Solar Fuel I: Rods and Stones

By Douglas L. Smith

"He who cannot store will have no power after four" is a favorite mantra of Nate Lewis (BS '77, MS '77), Argyros Professor and professor of chemistry, in a nod to the late Johnnie Cochran Jr. But the stakes here are even higher than the outcome of O. J. Simpson's trial for double murder—Lewis is talking about saving our civilization by running it on solar energy. Solar cells make electricity, and if you have a secret method for storing enough juice to keep the lights burning across America—or even the Los Angeles basin—while the sun is down, he advises you to "go out and buy electricity at night for a nickel a kilowatt-hour, sell it to all your neighbors for 25 cents a kilowatt-hour by day, and then fund your own solar-energy research institute with the profits."

The most promising avenue is to store solar energy in chemical form, as a fuel that can be used on demand. Plants do this via photosynthesis, turning sunlight, water, and carbon dioxide into sugar molecules containing lots of high-energy chemical bonds. Lewis and many other researchers worldwide are trying to take a leaf from photosynthesis, as it were, by finding a method for turning sunlight into fuel—without all the complex cellular machinery. "Photosynthesis is beautiful, and it has many lessons to teach us," Lewis says. "Birds with feathers are not good models for 747s, but we knew birds could fly, so we were inspired to build airplanes."

The easiest fuel to make is hydrogen gas, H_2. Simply stick a couple of electrodes in a vat of water and run the current through them—the positive electrode makes O_2, and the negative one makes H_2. "Burning" the H_2 in a fuel cell later gives you water back again, plus the stored electricity. Small, pilot-scale electrolyzer plants, as water-splitting facilities are called, already exist, but they use a lot of really expensive components, and they're not exactly the sort of thing you could put on the roof of your house, or even in your back yard. Lewis's goal is to repackage tons of complicated equipment into tiny assemblages of cheap materials that could be installed almost anywhere. “We’re going to make fuel from the sun with no wires,” he says. To do this, the light-harvesting and water-splitting machines have to be one and the same, not an assortment of components spread out over a solar farm.

The Hole Story

A conventional solar cell uses a cheese-and-cracker arrangement of two slabs of silicon, *n*-type and *p*-type, and the place where the cheddar meets the Triscuit is called the *p-n* junction. *N*-type silicon is “doped” with traces of phosphorus or arsenic, atoms of each of which have five electrons in their outer shell instead of silicon’s four. Four of the dopant’s electrons pair up with electrons in adjoining silicon atoms to make the chemical bonds that hold the slab together, leaving the fifth one at liberty, if given a little shove, to wander through the crystal. Similarly, in *p*-type silicon, the dopant—boron or gallium—has only three electrons in its outer shell. The dopant atom makes three and a half bonds, if you will, to its silicon neighbors, and the absent electron, or “hole,” is essentially a free-floating positive charge. An electron from an adjoining atom will readily jump over to fill it, moving the hole to the atom the electron abandoned.

When light hits the solar cell, the photons will penetrate some distance into the silicon before being absorbed, and the longer the wavelength (or the redder the light), the deeper they go. It takes 300 microns, or three-tenths of a millimeter, of silicon to soak up all the wavelengths that make up visible light, and herein lies the problem. When a silicon atom in either layer absorbs a photon, an electron gets excited, creating an electron-hole pair. If the silicon atom is in the *n*-layer, the electron can happily skate away to become part of the current, but then what happens to the hole? In order to maintain the charge balance and keep the current flowing, the hole has to make its way to the safe haven of the *p*-type layer, on the opposite side of the *p-n* junction, running a gantlet of onrushing electrons in the process. (Because they are out of their element, if you will, these much-set-upon holes are called “minority carriers.”) Any defect in the crystal’s lattice structure makes a natural trap where electrons and holes can recombine, dissipating their energy as wasted heat. Free electrons generated in the *p*-type layer face a similar struggle. The upshot is you need ultrapure—and therefore very expensive—silicon to make solar cells.

