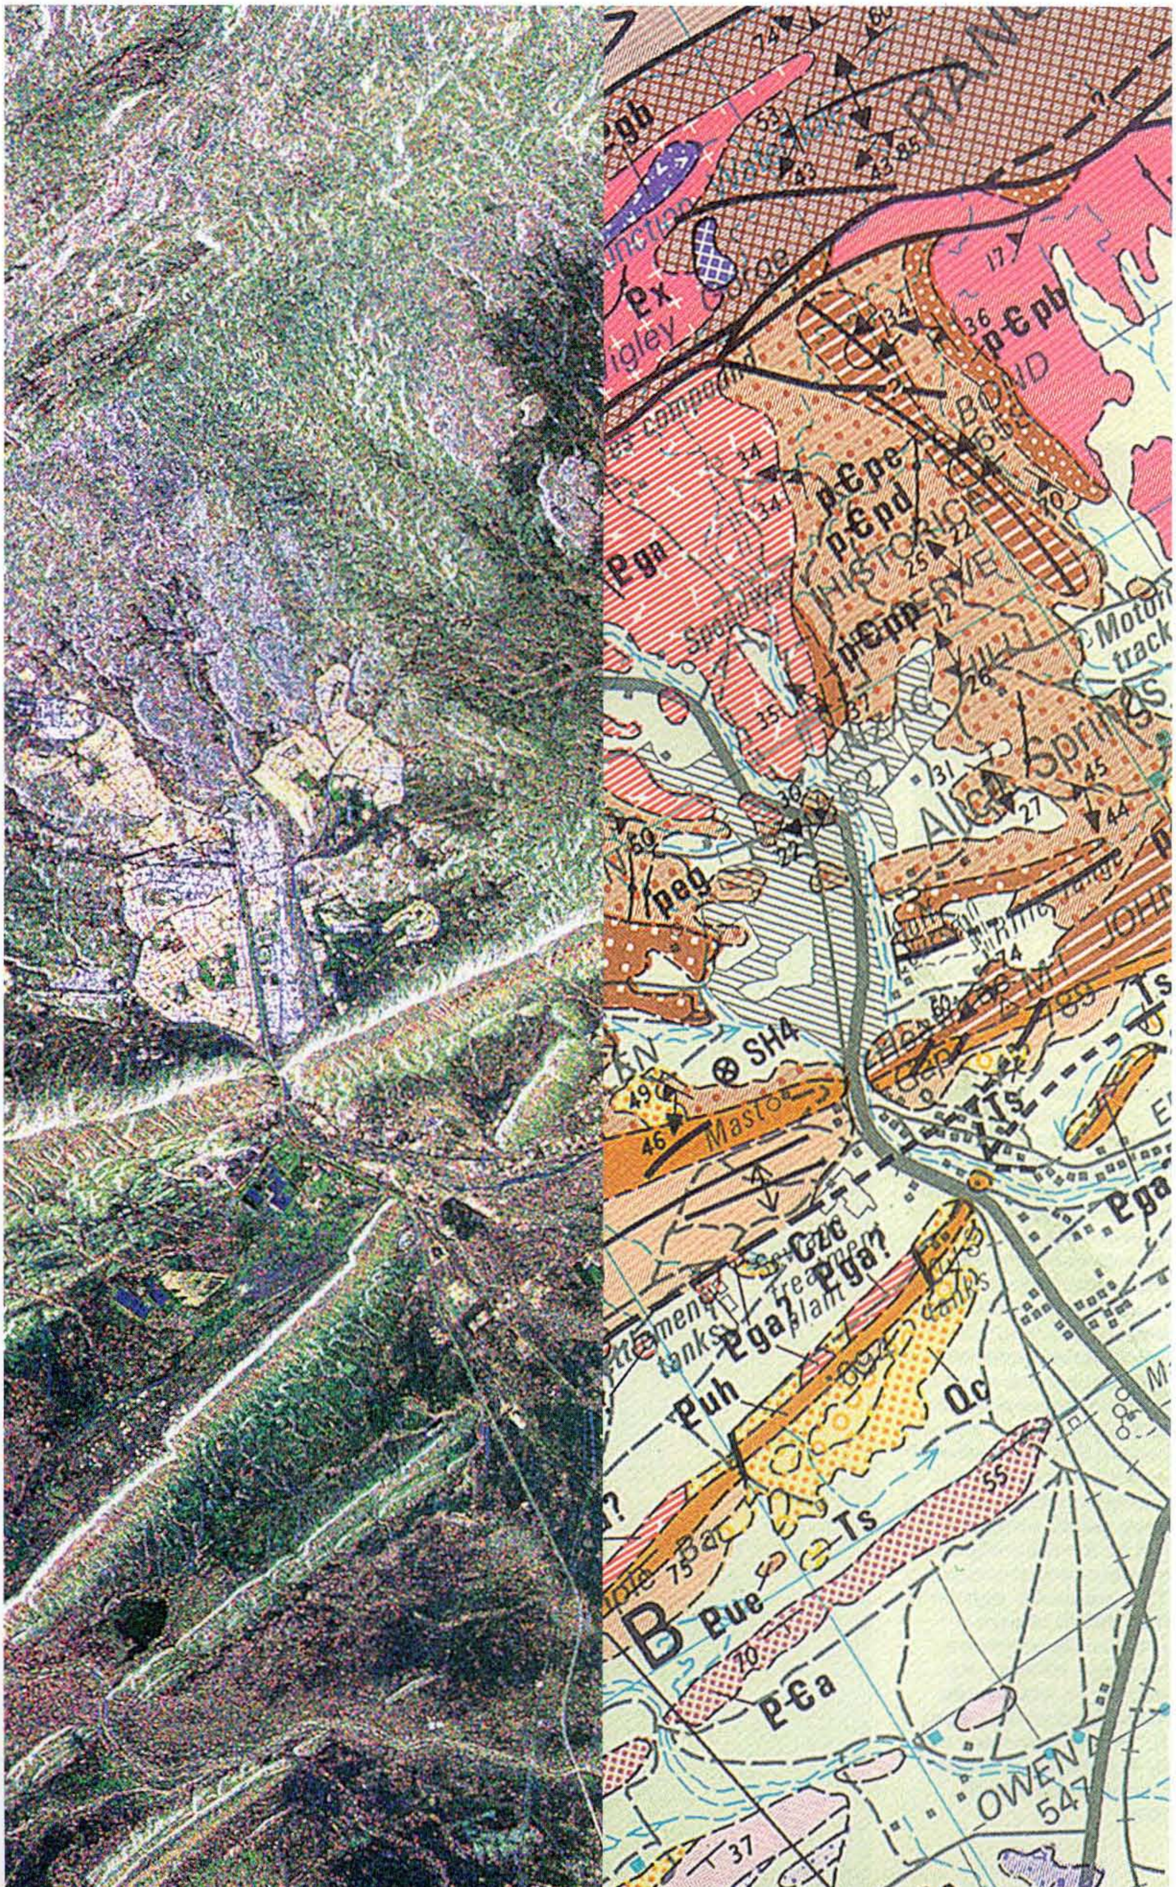

SH4
Ts
Qc
B
OWEN

Dreaming of Hypermaps

I have dreamed for years of an electronic alternative: owning a map whose center, scale, and content are determined by me, not by the maker of the map.

by Roy D. Williams

The vicinity of Alice Springs, Northern Territory, Australia, as seen by the space shuttle's Synthetic Aperture Radar, or SAR, (left) and the Australian Bureau of Mineral Resources, Geology, and Geophysics (right). The colors in the radar image are keyed to the region's mineralogy. The map helps us make sense of what we see in the radar view by marking faults (thick black lines), labeling rock types, and naming features, both natural and man-made. But maps don't always get things quite right, as you can see in the curve of the railroad in the bottom right corner.

I'm writing this on a laptop on an airplane, watching the Oklahoma landscape roll smoothly by. There's a lot to see, in my opinion, from airplane windows—geological faults and rock unconformities; the way water erodes rock, and how roads, farms, and cities form themselves around the resulting watercourses. Sometimes, with raking sunlight and a dusting of snow, you can see ancient villages or medieval agriculture (though not in Oklahoma). I love the nonstop from Los Angeles to London, looking over endless, endless Arctic Canada and the mountains, like broken bottles, that cut through the Greenland ice sheet.

With a little imagination, such a landscape springs forth from a paper map, especially a finely detailed, large-scale one. For this reason I have always enjoyed armchair traveling with the aid of a map; this is especially fun before a trip, and occasionally more fun than the actual trip. Thus, the favorite maps in my own collection are those that represent lands far from my own experience. I like to say the names to myself, to wonder what happens when that tiny road simply ends in the middle of a jungle, to speculate about who uses the quay that gives access to a tiny, Atlantic-battered island. There are others who share this passion: in the meat-market district of Manhattan there's a café whose walls are covered in old street plans of cities from around the world, stuck up with thumb tacks. (The time I was there I walked from map to map, peering at them over the heads of other customers, who had to lean aside to get out of my way.) A map can add color to a book or newspaper story by showing where a battle was fought, or a train derailed; where the world's rice is grown, or the territory of a vanished empire stretched.

