

Harry Gray and Mr. Sun enjoy each others' company.



Solar Fuel

by Harry B. Gray

In the next century, burning hydrocarbons—oil, gas, and coal—is going to be a no-no. And it won't be for the reason you think, which is greenhouse warming—it won't be because Cleveland will bake, and Pasadena will drown. We have to stop burning hydrocarbons as soon as possible because they're wonderful raw materials. We desperately need them to make dyes and drugs and T-shirts and chairs and automobiles—it's crazy to burn them. So sometime in the next century you'll see a massive conversion of this fossil-fuel-burning world of ours into a world that burns clean fuel. There are really only two clean fuels: hydrogen (H_2) and electrons. Hydrogen is clean, because when you burn it you get water back. (And, of course, the water can be split again to make more H_2 .) Electrons are clean only if they're generated cleanly, and the best way to do that in the long term is by nuclear fusion. I'll be overjoyed the day that fusion power comes to pass, but I'm a chemist, so this article is about chemists and other scientists who have tried to convert sunlight and water into oxygen (O_2) and hydrogen.

First, I'm going to give you a short course in solar energy. (In fact, this article is a bunch of short courses. You can either buy thick textbooks and get hernias from carrying them around, or you can read this.) There are three fundamentals of solar-energy conversion. The first is light capture—absorbing the sunlight, basically. The second is electron transfer—pushing a sunlight-excited electron off its home atom in order to harvest it. The third is catalysis—the efficient making and breaking of chemical bonds using the harvested electrons and the oxidized atoms they've left behind. Each fundamental builds on the preceding ones, so if we can do all three, we're most efficient. If we only capture light, we turn sunlight to heat, which can then create steam to run turbines to make electricity. We've known how to do that for a long time, but it's very inefficient. A few years ago, we took the next step, and fig-

We have to stop burning hydrocarbons as soon as possible because they're wonderful raw materials. We desperately need them to make dyes and drugs and T-shirts and chairs and automobiles—it's crazy to burn them.

ured out how to do both light capture and electron transfer. This requires light-sensitive semiconductors—photovoltaics such as silicon that convert sunlight into electricity. We now have durable silicon devices with reasonable efficiencies—say, 10 to 15 percent of sunlight converted to electricity. In fact, we could convert Los Angeles into a solar city right now, and we probably should. But if we make electrons, we've got to use them on the spot or lose them. We can store them on a small scale in batteries, of course, but my point is more general: we've got to be able to store the energy on a larger scale for use later. So the goal is to do all three steps, so that we can store the electrical energy in chemical bonds—the hydrogen-hydrogen bond, in this case.

Now for the short course in photosynthesis: photosynthesis is extremely simple. A good book on photosynthesis runs about 1,600 pages. You can't read 1,600 pages without falling asleep, so you may as well take my word that photosynthesis is easy to understand. You start with sunlight, CO_2 , and water, run them through a green leaf, and you get O_2 and carbohydrates. God took the three-step solar-energy scheme, and optimized each component. Chlorophyll is pretty good at capturing the sunlight that reaches the Earth's surface. The electron acceptors—the organic molecules that siphon off the electrons after the light's been captured—are nearly 100 percent efficient. And God made something I'm still marveling at—a beautiful manganese cluster that catalyzes the evolution of oxygen. Photosynthesis is 6 percent efficient overall in converting sunlight to stored chemical energy. You may not think that's much, but it's been enough to run this planet for a very long time.


We know that we can turn a leaf that makes carbohydrates and O_2 into a water-splitting system that makes H_2 and O_2 , because Eli Greenbaum at the Oak Ridge National Laboratory in Oak Ridge, Tennessee, has done it. (Greenbaum—what a

wonderful name for a guy who's doing artificial photosynthesis!) Eli's no fool—while the rest of us were trying to work up these incredible catalysts from scratch, he figured he could take God's invention, and just add one component—a catalyst that couples protons and electrons to make hydrogen. That catalyst is platinum, so Eli got some leaves, took the chloroplasts—the photosynthetic organs—out of them, filled a little Baggie full of chloroplasts, put in a platinum solution, dried off the chloroplasts, added some fresh water, shone sunlight on them, and boom! hydrogen and oxygen. It's beautiful. So we know this can be done, because Eli's done it. God plus Eli—God did all the hard parts, and Eli added a platinizing solution to a Baggie. And guess what? Eli picked up a percentage point of efficiency! (How'd he do that? There must be an error somewhere in the paper.) Eli's water splitting is 7 percent efficient, which corresponds to a semiconductor making electrons at 10 percent or higher efficiency because the hydrogen is stored chemical energy. This is incredible. What's the catch? Why don't we all quit and go home?

