard the space shuttle *Endeavour* measured that topography, and I'd like to tell you something of how the mission was conceived, how it was done, and how the data that we got from it are being used.

But first let me somewhat rescale your thinking about the earth's topography. If you were asked to name the highest point on the earth you'd probably say Mount Everest, which is 29,029 feet (8,850 meters) above sea level. But the top of Everest is not the farthest point from the center of the earth; that's an honor held by a volcano in Ecuador called Chimborazo. Although it's only 20,561 feet (6,267 meters) above sea level, the top of Chimborazo is farther from the center of the earth than the top of Everest, because the earth bulges at the equator. This bulge results from the earth's rotation on its axis, which is pretty fast: if you're reading this at the latitude of Pasadena, California, you're actually speeding toward the east at almost 900 miles per hour—so it's a good thing the atmosphere is going with you. Since the earth is not infinitely rigid, it bulges at the equator (picture a spinning ball of Silly Putty), and by a surprising amount. The difference between the diameter of the earth measured through the poles and the diameter measured through the equator is 25 miles (40 kilometers), almost the length of a marathon.

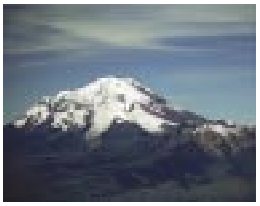

Chimborazo

Since Chimborazo is very near the equator, the top is actually 7,054 feet (2,150 meters) farther from the center of the earth than the top of Everest, up there at latitude 28 degrees north. (This information could be helpful if you're ever a contestant on *Jeopardy*.)

The ocean shares this equatorial bulge, but it also has other lumps and bumps all over the place, resulting from disparities in the earth's gravity field caused by variations in the density of its outer crust. You can see those bumps in the map of the earth at the top of the following page, in which I've color coded the oceans by height according to the earth gravity model. If you sail from the east coast of Africa to Malaysia, you're really sailing down into a hole that's about 100 meters deep and back up the other side.

The oceans cover about 70 percent of our planet, so they hide a lot of the earth's topography. In the picture of the earth at the bottom of the next page, the water has been stripped off and I've added maps of the other planets for which we have

Left: The color-coded, shaded relief map of California is a mosaic of 60 of the over 14,000 "cells" that will ultimately be generated by the SRTM when all the data have been reduced. Each cell covers an area of 1 degree longitude by 1 degree latitude. The flyover views around it are a combination of SRTM topographic data and Landsat photos, and show (clockwise from top right) the southern end of the agricultural San Joaquin Valley, Palm Springs, San Diego, the Los Angeles basin, Santa Barbara, San Francisco, and Mount Shasta, one of the highest volcanoes in the United States.

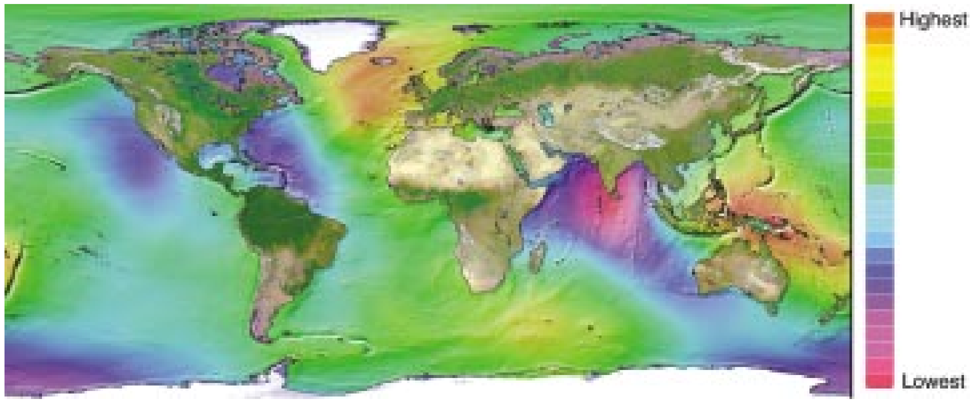

Sea levels are not uniform around the globe because of variations in the earth's gravity field. Color coding sea levels by height reveals "holes" like the 100-meter depression in the Indian Ocean.