Grow Your Own Aspen Forest

But what if, instead of using a flat plate of silicon, you made nanorods—teeny, tiny cylinders oriented so that the light shines down their central axes? The minority carriers would just have to reach the cylinder’s skin to make it to safety, and the odds of an electron falling into a hole drop drastically. “This is exactly what nature does,” remarks Lewis. “This is what leaves do. The cells absorb light over the long axis and transport the heavier particles—the water vapor and the nutrients—over the short axis. That’s what the retinal rod cells in your eyes do. This principle is found throughout nature, and yet it’s never been implemented in solar-cell design.”

Even better, you don’t need a layer of the opposite type of silicon—simply toss the rod in a beaker of salt water, which conducts electricity quite readily, and the liquid-solid interface acts as the finish line for the minority carriers’ sprint to safety. Once they cross into the liquid, they are ready to split water molecules into H_2 and O_2. All you have to do is coat the silicon surface with catalysts that will facilitate the reactions.

Harry Atwater, Hughes Professor and professor of applied physics and materials science, and Lewis tested the notion by growing what Lewis calls “an aspen forest” of nanorods, each some three microns long and about a quarter of a micron thick, on a conductive backing. The complex growth process used standard technology for making silicon chips, but for reasons that will become clear momentarily, Josh Spurgeon (MS ’06), a Lewis grad student working on this joint project, made the first nanoforest of cadmium selenium telluride instead. This material generates electricity by absorbing blue light, which takes about one micron’s worth of semiconductor. But such a photocell doesn’t work under red light, which takes three microns to absorb, because the minority carriers can only make it about a micron before being sopped up by a lattice defect. Spurgeon shone light in all the colors of the rainbow on a submerged nanorod array and on a CdSeTe slab made by the same method and compared the results. The slab and the rods both performed about the same under blue light, but the nanorods kept on soaking up photons all the way through the visible spectrum. “But we didn’t want to wait around for the rods to grow,” says Lewis. His grad student James Maiolo treated a silicon slab with hydrogen fluoride, etching so many holes in it that they overlapped to leave spires. “We wanted to make the equivalent of rods to see if they would work, or if the theory was missing something.” It wasn’t.

Light will travel a certain distance into a crystal before being completely absorbed. When a photon is absorbed by a solar cell, a pair of opposite charges in the form of an electron (e^-) and a hole (h^+) are produced. If either one has to travel too far through unfriendly territory, as the hole does in this traditional slab design (below left), it is likely to recombine with its opposite and no current will flow. In the nanorod design (below center and, blown up, below right), the light can penetrate to any depth and the holes can still escape out the sides and into a conductive liquid.

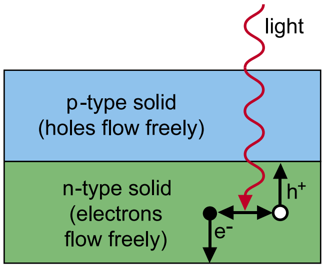

Slab design

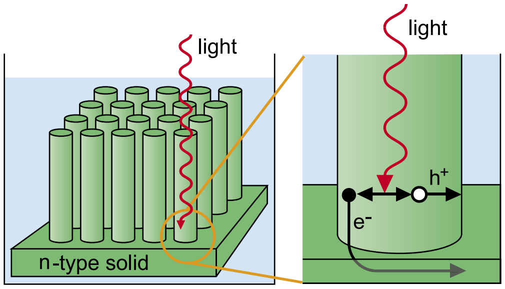

Rod design

Not So Fast . . .

If making hydrogen from sunlight were that simple, however, you wouldn't be reading this article. There are still three big problems to solve. The first is really only a problem when you try, as Lewis and Atwater are, to make fuel from the sun without using wires: H_2 and O_2 production need to be kept separate. This makes collecting the hydrogen easier, and prevents the gases from recombining to make water as they form. All fuel-producing solar cells to date isolate the oxygen-producing electrode from the hydrogen-producing electrode with either distance or a membrane. The simplest design uses a single semiconductor to absorb the light. In this case, it's an *n*-type semiconductor that acts as the oxygen-producing electrode. The holes leap into the drink and suck electrons out of the water molecules in a process called oxidation, creating O_2 gas and protons—the electron-deprived nuclei of hydrogen atoms. The electrons go through a wire to the hydrogen-producing electrode, usually platinum metal or the like, where they rejoin the newly produced protons that have diffused through the water (or across the membrane) to make H_2—a process called reduction.