Old maps can be a lot of fun too. A few years ago, I was living in Oxford, quite close to Holywell Church, which is on Holywell Street. After a few exploratory sessions, I could find no evidence of a well, holy or otherwise. This seemed like a challenge, since it must have existed at some time, so I decided to try to nail it down. I wheedled my way into the Map Room of Oxford University's Geography Department and dug up some town-planning maps from 1862. These maps were at the scale of 1:1250—in other words, an inch on the map equaled roughly a hundred feet on the ground. At this scale you could see everything! Next to the church, at a distance equivalent to perhaps 20 feet, the map was annotated "Ancient Well" in gothic characters. I rushed back to the church to check it out. The ground showed no evidence of anything ancient, just a compost heap. But the churchyard wall contained some extra angles, implying that the builders had been making space for something—presumably the well. It was quite satisfying to feel that a tiny scrap of very unobvious history had been unveiled.

Maps as objects are fun to collect and pore over, but maps as information providers sometimes leave something to be desired. The area one really wants to see so often seems to be near a corner of the map, or the journey one imagines continues off the edge; besides, paper maps are really very awkward to file and store.

I have dreamed for years of an electronic

The map showing the entire Earth would be stored within the display itself and instantly accessible, but to plan a hike in the Macdonnell Ranges deep in the Australian outback, the hypermap would need to use the Internet to get the data from a machine in, say, Alice Springs.

alternative: owning a map whose center, scale, and content are determined by me, not by the maker of the map. I envision a high-resolution display that can show custom maps, not just the ones that I happen to have bought for earlier expeditions, or the ones that someone else has decided would be useful enough to have printed. Such a display is generally called a Geographic Information System, or GIS, and is already in use in specialized forms in urban planning, epidemiology, seismology, and many other fields. I would like to develop something rather more general, however, that I call a hypermap. A hypermap would integrate data from many different sources, making the information available to professionals and amateurs alike.

My personal hypermap would be wall-mounted and backlit for armchair use; I'd also like a handheld version, the size of a legal pad, to look at while lying in bed. I would enter the map by selecting a point and having the scale double. After ten iterations, my point of view will have been reduced from a flyover of the globe to a ramble in the countryside. Naturally, things would slow down as the scale increases; the map showing the entire Earth would be stored within the display itself and instantly accessible, but to plan a hike in the Macdonnell Ranges deep in the Australian outback, the hypermap would need to use the Internet to get the data from a machine in, say, Alice Springs. As I looked at the map, its software wouldn't simply sit waiting for my next command, but would instead prefetch data on the surrounding areas, filling its memory in a spiral pattern, on the assumption that I would soon want to look at the land just beyond the frame of my display. Also, perhaps, software agents would be scouring the Net for other maps covering the same area, but containing different information.

For each map I called up, especially if I had set the prefetching at a voracious level, there would be a disbursal of micro-cyber-cash from my credit-card account to various data purveyors, some in Australia. Inevitably, my Net provider would sell this information to a database company specializing in travel-related matters. Over the next few days, junk mail (both electronic and paper) would arrive, advertising the joys of an adventure vacation in Australia. Through automated database correlation, yet another company would have narrowcast to me with piercing accuracy, simultaneously helpful and eerie.

The hypermap would be ideally suited to do what is known in the data trade as fusion: taking different data sets and combining them to produce something new and, one would hope, more informative. Fusion is the essence of mapmaking. A cartographer creates a paper map from survey data and aerial photographs. Information is added from other maps—the cartographer uses a pen or a mouse to draw in roads and county boundaries, and writes or types names such as "Wyre Piddle" next to a beautiful English village that sits next to a small stream. Information is cross-checked (the cartographer compares the scattered height data from the survey with the contour lines drawn from stereoscopic pairs of aerial photographs) and updated (a reservoir's shoreline is redrawn because the dam is now higher than it was). The hypermap should do