reasonably good topographic data for comparison: the moon, Venus, and Mars. I've color coded them all in the same way, and used exactly the same scale for all four surfaces. You can see they look quite different. On the right of each panel is a bar chart showing the percentage of the planet's surface at each elevation, with the highest elevation at the top and the lowest at the bottom. The moon, Venus, and Mars have quite a broad range of elevations (if you ignore the spikes on the Mars bar chart, which are just the result of artifacts in the data set) but the earth is unique—it's split into two. That spike near the top of the bar chart is the range of elevations of the continental surfaces, and as you can see, it's very narrow. There's not much difference in height between one land surface and another—the continents are pretty smooth. You can also see this in the color-coded map—the earth doesn't have as many yellow and orange areas as the other planets, does it? The lower bump in the earth's bar chart, the range of elevations of the ocean bottom, is different. It's much broader, like that of the moon or Venus. And that's because the ocean floor is much more uneven.

The fundamental reason for this dichotomy is age. The continents float around like corks on the outer crust, bumping into each other, moving apart; and they're very old, nearly the same age as the earth itself, about 4.5 billion years. Although new mountains do get pushed up on the continents in various ways, 4 billion years of erosion by rain, wind, and glaciers have weathered them down so much that they're almost flat. The ocean basins, on the other hand, are quite young. In places like the mid-Atlantic ridge, or over hot spots like Hawaii, they're constantly being reborn; stuff is coming up from the upper mantle of the earth, displacing older rocks, and these older rocks eventually plunge back into the mantle. The average age of the ocean basins is only a few hundred million years, which is young in comparison to the elderly continents. And since it doesn't rain on them, they don't erode, so topographically they're very much rougher.

If you could hold the whole world in your hands, would you feel any of this topography? Well, if the earth was the size of a desk globe, Everest would be about three-tenths of a millimeter high. You wouldn't feel it. In fact, at this

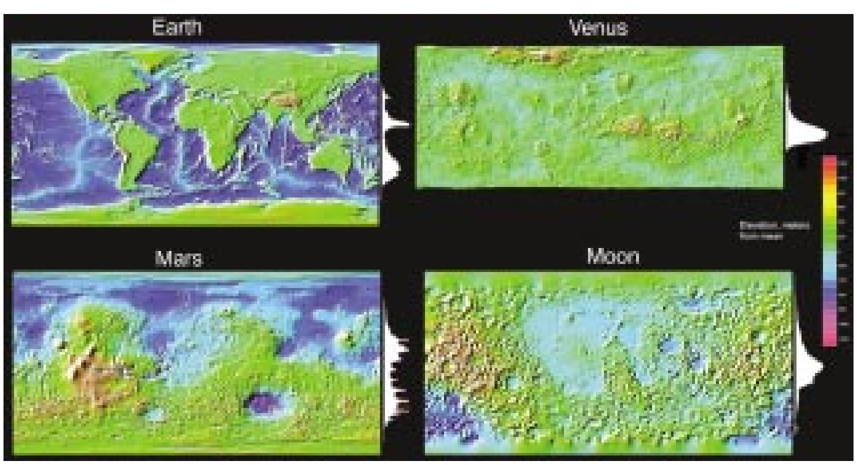

The four planet-sized bodies in the solar system for which we have good topographic data, viewed at the same vertical scale and with the same color coding. Bar charts of elevation are on the right of each panel. Earth's continents clearly have the smoothest surface (that we know of) in the solar system.

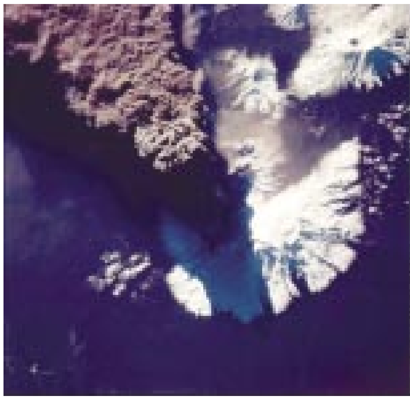

The volcano Kliuchevskoi on the Kamchatka peninsula erupted in October 1994, and was photographed (far left) by the shuttle astronauts and the Spaceborne Imaging Radar (near left), SRTM's predecessor. The radar saw right through clouds and volcanic plumes, and it didn't matter that the sun was setting at the time.

scale, the earth would feel like a smooth plastic beach ball!