The reason we can't quit—and if you don't remember anything else, I want you to remember this—is that efficiency by itself is not enough. If you read about a process that's 7 percent efficient, you should then ask the following question: how long does it last? Will it keep going? The answer here is, no! God made chloroplasts full of wimpy organic molecules that break apart. After being exposed to light for several hours, Eli's platinized leaf poops out and it's history. God doesn't mind, because if any part of the photosynthetic system breaks down, the leaf just makes more. (The leaf is a wonderful synthetic chemist!) But in artificial photosynthesis, we can't do that. We've got to make something that lasts forever. Not only does it have to be efficient, it's got to be durable. Rugged. No wimpy organic molecules.

I'm proud to be an inorganic chemist. Inorganic

Inorganic chemists are the Marines of chemistry. We knew we could make a simple, inorganic, artificial photosynthetic system that would last forever.



Above: How to turn sunlight into hydrogen without wimpy organic molecules. At the top of the cycle, ultraviolet light kicks a cerium atom (Ce^{+3}) into an excited state ($*Ce^{+3}$), which promptly sheds the extra energy and an electron by jolting a hydrogen atom off a water molecule, becoming Ce^{+4} . A ruthenium-oxide catalyst scrounges replacement electrons from other water molecules, liberating oxygen. (The complete set of reactions is too complex to show.)

chemists are the Marines of chemistry. We knew we could make a simple, inorganic, artificial photosynthetic system that would last forever. We didn't need CO_2 inputs, because we didn't want to make carbohydrates—they're too complicated. All we wanted to do was run water in, and run hydrogen and oxygen out. We made our system out of cerium, which is a tough, macho metal. We hit Ce^{+3} with deep ultraviolet light, and it split water to hydrogen. This left us with Ce^{+4} , which with a ruthenium-oxide catalyst converted water to oxygen and gave Ce^{+3} back. Our system was all rugged metals, with no organic components—it would never run down. Unfortunately, it was not very efficient—less than 10^{-10} percent efficient in practice! The problem is that we don't get deep UV light down here on the Earth's surface, because the ozone layer in the stratosphere acts as a filter. Out beyond the ozone layer, this system will be very useful for space travelers and space colonists, but here on Earth, it's back to the drawing board. We hit the beach, all right—we made the system so rugged that if it covered the entire surface of the Pacific Ocean, and

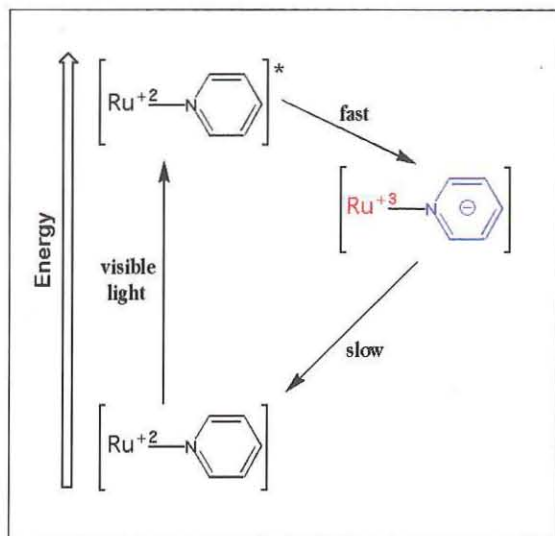
we irradiated it for a billion years, we'd only make enough hydrogen to drive your car a couple of blocks! So you've got to have both durability and efficiency.

We know, in principle, how to do this. We have to build a molecule that will absorb light efficiently (step one of the grand photosynthetic plan). The absorbed energy pushes the molecule into an activated state, denoted with a star, in which all the light energy has been imparted to one electron, kicking it into a more energetic orbit. While the electron is running around frenetically, it's more loosely attached to its home atom, which gives us a window of opportunity to do step two of the grand plan—electron transfer. A properly designed molecule will encourage the excited electron to forsake its home atom and visit some other part of the molecule, creating a positive charge—a "hole"—on the jilted atom, and a negative charge wherever the electron winds up. This negative charge, this electron, can then be diverted to a catalytic center to make hydrogen, and the hole can find another catalyst to make oxygen—step three.