The second problem with splitting water is fundamental chemistry: oxidation is brutal. Strong oxidants tear chemical bonds apart—fire is an oxidative process, and think of what would happen to your favorite shirt if you spilled concentrated bleach on it. Since the goal is to get rid of the wire—and, incidentally, the very expensive precious-metal electrode to which it's attached—the hydrogen-producing electrode needs to be a *p*-type semiconductor fused seamlessly with the *n*-type to inject electrons into the water. Unfortunately, known *p*-type semiconductors don't last long when exposed to strong oxidants. This includes elemental silicon itself, says Lewis. "Silicon is not stable under oxidizing conditions in water. It makes sand—SiO_2."

The third problem we will come to shortly.

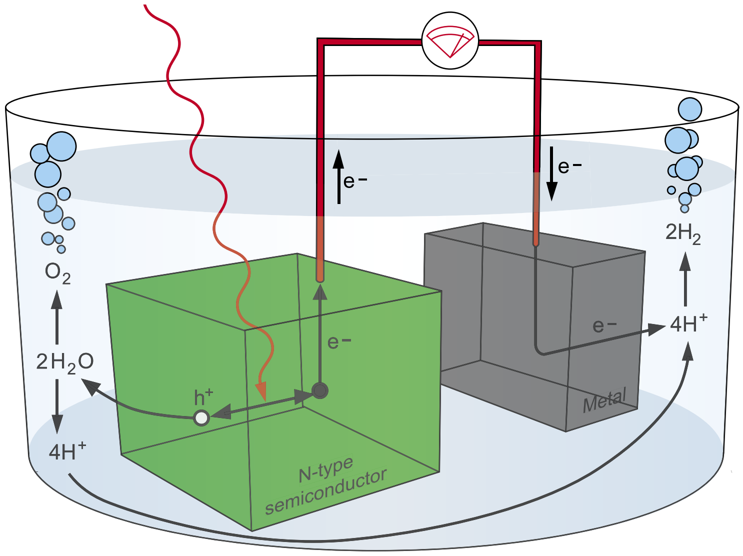

In this water-splitting solar cell, the light absorber is the oxygen-making electrode, an *n*-type semiconductor. The surrounding water acts as the *p*-type layer, accepting the holes, which convert water molecules into O_2 and H^+.

Fool's Gold and Rust

Proving that the nanorod idea worked allowed Lewis's and Atwater's groups to consider using nontraditional materials to collect light. "For instance," says Lewis, "iron oxide has been rejected in solar cells because it's very defective as found in nature. You need three microns of rust to absorb sunlight, and the carriers can only move 20 nanometers, so the other 2.99 microns are just wasted. But if we make rods 20 nanometers wide and three microns long, we can still move the carriers sideways. The same thing occurs with fool's gold, which is iron pyrite, FeS_2. So we are now trying to make solar cells out of fool's gold and rust, and you can decide whether we're completely crazy or only partially crazy."

The good news is that there's a whole class of defect-riddled minerals that can absorb light to make electricity while withstanding the oxidative assault. Some of them, such as titanium dioxide, or TiO_2, the coloring agent in white paint, are also in abundant supply and available at rock-bottom prices—two absolutely fundamental considerations if the technology is to be adopted on a global scale. "They are very stable materials. These literally *are* rocks; they're oxidatively inert, they're photochemically inert," says Lewis. "The bad news is that they would work better if we were inhabiting the moon instead of the earth, because they have a band gap where they start to absorb light that is about three electron volts or higher." The band gap is the energy difference between the state that an electron is in when it is firmly attached to its atom—the so-called "valence band"—and the state it gets kicked into once it absorbs a photon and gains enough energy to hop from atom to atom—the "conduction band."