At the Jet Propulsion Laboratory I work on remote sensing of the earth and planets using imaging radar. The word radar usually calls to mind big rotating parabolic dishes and air-traffic controllers staring at round screens with blips on them. Well, we've gone a bit beyond that—we can now use radar to take pictures of the surface of planets. Modern imaging radars take pictures that are often indistinguishable from photographs.

The imaging radar we flew on two space-shuttle missions in 1994 consisted of a phased-array radar antenna made up of many hundreds of little individual transmitters and receivers (called TR modules) distributed around the face of a structure looking a whole lot like a billboard (right). By adjusting the phases of those little transmitters, the direction in which the radar beam points can be changed. The shuttle flew upside down with the "billboard" pointing off to one side, and the radar beam swept out an image swath along the ground. It wasn't a continuous beam, of course, but consisted of a series of pulses. A pulse about 33 microseconds long was emitted, hit the ground, and bounced back. The echo was recorded, and another pulse emitted. This radar pulsed about 1,500 pulses a second, considerably faster than the human persistence of vision; if it were a series of light flashes it would look like a continuous spotlight. Where the ground was rough, as in a city (buildings, corners, and metal things reflect radar very well), a lot of the energy got reflected back toward the antenna, and we coded that as bright in the image. When radar waves hit surfaces that were flat and smooth, like freeways (eight lanes of concrete), they bounced off in a forward direction, and very little got back to the antenna, so they looked dark.

Radars can see right through clouds, dust, and smoke because they use wavelengths of centimeters or longer, in this case 5.6 centimeters. That's about the length of your little finger, much bigger than the particles that make up clouds, smoke, or even the smog that sometimes obscures the Los Angeles basin. When ocean waves hit a rock that's much smaller than the wave, they don't even notice, they just march right on by, and it's the same with radar waves: they go right through the clouds. It also doesn't matter whether the sun is shining or not, because the radar provides its own illumination, and images can be taken day and night. During the second of the two 1994 flights, there was a volcanic eruption on the Kamchatka peninsula. The picture top left is a photograph taken by the astronauts as the shuttle flew over, and next to it is a picture taken by the shuttle radar. The radar could see right through the smoke plume.

The crew of *Endeavour* visiting JPL's spacecraft assembly facility before the February 2000 mission. Behind them is the huge "billboard" radar antenna, a veteran of two previous flights in 1994.

The two 1994 missions were very successful, but as is typical with imaging radars, the width of the swath was only between 48 and 64 kilometers (30 and 40 miles), not really giving the sort of wide-angle view that should be attainable from space. So one of the engineering geniuses on our team at JPL thought of a better way to do the imaging, and as we were also running a few research and development experiments on the flights, we tried

This 250-kilometer-wide picture of ice floes entrained in ocean eddies in the Weddell Sea near Antarctica was the first ScanSAR wide-swath radar image ever acquired from space. Taken by the Spaceborne Imaging Radar in 1994, it paved the way for the SRTM.

The SRTM mission astronauts, from left: mission specialist and payload commander Janice Voss, mission specialists Mamoru Mohri (Japanese Space Agency), Janet Kavandi, and Gerhard Thiele (European Space Agency), pilot Dom Gorie, and commander Kevin Kregel.