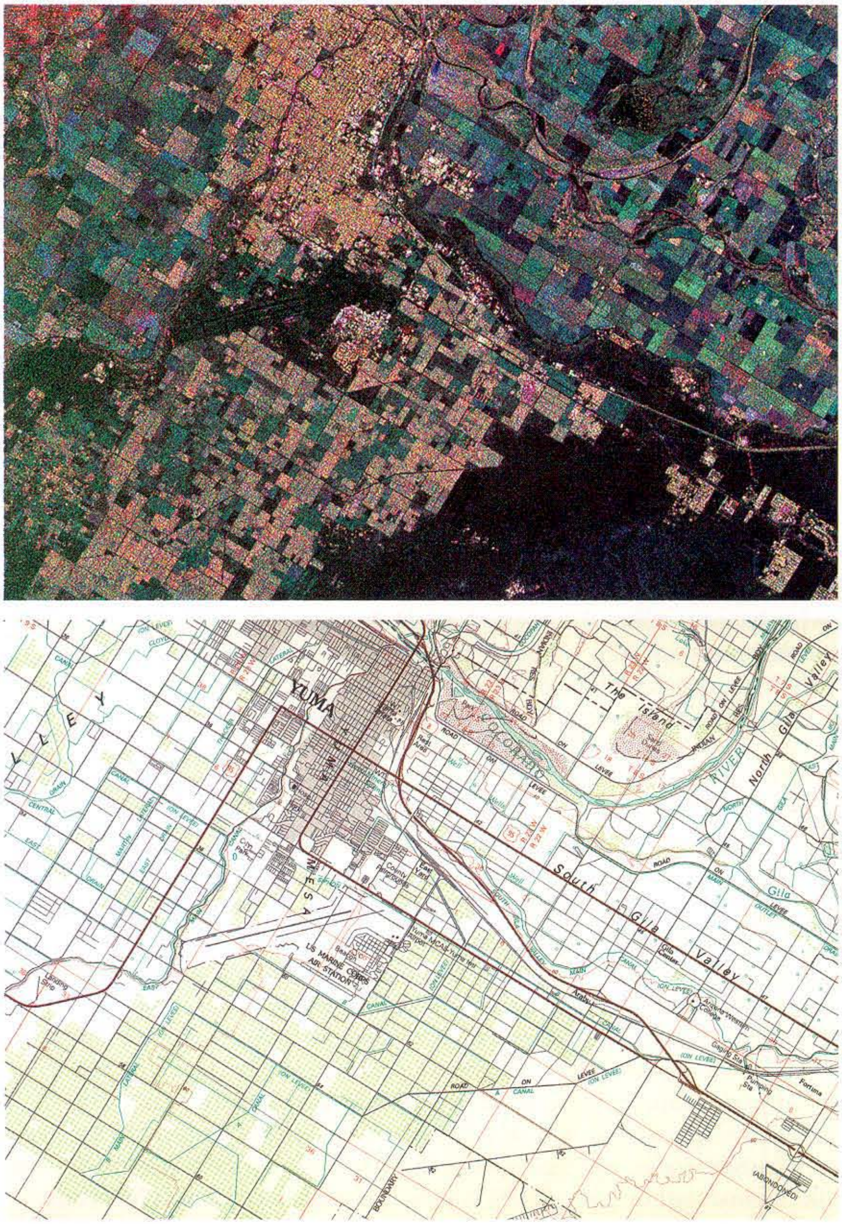

Yuma, Arizona, as seen by SAR (top) and on a U.S. Geological Survey map (bottom). Note how the course of the Colorado River has changed since the map was drawn.

all this as well, while combining a high-resolution photograph's intense feeling of reality with a cartographer's knowledge of the geographical significance of the landscape. Some maps already approach this—on my office wall hangs an image of Los Angeles as seen from space. The green, forested mountains, the gray-brown urban area, the brilliant white dry lake beds in the red-brown Mojave desert, the San Andreas fault, and a lot of other physical features are shown in absorbing detail. But major roads and place names are also marked, allowing us to get our bearings. We can relate what we see anew to what we already know, and thereby create knowledge from data.

The personal GIS already exists, but without the use of remote servers and the possibility of junk mail: there's a CD-ROM called *Street Atlas*, which gets the most use of the dozen or so CD-ROMs that my wife and I have bought since getting a CD-capable machine. (*Street Atlas* even has the scale-doubling feature, though I would like to point out that I thought of it before I got the CD!) In one sense, these little plastic disks hold a great deal of data: one of today's CDs can hold enough novels to read one a week for ten years, and soon it will be possible to pack a lifetime's supply into an alluring plastic rainbow. But for storing images or geographic data the CD-ROM seems much smaller—more like a single drawer than a library—simply because a page of image takes up a lot more bits of data than a page of text. Consequently, *Street Atlas*'s maps do indeed cover the whole United States, but the streets are really just named geometric lines, all the same width. And there's no indication of any texture to the landscape: no forests, ruins, salt marshes, glaciers, battlegrounds, or wind farms.