Unfortunately, three electron volts is borderline ultraviolet. Earth's ozone layer screens out most sunlight at higher than that energy, so the efficiency of those materials is limited to a few percent at best. The light that actually makes it down to us starts at about 2.75 electron volts on the purple end, hits peak intensity at about 2.35 electron volts in the blue-green region, and then slowly trails away down through the infrared. The optimal band gap for catching everything with a single material is about 1.4 electron volts, which is actually in the near-infrared. It only takes 1.23 volts to split water, but there's a catch—and that's the third problem.

The band gap not only has to be big enough, it has to be in the right spot, energetically speaking. For example, iron oxide's band gap is 2.2 electron volts, but the electrons thus liberated aren't very energetic. From an electrochemical point of view, making H_2 from water and making O_2 from water are two entirely separate reactions; worse, they operate at different energy levels. If you use an iron-oxide photon absorber, "you can still make oxygen, but you can't make hydrogen any more,"

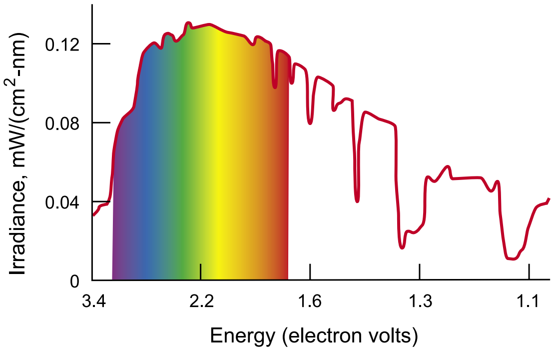

This plot shows the intensity of the sunlight reaching the earth's surface versus the amount of energy per photon in electron volts.

Lewis says. "Absolute energies count. The electrons that are promoted to the conduction band aren't reducing enough to reduce water to make hydrogen. It's like climbing a wall—going three feet up doesn't put you a foot over the top if you start out four feet down."

ALL THE OXIDES THAT ARE FIT TO PRINT

Just because no known oxidatively stable material has a big enough band gap at the right energy level doesn't mean that such a substance doesn't exist. The quest for something that would fill the bill was taken up in 2006 by Rhodes Scholar Todd Gingrich (BS '08) and grad student Jordan Katz (PhD '08), as detailed in *Caltech News* 2008, No. 1. The idea was to test tiny samples of thousands of candidates as fast as possible by modifying a $140 inkjet printer to spit swatches of metal salts dissolved in water onto thin glass plates. The print head mixes the salts in exacting proportions, just as it would faithfully render a sunset from cyan, magenta, yellow, and black inks, but the output looks more like a color-calibration page than a vacation photo—up to 250 tiny test squares on a single sheet. A 500°C oven bakes the dissolved metals into metal oxides. "We choose our metal oxides in an Edisonian educated-guess fashion," Lewis says. "We see which ones give us the right energy levels to reduce water to hydrogen, and then we investigate which ones stay around after one day of operation. We're building up a database to try to give theorists a guide to help us find the next generation of materials."

Lewis's lab is sticking to metal oxides because, well, they're already as oxidized as they can get. "The chemist's intuition is that whatever oxidizes water to oxygen will be sufficiently oxidizing that if we don't have an inherently stable, nearly ionic chemical bond such as a metal oxide, we're going to eat our semiconductor alive." This tactic has uncovered a couple of promising materials already, but Lewis is not betting the solar farm on one printer—he's also working on a Plan B.

Things would be greatly simplified if one material didn't have to do everything. Nature uses such a divide-and-conquer strategy—a plant's photosynthetic system uses two light-absorbing components that together provide the oomph to drive the reaction. The photosynthetic machinery within such a cell straddles a membrane that isolates and protects the hydrogen-producing apparatus from the harsh oxidants on the other side. In a membrane-based, two-part inorganic system, the hydrogen-producing side could be powered by metal sulfides. Metal sulfides have chemical compositions similar to the active catalytic sites in the enzymes used by a family of photosynthetic microbes called green sulfur bacteria to eke out a living in low-light, oxygen-free environments. These hardy bugs have even been found in hydrothermal vents called black smokers on the deep ocean floor, where they subsist on the dim glow of the hot vent itself. The oxygen-producing side could use good old iron oxide, and if the two types of nanorods could somehow make electrical contact through the membrane, the system wouldn't need wires.