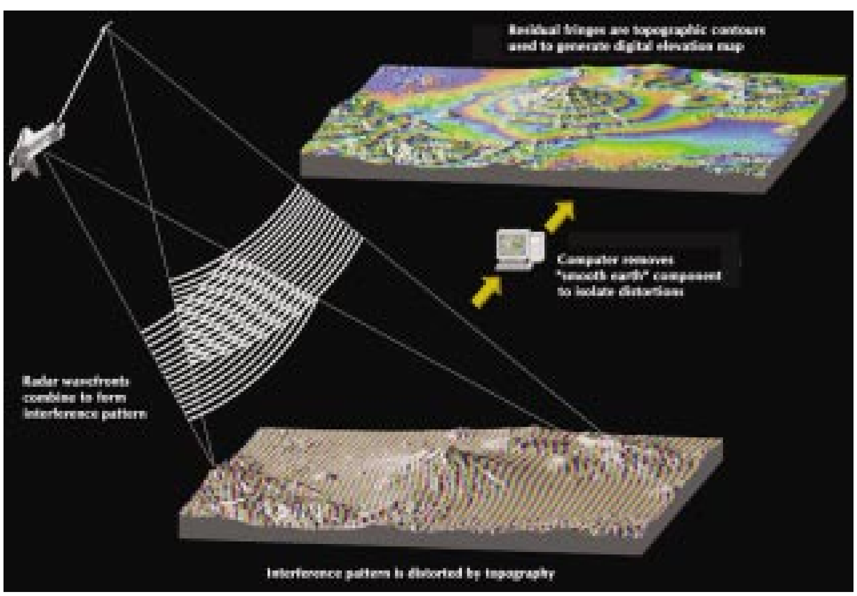

An interferometer can be visualized as two sets of lock-stepped (coherent) radar wavefronts combining to form interference fringes, which overlay the radar image and are distorted by local elevation differences. Isolating and displaying only these distortions produces elevation contours, which can be color coded. Recognize Mount Shasta?

his idea out. Instead of having the beam sweep along in a continuous fashion, we triggered the TR modules to electronically wiggle the beam back and forth, perpendicular to the direction of flight. In the radar world, we call this ScanSAR (Scanning Synthetic Aperture Radar), because the beam scans across the swath the way an electron beam sweeps across the face of a cathode-ray tube to make a television picture. You point the beam at one spot on the ground for about a tenth of a second, then zip it over to another place and point it there for a while, then you zip it over to a third place, then a fourth. By looking at four such subspots, you can sweep out a much wider swath of data. Shown left is the result. It's a historically significant image because it's the first ScanSAR image that we know of ever acquired from space. Instead of a mere 50 kilometers, the image is 250 kilometers across. It's a huge area.

On those 1994 shuttle flights we also experimented with a relatively new technique called radar interferometry, which is not too dissimilar from optical interferometry. Our imaging radar uses "coherent" radar waves, similar to laser light, in which all the electromagnetic waves are in step with one another like marching soldiers. The radar beam can be pictured as a set of concentric wavefronts centered on the antenna on board the shuttle, with a 5.6-centimeter separation between the fronts. If there's a second antenna some distance away with similar wavefronts coming from it, a series of interference fringes is generated where the two sets of arcs overlap, as shown in the diagram on the left. If the surface of the earth were perfectly flat, these fringes would be nice, straight lines. But if there are lumps and bumps on the surface, the fringes get distorted, and by measuring these distortions we can deduce the topography.

With a nontrivial amount of computer processing we can turn these fringes into a digital elevation map, which is like a picture made up of rows of dots, where the brightness of each dot represents a topographical elevation. The brighter the dot, the higher the elevation. If we display this map as a picture on a computer screen, it resembles an X ray, fuzzy and indistinct. But some simple filtering with an image-processing program like Photoshop can turn it into a shaded relief map, and the elevations can also be color coded. Then it starts to look a lot more familiar.

The leftover hardware from the 1994 space shuttle flights was qualified to fly again, but there were no plans to do so and it was sitting in storage down in Carson in an airproof box. As we were casting about for ideas as to what to do with it, one of our more creative engineers at JPL, Ed Caro, had a brainwave: what if we borrowed one of the long, extensible masts that were being built to hold the huge solar panels for the International Space Station? We could attach one end of the mast to our existing antenna structure in the

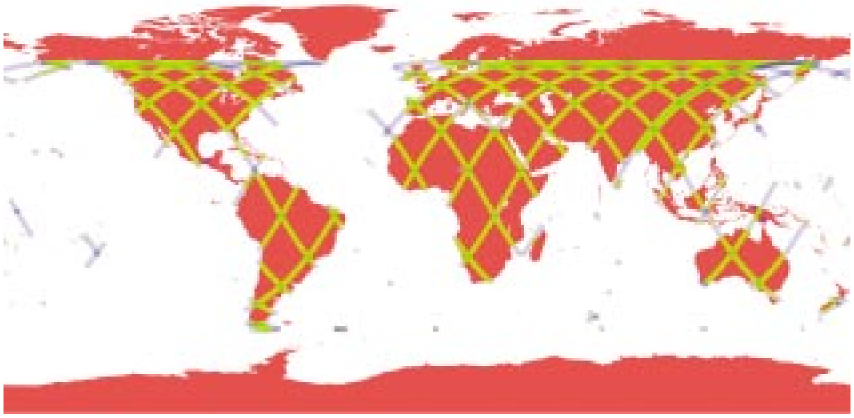

The radar swath coverage as the shuttle orbits the earth, shown for a single day (day 5), left, and for the entire mission, lower left. Land not mapped is red, which the swath coverage turns light green over land and blue over water, then darker and darker green or blue as terrain or water is covered more than once.