There's another, more compelling reason why a physical data repository—even a roomful of CDs!—cannot be adequate for the full efflorescence of the hypermap. The problem is that a single physical object is tangibly limited; no matter how much data it holds, it is obviously finite in extent. In contrast, a networked hypermap can explore a potential infinity of data and can go anywhere in the world to get it. True, the amount of digital geographic data in the world is also finite, but the difference is that it is growing. The hypermap concept carries with it the idea that as data is newly minted from cartographers and orbiting satellites, then all the world's hypermaps immediately gain the new depth. Thanks to airplanes and satellites, the quantity of geographic data is increasing even faster than the violent increase in computer speed that's changing society so quickly. When up-to-date maps of essentially infinite detail are combined with

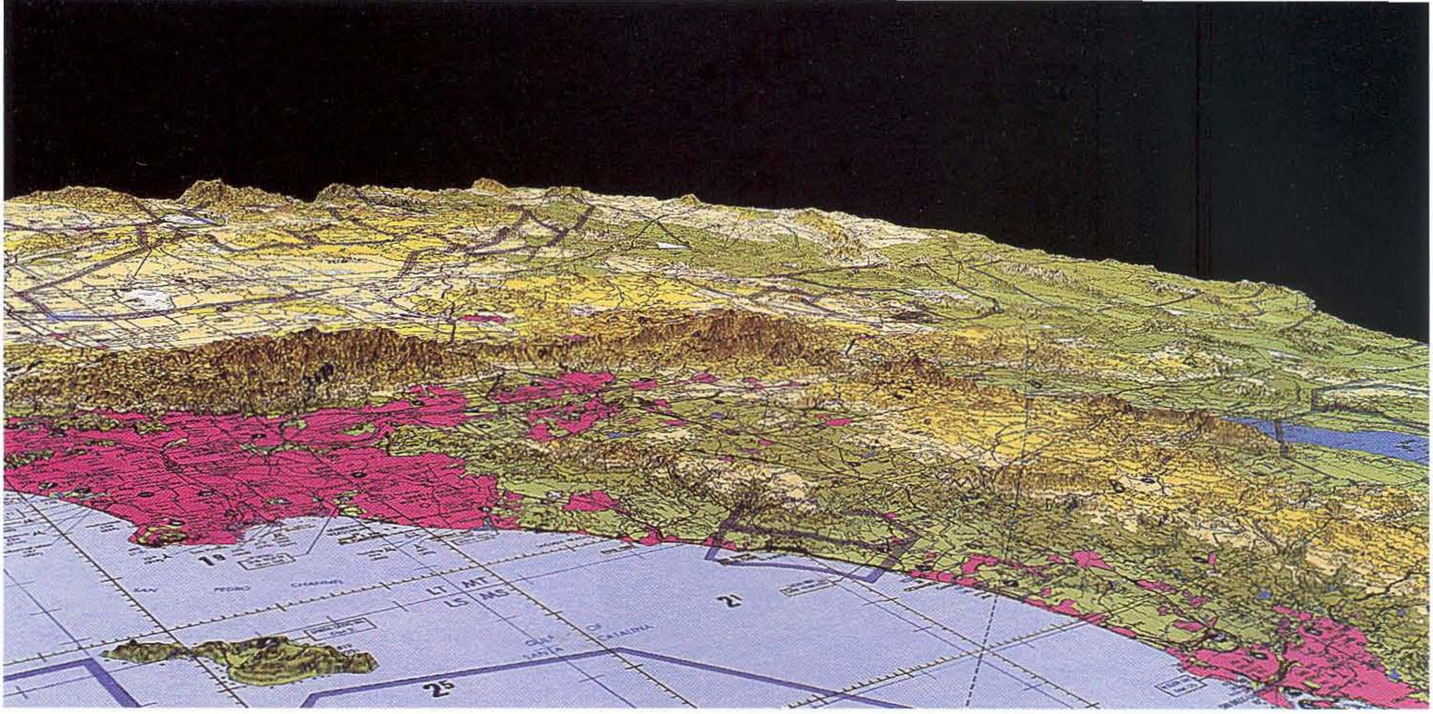

The view out the right-hand cockpit window of a virtual space shuttle descending toward Los Angeles International Airport. We're about 62 miles high and 138 miles southwest of the airport. Catalina Island is in the foreground, LAX is at the very left, and San Diego can be seen at the far right. This relief map combines an air-traffic control chart issued by the Federal Aviation Administration with digital elevation data from the USGS. The vertical relief has been exaggerated by a factor of 3.5.

cellular phones and high-accuracy geographic positioning systems, can the concept of an expedition into the wilderness survive?