"This gives us the ability to retain the chemical stability of two separate materials," says Lewis, "while not also asking one single material to deliver the full energy needed to split water. The design

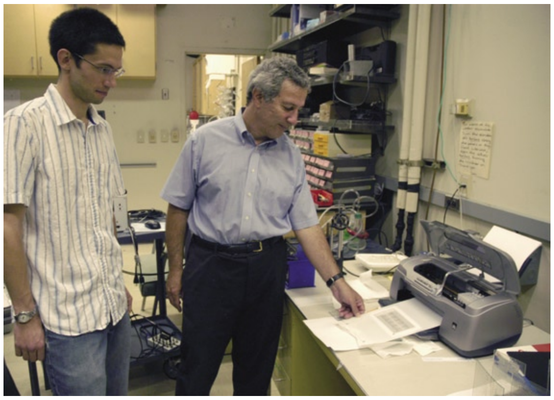

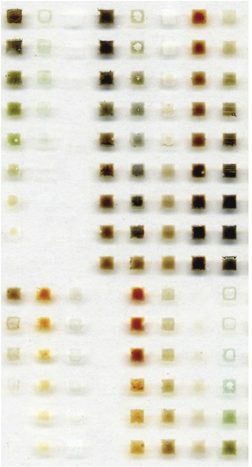

Top: Katz (left) and Lewis watch a glass plate come out of the printer. Right: A portion of a typical plate.

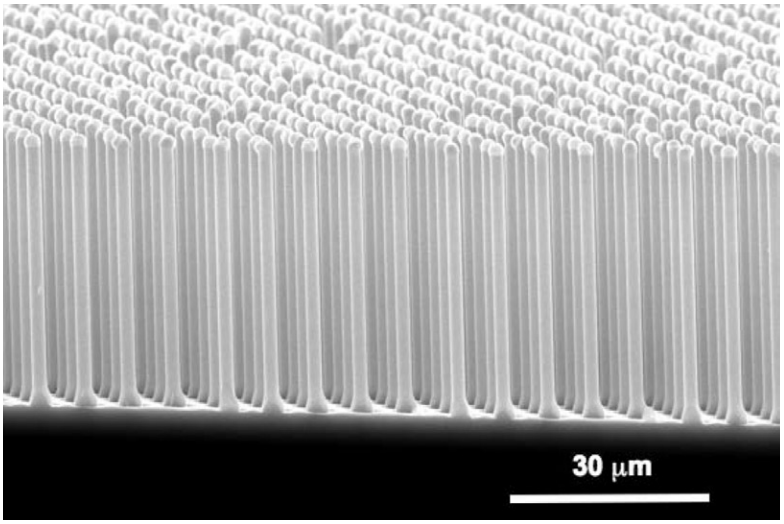

A cornfield of silicon nanorods grown from seeds of gold. The scale bar is 30 microns, or millionths of a meter.

is actually the inverse of a fuel cell, where hydrogen and oxygen come in, protons go across the membrane, and out comes electricity. Instead, in our cell, in comes the energy source, protons go across the membrane, and out comes chemical fuel. We're driving it in reverse."

Bringing in the Sheaves

This all sounds horribly complicated, but the Lewis lab has hit on a very clever way to make the integrated system. The process can be done on any scale, which is vital for any technology intended to supplant fossil fuels, and if the work with fool's gold—an iron sulfide—and rust is successful, the raw materials would be dirt cheap, which is even more important. For the moment, the people putting the pieces together are using silicon nanorods, because the industrial techniques for making them are well established.