But time is working against us. If we don't separate the active state into positive and negative charges in a few picoseconds (trillionths of a second), the electron sheds its excess energy and returns to the ground state. And once we separate the charges, we have to hold them apart for at least a thousandth of a second—a millisecond—or possibly a second, long enough to interface with the catalysts. If the electron falls back into the hole, we just get a little heat and we're toast. So we have to make the recombination rate slower than the rate at which we can siphon off the electrons. How in the world are we going to do this?

Rudy Marcus, who's here at Caltech, won the Nobel Prize in 1992 for telling us how to start this process. [See *E&S*, Fall 1992.] About 40 years ago Rudy predicted what's now called the inverted effect: that you could build molecules where the recombination rate would go up for a while as you increased the energy of the reaction, and then come back down. Everybody knows that if you put more energy into something, it should go faster and faster, but very few people thought that at very high reaction energies the rates would start going back down again. But in the mid-'80s, a group at the University of Chicago and the Argonne National Laboratory was finally able to verify the inverted effect. After that, we figured out how to build inverted-effect molecules, and now we've got tons of them.

On the opposite page is an example that uses ruthenium and pyridine. (You'll notice that we've gone back to wimpy organic molecules. We figured, why not—let's throw a bone to the organic chemists. So I apologized to all the organic chemists and said we're gonna need your wimpy molecules back again. But this time we put them on some nice, firm metals.) The ruthenium-pyridine



Above: This ruthenium (Ru^{+2})-pyridine (the hexagon) complex obeys the Marcus inverted effect. When it absorbs light, an electron immediately jumps from the ruthenium to the pyridine. And there the system pauses for a moment—the charge-separated state (negative in blue, positive in red) is slightly less energetic than the initial excited state, but the electron has too much energy to fall back into the hole very fast.


complex absorbs visible light, which helps with the light-capture part of the problem. The electron transfer from the excited state to the charge-separated state (with the electron on the pyridine, and the hole on the ruthenium atom) is very, very fast—about 50–60 femtoseconds (a femtosecond is a thousandth of a picosecond!). And the charge-separated state lingers for about a microsecond—a millionth of a second. So Rudy got us part way, but we need to buy another factor of a thousand to a million in time—to reach a millisecond or a second in separation time—to give the energized molecule enough time to make hydrogen and oxygen.

To figure out how to buy this extra time, Jay Winkler [PhD '84], who's a Member of the Beckman Institute, I, and several other researchers in the Beckman Institute Laser Resource Center here at Caltech have been studying electron tunneling over long distances. [See *E&S*, Fall 1991.] In general, these experiments involve putting an electron donor and an electron acceptor on opposite ends of various kinds of molecules—for example, a protein structure called an alpha helix (shown in yellow in the figure at right), or another protein structure called a beta strand (shown in green). Then we zap the electron donor with a laser beam, kicking an electron loose, and we measure how long it takes the electron to arrive at the acceptor. We've also varied the molecular geology, as it were, within a given structural class, and we've shown that electrons tunnel through some features faster than others. This means that electrons can't go through empty space—they have to tunnel through chemical bonds. And you can see for yourself that if electron tunneling goes through bonds, it's going to take forever to go through the twisty α helix, while the β strand is a straight shot. So, in relative terms, a β strand is a conductor, and an α helix is an insulator.

And, of course, God figured this out before we did. The structure of the natural photosynthetic


So I apologized to all
the organic chemists
and said we're gonna
need your wimpy
molecules back again.
But this time we put
them on some nice,
firm metals.

reaction center is known, and guess what? There are α helices everywhere you look! Light hits a pair of chlorophyll molecules, creating an excited state that does electron transfer to a charge-separated state. Then the electron hops around from molecule to molecule and finally ends up on a quinone. And between the quinone and the first chlorophyll pair is a long stretch of α helix, so the




A molecular drag race. Each electron donor (D) is separated from its acceptor (A) by some 20 carbon-atom diameters as the photon flies. The electrons, however, have to tunnel along the chemical bonds shown in color. It takes an electron 1–10 seconds to travel the corkscrew α helix (yellow, at left), but a mere thousandth of a second to zig-zag down the much straighter β strand (green, at right).