payload bay, and put a second radar antenna at the other end. With both antennas pointed at the ground, we would have a proper interferometer. Now, in principle, you don't need both antennas at the same time: when we did this interferometry in '94 we did it by flying back to almost the same spot in space and repeating the scan on a separate day. But if there are any atmospheric changes in those two days (like wind, rain, or dust storms), the surface "decorrelates" and interferometry doesn't work very well. If both antennas are there at the same time, however, it should work perfectly.

But as bold an idea as that was, it wasn't really bold enough because nobody doubted that it would work. Basically everyone said, Well, so what? We know interferometry works, and there's no need to demonstrate it again. Then I had an idea. This was probably the only good idea I have really ever had, and like most good ideas it wasn't inventive or creative—I just stole two other ideas, interferometry using Ed's mast, and ScanSAR. What if we make that secondary antenna at the end of the mast also a phased array, so that the two beams can scan in unison? Both beams, scanning at the same time, would be doing interferometry over a swath we calculated could be 225 kilometers wide. Now that's an important number I want you to remember.

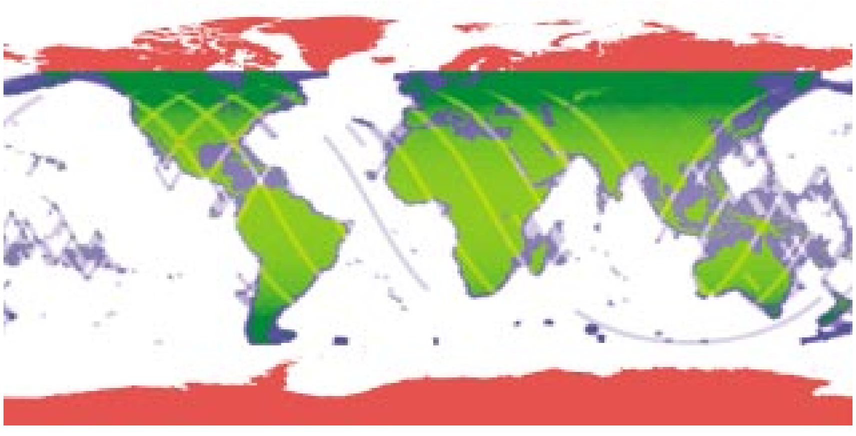

How much of the earth's total landmass could we measure with that? After resurrecting some software I wrote in graduate school, I calculated all the exact repeat orbits around the earth possible for the highest orbit inclination the shuttle could reach (a repeat orbit is one that brings the satellite back over exactly the same spot again after a certain number of revolutions). Although there were a large number of solutions, some fundamental physics limited the number of choices. The lowest the shuttle could go was about 200 kilometers above the earth; below that, there's too much atmosphere and it doesn't stay in orbit very long. Also, our payload would be fairly heavy, so the orbiter wouldn't be able to get up very high into space, ruling out everything above 250 kilometers. Further, the orbiter could only fly for about 10 to 12 days, after which the astronauts would start to run out of fuel, air, food, and so forth (not a good thing), so with a margin for error, the mission could not go beyond 11 days. This left about half-a-dozen solutions, with the obvious choice being an 11-day mission at a 233-kilometer altitude. It turned out this pattern would repeat exactly in 159 orbits. (To get an idea of what this would look like, imagine wrapping twine around a ball, with the twine perfectly equally spaced, 159 times.)

The key figure I needed to know was how far apart the ground tracks—the tracks of the orbiter projected onto the earth—would be as it went round in those 159 orbits. Would it be more or less than the width of the radar swath? If it were more, there would be gaps between the swaths, and the idea wouldn't work. It's a fairly simple calculation: divide 159 into the circumference of the earth and multiply by the trigonometric sine of the angle at which you cross the equator, which is almost exactly 60 degrees.

Two views of the mast: just starting to deploy from the canister, left, and fully deployed, right—the technician's still in the picture, just too far away to be seen. Below: Our home planet as it's typically seen from space—cloud covered.