So we zoom in closer and closer and the map's scale gets larger and larger—what happens when we let our imaginations play with this system? Suppose the map becomes three-dimensional, like those plastic maps that you can run your fingers over to feel the mountains. Perhaps the viewer will use a virtual-reality helmet to land an airliner at LAX, or loop-the-loop in Arches National Monument, or fly through the glacial canyons of Antarctica and over Everest.

An existing system of some interest here is the "Virtual Los Angeles" project—created by William Jepson and the rest of the Urban Simulation Team from UCLA's Department of Architecture—a virtual model of large tracts of the city, complete with trees and graffiti. Graduate students shoot video footage of the streetscape, which is fused with satellite and mapping data into a seamless, realistic, textured urban landscape. The user can then "walk" around within the model by means of a mouse, a joystick, or a virtual-reality helmet. The city is intensely real, yet eerily deserted—the streets and sidewalks are practically empty because vehicles and pedestrians still take too much computing power at this point to be included profligately.

The Internet is providing this kind of three-dimensional experience today, meaning that all you need is a fancy computer, rather than having to know the right people in addition to having a fancy computer. At the moment, one can tour, among other places, an Italian castle, Jerusalem's city hall, and the city of Berlin. The new protocol that enables this to occur is called Virtual Reality Modeling Language, or VRML; when you download a VRML document, your computer opens a 3-D "browser" that reads the file and allows you to explore the space encoded within it—turning, twisting, accelerating, panning, and zooming at your whim. You might even meet representations of other people, known in the trade as avatars. (If I am ever virtually represented in a 3-D space, I would like my avatar to be the boot token from the Monopoly game.) And VRML isn't just for architectural touring or city planning; it's a method of transmitting any kind of three-dimensional data for interactive exploration. Other Net sites allow the user to wander about within molecular structures, the fruit fly's nervous system, and galaxies, to name a few.

So how can we get closer to the hypermap? Where will all the data come from, and who will pay for it? If the science budget can survive congressional attack, the Earth Observing System (EOS) will be operational in the next few years. EOS is part of NASA's "Mission to Planet Earth"; even if it gets cut, as seems likely, I hope it will just be delayed for a year or two until the next Congress reinstates it. One rationale for EOS is to provide the data needed to predict the effects of global warming in specific, quantitative, local detail for long-range planning purposes. The current climate models, even those that run on the fastest supercomputers, have as inputs scattered observations supplemented by sharp guesses, and give vague, global predictions as outputs. When the EOS data start coming, there

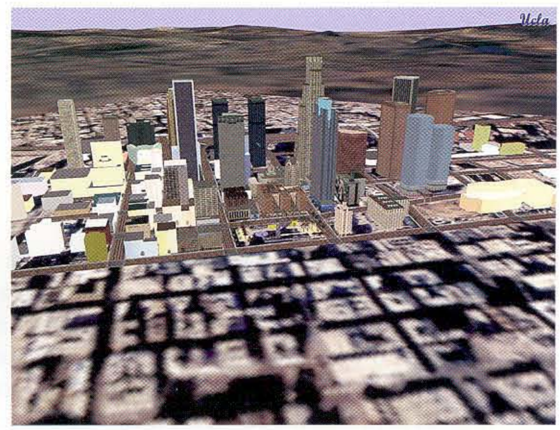

A flight into Jepson's Virtual L.A. model, starting over Catalina Island and then hooking around to land at the intersection of 5th and Hill from the east. From here, one could stroll around downtown. In this version of the model, every building and streetscape in a one-square-mile area is modeled in three-dimensional detail; the surrounding city is roughed in from LANDSAT photos. Detailed models of several other parts of the city, including the Pico-Union district, Playa Vista, and part of South-Central L.A., have also been created, primarily for urban-planning studies. Jepson also models streetscapes that don't exist any more—another project shows the Forum in Rome; by pushing a time control back and forth, one can watch the landscape evolve over the centuries.

will be much more sharply detailed predictions based on a much firmer footing. EOS will bring in satellite-based remote-sensing data about Earth's land, oceans, and atmosphere at a vast number of gigabytes per day. Supercomputer centers and data-handling warehouses are already deciding how to process and store the data: silos of tapes and disks will be needed, and the task of keeping it organized, catalogued, and accessible will be Herculean.