Right: Nature's photosynthetic center, as embodied in *Rhodobacter Sphaeroides* and elucidated by Erimler et al. Light hits the pair of chlorophyll molecules (red), sending electrons leapfrogging along the accessory chlorophylls (orange) and pheophytins (green) to the quinones (blue). The yellow curls between the quinones and the chlorophyll pair are α helices; the long tail from one quinone that appears to reach back up to the chlorophyll sticks out of the plane of the page and is not an electron-transfer route. The purple blob is an iron atom.



electron and the hole don't crash back together and make heat. The initial hops from the chlorophyll pair onward only go forward, never backward, because of the inverted effect. But by the time the electron arrives at the quinone, the inverted effect has played out, so you need an insulator—an α helix. It's beautiful.

The one remaining problem is the catalytic step—how do we split water after we've manipulated all these electrons and holes? Nature does it with the manganese system below, which no chemist has ever been able to make. And the manganese system evolves O_2 very efficiently, but we don't want to do that. Plants make O_2 and carbohydrates, but we want to make H_2 for



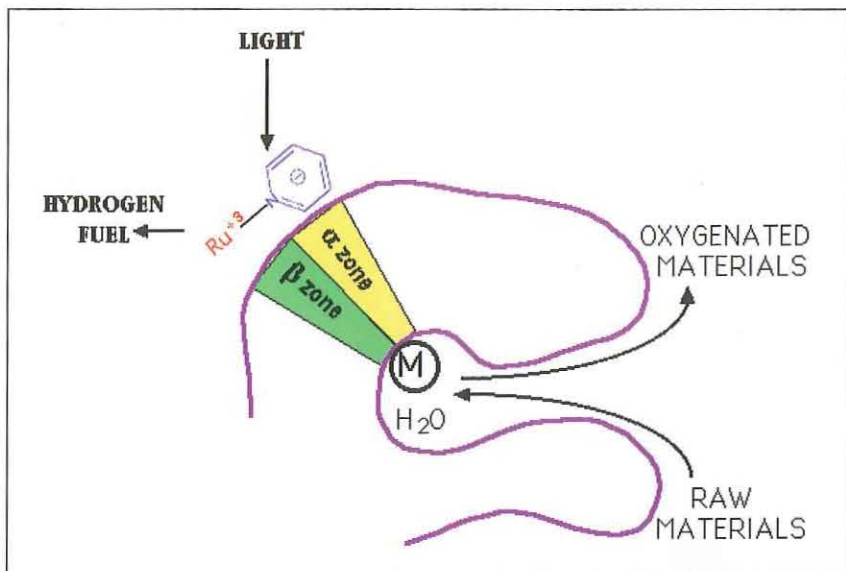
Above: The molecular scaffolding surrounding the oxygen-generating manganese system (the cluster of Mn's and O's at the bottom) is almost an inverted version of the photosynthetic center. Electrons follow the arrows from the P680 (a kind of chlorophyll) to the Pheo (a pheophytin) to the quinones Q_A and Q_B .

fuel without liberating O_2 . It turns out that O_2 causes the poop-out problem. Whenever we make O_2 , we run the risk of also producing an energetic form of oxygen called singlet oxygen—a very reactive molecule that oxidizes everything, including the organic molecules that make up our photosynthetic system. That's why we all eventually die—oxygen is good for us, but it's also bad for us. In the long term, we all get oxidized—some of us faster than others.

So we're now trying to figure out how to make catalysts that will take electrons and holes and make new materials as well as fuel. In a new and improved water-splitting scheme, oxygen won't be allowed to escape, but instead will be incorporated into molecules that can serve useful purposes. One possible system for doing this is shown at right. It's purely a conceptual drawing—nobody has made such a thing yet, but a lot of people are working in this area. In addition, we're going to have to learn how to get materials from carbon dioxide, water, and solar energy, because we're going to run out of hydrocarbons before long. People say we have enough coal for a few hundred years, but a few hundred years is nothing in the life of our planet.

I'm encouraged by our prospects, because we've made enormous progress in related fields—for example, the methanol fuel cell is here right now. I believe you'll see one in your car, and possibly in your house, in the next few years. It uses a ruthenium-platinum anode to oxidize methanol to carbon dioxide. This generates electrons that flow through a wire to run your car motor or your computer or whatever, and eventually return to the cathode and reduce oxygen to water. So you put methanol in and get electricity out.