The back of my envelope said 218 kilometers. I couldn't believe it! It meant that with our 225-kilometer-wide swath there would be no gaps, and we could completely map all the landmasses between +/- 60 degrees latitude in a single shuttle flight using just leftover hardware. This was so unbelievable, I was sure there had to be something wrong with the idea, so we spent a couple of weeks trying to figure out what that could be—and found nothing. In fact, all the usual bugaboos that make your life miserable in planning a shuttle mission just weren't there, and not by our brilliance, but seemingly by random chance.

For example, all the land on the earth except Antarctica was north of the southern extent of those ground tracks, so we would only need to have the radar look in one direction—north—which meant we could always stay in the same shuttle attitude. The attitude is the direction the shuttle is pointing as it flies along the orbit. This saves a lot of fuel, which we knew would be in short supply. When we planned the 1994 flights, a coin toss had determined that the onboard antenna would be tilted slightly toward the orbiter's starboard side, which meant that the attitude the shuttle would have to be in to look north on this flight was tail first. Now the attitude of the shuttle is important, because there's a lot of debris floating around in space that could hit the windows and the leading edges of the wings. There are many attitudes NASA won't let you fly in for very long, like nose first. Tail first was no problem, we were told: we could stay there the whole flight. What luck!

So at the end of 1994 we presented our brilliant idea to the Earth Science Enterprise at NASA in the hope that we'd get funding. The cost of the payload would come to about $100 million, about one-fifth of what we figured it would cost to do the same thing with a free-flying satellite. They were impressed, but suggested we search out another sponsor who could help defray the costs. This was pretty discouraging. After all, who has that kind of money and is interested in maps? Well, we got lucky again: the Defense Mapping Agency (DMA), a branch of the Defense Department that makes maps for the military, was looking for a way to make digital maps at high resolution—one data point for every area about the size of a football field for the whole earth. They'd done about 60 percent of the earth using satellite photos and other sources, but now they were stuck, and the main problem was clouds. At any one time only about 40 percent of the earth is covered by clouds, but they're not statistically distributed. Some places are almost always clear, while others such as northern South America, Indonesia, and islands in the Pacific are cloudy almost all the time. Camera-carrying satellites couldn't photograph through it, but our radar could. So the DMA, now renamed the National Imagery and Mapping Agency (NIMA), hopped right on board. Other participants in the project were the German Aerospace Center (Deutsches Zentrum für Luft- und Raumfahrt), who built part of the system, and the Agenzia Spaziale Italiana. The Shuttle Radar Topography Mission was on its way.

Artist's concept of the SRTM in space, as viewed from a nearby UFO. This painting is courtesy of Ball Aerospace, builders of the outboard antenna, who clearly view the SRTM as their antenna attached to some minor associated hardware—the shuttle!

The heart and soul of the project was the huge mast (opposite page, top). At launch it was folded up inside a canister about 3 meters long, and once in orbit the canister started to turn a helical screw mechanism that pulled the mast open and unfurled it one "bay" at a time. It was an engineering marvel: 60 meters long, with 76 bays in all, made of plastic struts reinforced with carbon fiber, with stainless-steel joints at the edges, and titanium wires held taut by 500 pounds of tension. One of our two radar antennas was mounted on the end.

On February 11, 2000, space shuttle *Endeavour*, with a six-person international crew, carried up the SRTM. They returned to Earth on February 22. The mast–orbiter combination was the largest rigid object that had ever flown in space, even bigger than Mir. In one day of mapping, *Endeavour* made 16 orbits, with the radar covering a 225-kilometer-wide swath each time (as shown on page 27). We imaged almost every landmass twice to get a more accurate result, although a few places were imaged only once because the astronauts were running low on fuel and had to stop a few orbits short. But those areas were in the United States, and had already been mapped accurately by conventional means. And, anyway, even once was enough to meet our accuracy specs. The most nerve-wracking aspect was that for the mapping to be successful, everything had to work correctly over the entire mission—and it did. At the end, we'd imaged 99.96 percent of our target.