> "Some day some fellow will invent a way of concentrating and storing up sunshine . . . When we learn how to store electricity, we will cease being apes ourselves; until then we are tailless orangutans." — Thomas Edison

The process starts with an ultrapure silicon wafer, but it can be reused over and over again, as grad student Spurgeon has shown—successive generations of rods grown on the same wafer look just as good as the first batch. The rod-growing technique was developed by grad student Brendan Kayes (MS '04), who is coadvised by Atwater and Lewis, along with Atwater postdoc Mike Filler and coworkers. It begins with a pattern of tiny gold dots that can be laid down on the wafer by any of various means, and which become the growth template. Their diameter is the size of the nascent nanorods, and the spacing between them is the spacing between the rods in the array. The rods are grown through a process called chemical vapor deposition, in which the gold, now in little molten globules, reacts with gaseous silicon chloride ($SiCl_4$) in a hydrogen atmosphere at a temperature of 1,000°C. When the solution becomes supersaturated with silicon atoms, the excess precipitates out at the bottom of each droplet, becoming one with the silicon wafer below. "The gold rides on top, as the silicon starts to deposit, and you continue that growth process," Lewis explains. "It literally grows like cornstalks in a field." Turn off the gas flow, and you can stop the growth at any height you like. Even better, says Lewis, "We can do this with copper. We can do this with nickel. We don't need gold as the catalyst."

However, there's a really good reason people have been making solar cells out of slabs instead of rods. If you picture the layer of atoms that forms the surface of a slab, a silicon atom sitting there can have three of its four bonds—which point to the four corners of a tetrahedron centered on the silicon atom—easily connected to other silicon atoms in the surface layer. The fourth bond sticks out into space. These "dangling bonds" are prime locations for charge carriers to recombine. A slab has limited surface area, while a nanorod is basically all surface area—you couldn't ask for a better way to maximize the number of dangling bonds. "So if we don't have a way of tying up those dangling bonds, then we don't have much hope of actually making good devices out of them," says Lewis.

The nanorods end the manufacturing process with their dangling bonds capped with hydrogen atoms. Unfortunately, Si—H bonds are oxidatively unstable in air. "We do a lot of fundamental chemistry in our group, so we developed a method using phosphorus pentachloride or PCl_5, and zit medicine, benzoyl peroxide, to convert the silicon-hydrogen bonds to Si—Cl bonds," says Lewis. Another set of steps replaces every chlorine atom with a methyl group, CH_3. "The silicon-carbon bonds are sufficiently strong that they fool those surface silicon atoms into thinking that they're just like the bulk material. There are no dangling bonds left at the surface—less than one electrical defect in every 100,000 surface atoms." This work was done by then-postdocs Hossam Haick and Patrick Hurley, along with collaborators Peidong Yang of UC Berkeley and his grad student Allon Hochbaum.

Solar Panels You Can Unroll

For the next step, postdoc Kate Plass has developed a remarkably simple process for embedding the nanorods in a membrane. She uses a type of silicone rubber called polydimethyl siloxane, or PDMS—a common waterproof sealant that, says Lewis, is "more affectionately known to people on the street as fish-tank goop.” Plass carefully pours a thin layer of PDMS over the nanorods until she almost, but not quite, covers their tips, leaving them exposed to make contact with the water later. She gives the surface a close shave, literally, with a razor blade, and then she peels the goop off in one smooth, sharp motion, like doing a hot-wax treatment to get rid of unwanted body hair. (Howling and jumping around like Mel Gibson in *What Women Want* or Steve Carell in *The 40-Year-Old Virgin* is optional.) Just as the hairs go with the wax, the nanorods go with the goop, shearing off neatly at their bases and leaving the silicon wafer ready to take another round of metal dots. The liberated nanorods retain their original parallel orientation and regular spacing, only now “they are in a piece of plastic that we can roll up,” Lewis says.

Plass has so far made nanorod-embedded membrane squares more than a centimeter on a side. Scaling the process up should not pose any great difficulties, but it will take something more than just building very large tweezers.

Before being embedded, the *p*-type silicon nanorods in a demonstration system would be coated with cobalt, which acts as a catalyst for H_2 production. Postdoc Steve Maldonado is working on that angle. The *p*-type membrane would be laminated back-to-back with a second membrane containing *n*-type, oxygen-producing nanorods of iron oxide coated with a cobalt-oxide catalyst; sandwiched between the two would be an electrically conductive layer.