And the data aren't just digital photographs, but the outputs of other sensors that have nothing to do with visible light and work instead in infrared or microwave frequencies. There's a strong analogy here to astronomy, which was confined to optical observations until the arrival of radio telescopes. Today, a burgeoning family of telescopes observes the entire electromagnetic spectrum, neutrinos, and soon even gravity waves. The invisible emissions captured by these instruments have provided a new view of the universe, revealing it to be a violent and capricious place, in sharp contrast to the quintessentially perfect "music of the spheres" of the medieval imagination. In the same way, the wide availability of high-resolution geographic data will change our view of Earth, making it at the same time more familiar, more mundane, more complex, and more precious.

One source of such high-volume data is Synthetic Aperture Radar (SAR). Parts of Earth that had previously been difficult to see with visible light because of almost continuous cloud cover, can now be seen clearly by radar. SAR can see through clouds, vegetation, and sometimes even

SAR can see deep enough into sandy desert to discover an ancient ghost city on the Silk Road, and can espy eco-friendly farming taking place beneath the canopy of the Amazon rain forest. SAR can measure the moisture content of Kansas cornfields, and differentiate spruce from birch in the Russian taiga.

a few meters of sand. In the Andes, volcanoes have been discovered that were previously unknown, due to their inaccessibility at ground level and to being shrouded in clouds and fog when looked at from above. SAR can see deep enough into sandy desert to discover an ancient ghost city on the Silk Road, and can espy eco-friendly farming taking place beneath the canopy of the Amazon rain forest. SAR can measure the moisture content of Kansas cornfields, and differentiate spruce from birch in the Russian taiga. SAR can trace the movement of Chilean glaciers, document the destruction of African gorilla habitat, probe the geology of Hawaiian volcanoes, determine the vintage of Antarctic sea-ice, and monitor the recovery of Yellowstone from forest fires.

In order to see so much so clearly, there is of course a price to pay—the raw data from the satellite are not directly visible, but need to be processed by a supercomputer before becoming intelligible. Such a project has been under way at Caltech and JPL for two years now, and it's a massive endeavor involving many people. We feed Intel and Cray parallel supercomputers with tapes of raw data and receive multichannel color images in exchange. Every pixel in a color image conventionally represents three data channels, encoded in the colors red, green, and blue, which correspond to the three kinds of receptors in the human eye. But SAR takes data at eight or more channels, leaving a choice of how to throttle the flow down to only three. Such filtering choices can be made to emphasize different aspects of the terrain; for example, to identify types of trees or the composition of volcanic lava, or to gauge the quality of potential ski slopes.

For the armchair explorer, part of the exhilaration of this remote-sensing data is that it has not been processed and digested by a human, only by a computer. There may be the ruins of a hidden city, barely visible without contrast enhancement, or a lake forming where none was known before. "Traveling" by means of SAR data carries the possibility of discovery, much like that offered the patient comet-watchers, roving the sky with binoculars. By comparison, making the trip by paper map will feel like looking at a star atlas instead of looking at the sky. SAR data are not simple pictures to be examined, but can be reprocessed in many ways—just as statistical data can be massaged and processed to highlight, emphasize, and maybe even cheat. When we combine SAR data with conventional maps, we can see correlations and associations that were previously hidden, thereby creating knowledge, and—who knows?—perhaps a scrap of what we all crave: insight.

With my colleagues Thanh Phung and David Payne of Intel Corporation and Ellen O'Leary of JPL, I am developing a pilot version of a World Wide Web–based hypermap, called SARA, for Synthetic Aperture Radar Atlas. SARA actually lives on a supercomputer here at Caltech, but you can get to SARA's Web site via an ordinary Web browser, such as Netscape. SARA welcomes you with a map of the world, on which you click in the general area you wish to visit. This in turn brings up a closer view of that region, in which the available SAR data sets are highlighted in red. Clicking on a red zone brings up a "thumbnail" black-and-white SAR image. These images are compressed eightfold from the actual SAR data, meaning that the smallest details one can see in the thumbnail image are eight times bigger than the smallest details one can see in the actual data; similarly, the color channels of SAR data are replaced by one channel rendered in shades of gray. Thus the volume of data to be transferred to your computer is a mere 1/512th of the actual SAR data set—a necessary concession to the speed of the average modem. If you're directly connected to the high-speed part of the Internet, you can then call up the real SAR images, set the three color channels to show you what you want to see, and zoom in to bring up the spatial data that got compressed out of the thumbnail version. This is currently unrealistic for home or high-school use, but soon, we net-mongers hope, higher-speed networking and even-faster cheap computing will make its way to the domestic hearth. SARA was demonstrated