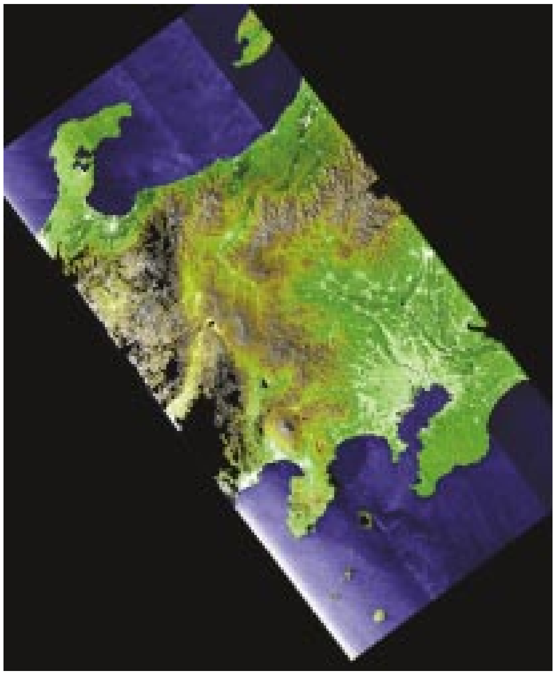

Section of a typical 225-kilometer-wide SRTM data swath, this one covering central Honshu, Japan. The lighter region around the bay to the right is Tokyo; the black regions are data gaps that will be filled in by adjacent swaths.

The mission collected 12 terabytes (12,000,000,000,000 bytes) of data, about the same as the amount of information contained in the U.S. Library of Congress. The ongoing plan is to process every single byte; we have something like the tenth largest supercomputing facility in the world at JPL, and it's processing these data full time. Even with all that computing power, it's going to take until late 2002 to finish. The intention is to process all the swaths, put them together in a mosaic, chop that into little squares of 1 degree longitude by 1 degree latitude, and deliver them to our users, who are scientists, the public, and NIMA.

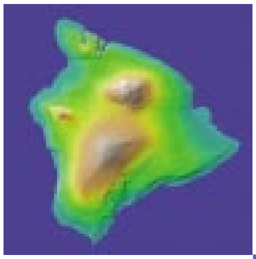

Topographical maps of the Big Island, Hawaii (left) and Sicily (below), color coded for height.

We can do quite a lot of interesting things with the data. For instance, we can create perspective views by sandwiching an optical photograph taken from orbit by Landsat with a digital elevation map of the same area. To give more detail, aerial photos can also be superimposed, and the elevations color coded for height. We can also generate video to give the feeling of flying over the area. And this brings to mind one of the prime civilian, nonscientific uses for these kinds of data in the future, enhanced ground-proximity warning systems for aircraft.

The most common element in plane crashes—so common it has its own acronym—is CFIT, or Controlled Flight into Terrain. This means that even though the plane is working normally and the pilot is in control, for some reason it just gets

Two views of Costa Rica, a country often covered by clouds. On the left, the Caribbean Sea in the foreground is separated from the Pacific by the central mountain range; on the right, the capital, San José, can be seen nestling below the volcano Irazu.

flown into the ground, generally because the pilot couldn't see where the ground was. Well, with the advent of the Global Positioning System (GPS), an airplane can know exactly where it is in three-dimensional space, but the plane also needs to know where the ground is, and that's where the digital topographic map comes in. With such a map a pilot could have a virtual-reality screen in the cockpit showing exactly what it would look like out the window if it weren't for those pesky storm clouds, or just plain darkness. It might not be cheap to install such a system in all existing aircraft, but now that we've produced this data set there's technically no reason it can't be done.

Wireless communication is another important application. I'm sure many of you get annoyed when you're driving around and your cell phone cuts out because the terrain blocks the phone from the antenna. Now, suppose I gave you the task of deciding where to put antennas to get the best coverage of Los Angeles. It would be pretty tough—you'd almost have to do it by trial and error, by going out and doing tests in different places. Well, with the digital elevation map it's easy. Just as an exercise to demonstrate how easy, I sat down and wrote a computer program in about 20 minutes to figure out what the coverage for Los Angeles from Mount Wilson would be. Mount Wilson is one of the highest accessible points overlooking the Los Angeles basin, which is why there's a sizable "antenna farm" up there. But when the illumination pattern from Mount Wilson is overlaid with a Landsat photograph coregistered with a digital elevation map, you can see where the reception shadows are (below). The

Satellite view of Los Angeles without and with the illumination pattern from the antennas on Mount Wilson. The mountains cast "shadows" where television and radio reception are typically not very good. Calculating antenna coverages like this is quite difficult using paper maps, but trivial with SRTM-style digital elevation maps.