The idea is that eventually this technology might be developed to the point where the solar-membrane laminate could be sandwiched between two more layers of plastic that would allow water to bathe the nanorods. The entire assembly would then be manufactured in large rolls that would be sold at home-improvement stores, the way rolls of insulation are sold today. (The main difference, of course, would be that you’d unroll this stuff on top of your roof, instead of in the attic underneath.) With an eye toward this, the Lewis lab is starting a collaboration with Mory Gharib (PhD ’83), the Liepmann Professor of Aeronautics and professor of bioengineering, to measure the membranes’ mechanical properties to discover how far one can be bent without fracturing the rods.

There is one last hitch, however. Remember those naked protons that got stripped off the water molecules and that have to cross the membrane and reunite with their electrons to make H_2? Fish-tank goop is impermeable, so a more porous membrane will have to be found. “Either that,” says Lewis, “or we’re going to have to poke lots of little tiny holes in it.”

“We haven’t yet built a full system,” says Lewis. “But we have a clear path toward being able, in two sheets of plastic, to at least demonstrate a prototype system that could take water to hydrogen and oxygen with sunlight. As Melvin Calvin said at my first Department of Energy meeting in 1982, ‘It is time to build an actual artificial photosynthetic system, to learn what works and what doesn’t work, and thereby set the stage for making it better.’” Now, only 25 years later, it may finally be on the brink of happening. A Caltech-MIT initiative called “Powering the Planet,” about which you will read more in the next article, has set itself the goal of creating a solar-fuel generator that uses Earth-abundant elements, needs no connecting wires, can be scaled up with existing manufacturing technologies, and is 10 times more efficient than photosynthesis.

Calvin wasn’t the first to have this idea. Lewis is also fond of quoting Thomas Edison, who in 1910 told Elbert Hubbard, in Volume 1 of the latter’s series of books called *Little Journeys to the Homes of the Great*, “Some day some fellow will invent a way of concentrating and storing up sunshine to use instead of this old, absurd Prometheus scheme of fire. I’ll do the trick myself if some one else doesn’t get at it.” (Apparently, it never quite made the top of his to-do list.) Edison went on to say, “Sunshine is spread out thin and so is electricity . . . the trick [is], you see, to concentrate the juice and liberate it as you needed it. . . . This scheme of combustion to get power makes me sick to think of—it is so wasteful. . . . When we learn how to store electricity, we will cease being apes ourselves; until then we are tailless orangutans. You see, we should utilize natural forces and thus get all of our power.”

Adds Lewis, “In the end, by far the biggest energy source available to humans is the sun. Now, we can’t afford to use it, but in the future we can’t afford *not* to use it. Nature figured this out, but from a ‘product’ point of view, photosynthesis is a failed solution; the fastest-growing plants, on a yearly basis, store less than 1 percent of the total sunlight that hits an average acre over a year. We can do better, we have to do better, and we will do better. The question is not if, but when and how. Someday, you’ll paint your house with solar paint [that’s another article for another day—*DS*], and instead of putting up solar pool heaters, we will roll out solar fuel generators.” □

Nathan S. Lewis (BS ’77, MS ’77), Caltech’s Argyros Professor and professor of chemistry, got in on the ground floor of the solar-fuel game as an undergrad working with Harry Gray (see following article). Lewis got his PhD at MIT in 1981, and earned tenure at Stanford before returning to Caltech in 1988. He’s been working on the photochemistry of semiconductor-liquid interfaces (including an effort to develop solar paint) ever since, with time out for such side projects as developing an electronic nose and helping to debunk cold fusion.

PICTURE CREDITS: 20-23, 25 — Doug Cummings; 23 — Bob Paz; 24 — Brendan Kayes; 25 — Michael Filler

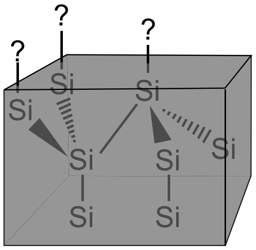

Every silicon atom has four bonds, not all of which are shown here. The “dangling” bonds stick out of the surface of the crystal, and *something* has to cap them off.

A piece of nanorod-embedded plastic.