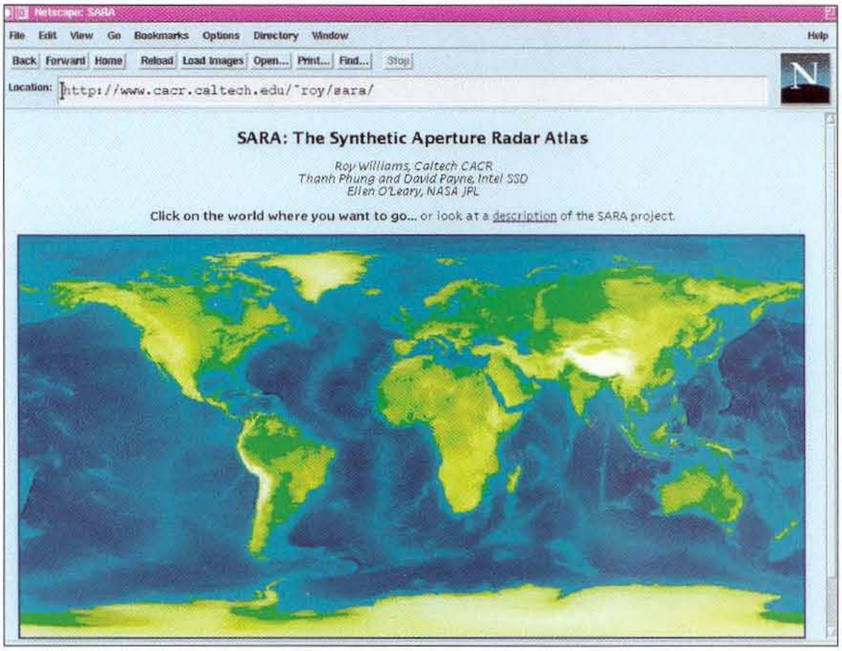

Click anywhere in this map of the world on SARA's home page, and you get a map of SAR data sets available in that region.

The taxpayers have already paid for the raw material, but not for the mining and processing that is necessary to understand it.

at the Supercomputing '95 conference in San Diego last December, running on an Intel Paragon machine, and, I think, was received with some interest—we had people standing in line to see our show.

The technical path to the hypermap is fairly well marked at this point. There are enormous challenges awaiting in database management, and in designing means of exploiting and presenting the data that will be useful to the expert and novice alike, but the routes to solving these problems look reasonably clear. More difficult are the political issues inherent in making available remote-sensing data owned by the government. Part of the difficulty springs from copyright issues; the laws governing the ownership of electronic documents are being written and rewritten even as this article is. But the most daunting issue is cost: the taxpayers have already paid for the raw material, but not for the mining and processing that is necessary to understand it. Geographic data produced at the public's expense has traditionally been made available for the cost of the duplication. In the old days, this meant that you wrote a check to cover the cost of the clerk photocopying the information and mailing it to you, or in the case of a U.S. Geological Survey map, the cost of printing the map. But now, the cost of duplication includes the cost of the electronic pipeline that brings the data to you—the networks, the disk farms and tape robots, the software to run them, and the people to keep it all going. They all cost money, and who will be paying how much for this service is probably the thorniest question of all. Significantly, the U.S. government has already decided that EOS data will be available not only to the priesthood of scientific digerati, but also to colleges, high schools, and individuals.

I hope that these issues can be resolved, and the hypermap brought to fruition, because it will deliver information that is not only useful but inspiring, enabling individuals to see the world through their own uniquely personal maps. In earlier times the translation of the Latin Bible into the languages of the common folk allowed fresh views on Christianity; now hypermaps will allow people to choose how raw data are to be processed and delivered to them, thereby minimizing the distorting lenses of those who would "interpret" the data for us. □

Roy Williams, a senior staff scientist at Caltech's Center for Advanced Computing Research, earned his bachelor's degree in mathematics at Trinity College, Cambridge, in 1979, and his PhD in physics at Caltech in 1983. After postdoctoral stints in England, he returned to Caltech in 1986. His non-map-related research interests include fluid-flow algorithms, differential equations, parallel software, and high-speed networking.

This is his second article for E&S.