Nyiragongo volcano in the Democratic Republic of the Congo erupted on January 17, 2002, sending streams of lava into the city of Goma on the north shore of Lake Kivu. More than 100 people were killed, and more than 12,000 homes destroyed. This picture shows the lava flows (red) by combining a Landsat satellite image, an SRTM elevation model, and image data from the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) on NASA's Terra satellite.

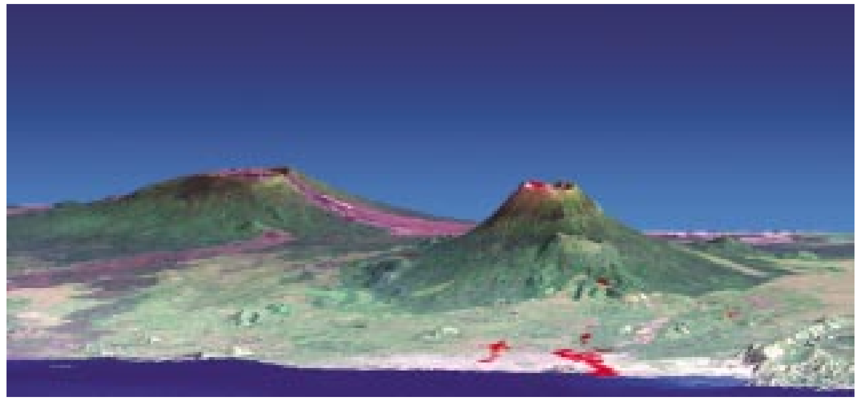

Palos Verdes peninsula casts a big shadow out over the ocean, and the Santa Monica mountains shade many parts of Malibu. I used to live in a place called La Crescenta, only 20 miles away from the transmitters on Mount Wilson, but I got no TV reception at all because that area is in the shadow of Mount Lukens. About 10 years ago I moved over to Pasadena, and now my reception is perfect, just as the map says it should be!

At present, we're still cranking those 12 terabytes of data through the computer and mapping the earth. What are we going to do with this in the future? The hardware that we flew on the SRTM is qualified for a number of additional flights, so it could be flown again, although there's no plan to do so—the 2000 flight was so perfect we can't think of any reason to repeat it. But there's another application of interferometry, called differential interferometry, that is truly amazing. If you know the topography (and that will be true for almost everywhere on the earth now that we have the SRTM data), and you have a radar image taken before and after an earthquake, for example, you can measure the movement of the surface of the earth with almost unbelievable precision. Our vertical resolution in the SRTM is only 10 meters (30 feet), but differential interferometry could detect earth movements of just a few centimeters. Earthquake movements could easily be measured from an orbiting satellite equipped with differential interferometric imaging radar. Some theories predict that the earth does funny things just before an earthquake, like bulging—remember the Palmdale bulge?—so it's possible that differential interferometric remote sensing could actually be used for earthquake prediction. Why not build a satellite that just flies along the San Andreas fault every day or so and checks for earth movements? In my opinion the money saved from predicting just one major earthquake correctly would be more than the cost of such a satellite.

There'll be a lot more interesting results coming out of this mission as we reduce the data. Take a look at our Web page at http://www.jpl.nasa.gov/srtm for updates, more background information, pictures of the world's topography, and flyover animations. □

PICTURE CREDITS:
22, 23 – Doug Cummings/NASA/JPL/NIMA; 24-28 – NASA/JPL; 28 – Ball Aerospace; 29-31 – NASA/JPL/NIMA

The Caltech campus in Pasadena with the San Gabriel mountains behind.

A specialist in radar remote sensing of the earth and planets, Michael Kobrick has worked at the Jet Propulsion Laboratory for 28 years, ever since earning his PhD in planetary and space physics from UCLA. He also has a BS in physics from the Rensselaer Polytechnic Institute, and an MS in astronomy from the University of Illinois. He was principal investigator in the early spaceborne radar experiments of the Apollo program, science manager for the 1990 Magellan mission to map Venus with radar, and has also spent several thousand exciting flight hours experimenting with airborne imaging radar systems. As project scientist for the Shuttle Radar Topography Mission—his brainchild—he is now in charge of the reduction of the data set and its release to end users. This article is adapted from a Watson lecture given on February 21, 2001.