In about 10 or 15 years your car will run on electricity. General Motors' EV-1, which uses a big lead-acid battery that goes only 50-60 miles on one charge, is already on the road. But the car of the future will have an advanced, lightweight battery, not a lead-acid battery; a methanol fuel cell; and, eventually, a solar-paneled roof, so that the battery can charge during the day while the car is sitting in the parking lot. The fuel cell will be for long-range driving—to San Francisco, say; if you're just driving around town, you'll run on sunlight. GM has a methanol-fuel-cell car on the drawing board now that will come out in 2004, and they aren't the only company that's getting into the game. Mercedes-Benz has plans to mass-produce 100,000 such cars a year by 2005, according to the *Los Angeles Times*, and Toyota, Chrysler, and Ford have all announced fuel-cell projects. These cars will not only be environmentally better, they'll be better in absolute terms. People are going to *want* these cars! Who wouldn't want something that looks like a Porsche, burns rubber, and gets (in the more advanced models that have solar roofs) over a thousand miles per gallon? So when they become available, people are going to




Above: The "T. rex" scheme for splitting water. (This just sounds more dignified than the "sock puppet" scheme!) The dinosaur's head is a shorthand representation of the sort of molecular scaffolding that surrounded the chlorophylls and quinones on the previous page.

Sunlight hits the ruthenium-pyridine complex as before, briefly separating an electron (blue) and a hole (red) under the Marcus inverted effect. Meanwhile, in T. rex's mouth, the metal atom (M) strips an electron from a water molecule while catalyzing the conversion of oxygen to oxygenated materials. The liberated electron tunnels up the green β strand to the surface and falls into the hole on the ruthenium, leaving the electron on the pyridine free to make hydrogen fuel. The electron doesn't tunnel backward because of the insulating α helix (yellow) between the pyridine and the metal.



Designed for suburban commutes and urban drives, the EV-I is already on the road.



be knocking each other down to get them. There's going to be a big rush, and then people will say, gee, while we're at it, let's convert Los Angeles into a solar city.

And I'm not just talking about places like L.A. and Tucson—everybody's got sunlight. Every year since the mid-'70s, physicists, chemists, and engineers have improved solar efficiencies as well as the durabilities of solar materials, and we're almost to the point where we can get good solar conversion, using layered semiconductors, even when we can't see the sun through the clouds. In the next 10 years we'll have solar power generation working so well that it will be a shame if we don't use it widely on the planet.

In my view, the overall change to solar fuel will take a dramatic demonstration of a vehicle (or a stationary power plant) that runs much better than anything else has ever run, because changing the infrastructure in this country is going to be very, very difficult. There's great resistance to change, and it's going to cost a lot of money. Right now, we could put silicon on the roof of your house. You'd sell electrons to the power plant during the day and buy them back at night, and at the end of the month you'd get a check instead of a bill. But if we put a solar roof on every house in Los Angeles, there'd be a huge capital cost, so there'd have to be tax credits to help out. All sorts of people would have to clamor for change, in order to get the appropriate tax legislation passed. On top of all this, Southern California Edison would have to agree to buy electrons from some of its customers.

I hope we will phase out fossil fuels in a rational way, but I'm afraid the change won't be orderly. We won't say, "Gee, we need these hydrocarbons for materials. We're burning the most valuable resources we have. This is really stupid." We should be using clean fuels wherever we can right now, and we should be learning how to convert the rest of our fossil-fuel-guzzling technologies before we're forced into a corner by some war or other disaster. It's going to cost a lot of money to make this conversion, but after we do it, we'll have a planet that runs on clean fuel—an ecologically sounder, infinitely more sustainable place to live. □

Harry B. Gray, Beckman Professor of Chemistry and director of Caltech's Beckman Institute, came to Caltech from Columbia University as a visiting professor in 1965, and returned for good in 1966. For his research in bioinorganic chemistry and inorganic photochemistry, he received the National Medal of Science from President Reagan in 1986 and the Priestley Medal from the American Chemical Society in 1991. His work on electron tunneling in proteins is supported by the National Science Foundation and the National Institutes of Health. This article was adapted from a SURF (Summer Undergraduate Research Fellowship) seminar.