THE HUBER MILLS
digital
power report
Powering the Telecosm
September 2000/Issue 6
Published by Gilder Publishing

Fuel Cells I: Small and Cool

The most interesting technology we’ve come across in hydrogen generation systems has been developed by Proton Energy Systems

Sure, the fuel cell is a Powercosm technology. But whose cell? The long and winding fuel cell trail leads us to the edges of the technology curves, and to the least-hyped units. If you like the technology at Ballard Power (BLDP**)**, Global Fuel Cell (GFC), Plug Power (PLUG), Avista Labs (AVA), H Power (HPOW), or Analytic Power (IPO later)–stick with it. If you want to go long on greening the planet. Just don’t call it a Powercosm technology. We’ve tried. It can’t be done.

The government subsidizes fuel cells to save the earth. Investors invest because green virtues are bound to sell, somehow or other, sooner or later. Maybe. But this letter is about the kind of power that will save rooms and buildings full of chips and lasers, and huts full of fiber-to-wire switches and RF transmitters. We’re after technologies that will keep the silicon hot, not the planet cool.

FuelCell Energy (Danbury, CT, FCEL) has one. It comes in a forest-green can, 12’ high and 12’ around. Think of it as a humongous, very hot (1200°F) battery, good for 1,000 kW of power. “Big and hot” defines one promising path for the fuel cell. “Small and cool” defines the other. Manhattan Scientifics (MHTX) is doing some promising (albeit, still speculative) work with a nano-fuel cell plant, one-millionth the power (1 W) and the size of a credit card.

The space in between–where most of the fuel cell crowd has clustered–may prosper for other reasons, but it isn’t about to prosper at the heart of the Powercosm. For now, the technology to watch in this middle ground is technology that can provide the fuel–hydrogen–on which medium-size, mid-temperature fuel cells depend. Proton Energy Systems (Rocky Hill, CT, still private, but with an IPO scheduled for late September) has developed one promising unit.

Our first six issues looked at technologies that, though far from fully mature, are mature enough to propel significant commercial sales in the marketplace today. The same cannot be said of any of the fuel cell technologies we discuss here, not even the ones we most like. We’re picking among distant technology futures here. Some are quite a bit less distant (in our view) than others. It is equally important to emphasize that “fuel cell” defines a huge range of technologies, engineered to generate anything from a single watt to a megawatt. Doing cursory justice to the technology requires more space than even this double issue allows. These two issues represent our first excursion through this space, not our last.

Old Technology

The most remarkable thing about fuel cells is that there are so many of them, and that they’re so widely hyped, and that the technology is so neat and that … so few of them are practical outside the laboratory. Only one company sells real fuel cells for real dollars to earth-based customers today. To get a piece of that action, however, you have to buy into a much larger business package, one that includes elevators, air conditioners, and helicopters, because the company is United Technologies (UTX). UT’s ONSI division delivered its 200th commercial fuel cell last March.

These are real machines, with a solid record. Two 200 kW ONSI units situated on the fourth floor of the Conde Nast Building, at Four Times Square in the heart of Manhattan, power a huge neon sign on the building’s façade. ONSI would join our PowerPanel if it were an independent company; we’re hoping for a spin-off some day. FuelCell Energy is on the brink of becoming the second manufacturer selling fuel cells to real commercial customers. And that’s the core of its business. All the rest of the pack are still in "beta" or "alpha." Or even further back up the R&D trail.

Which is what you'd expect in an industry rolling out brand new technology. Except that the fuel cell isn't new at all. The first unit was developed in 1839, long before the internal combustion engine, when British physicist Sir William Grove discovered that the electrolysis of water–using electricity to break down the powerful bond that unites hydrogen and oxygen–works in reverse, too. Allis-Chalmers put a fuel cell into a farm tractor in 1959. NASA has been putting them on spacecraft ever since the Gemini program. The oxygen tank that exploded on Apollo 13 was there to feed the fuel cell. As NASA learned the hard way, the fuel cell devil is in the details.

Few have mastered them. Fewer still are anywhere close to mastering the details that are important in the Powercosm–the details that make a fuel cell not just utterly green, but utterly reliable too.

The Short Wire

For Powercosm purposes, the fuel cell starts out in just the right place–on the premises, next to the load. Distance, recall, is what makes the hundred- and thousand-mile trip for electrons on the grid doubly unreliable. The grid's thousands of miles of exposed wires come under frequent assault, from lightning, ice, tree limbs, and cars that bring down poles. The grid's own customers undermine it too–by abruptly adding or subtracting large loads that send spikes and sags rippling up the line, and by occasionally pushing total demand beyond the grid's capacity to supply. An independent power supply on your own premises adds a short-wire alternative, and is thus the key to boosting overall reliability.

How much independent power? Fuel cells from Ballard, Analytic Power, Plug Power, and H Power span the 5 kW to 250 kW range. Those under development by Siemens Westinghouse Power and Atek Corp range from 25 kW to 25 MW. FuelCell and ONSI build for the 200kW to 2 MW space–big compared to most of the rest, though still tiny by utility standards. These are all nice numbers. The lower (<100 kW) range devices match the kinds of electrical loads now created by tens of thousands of wireless base stations, apartment blocks, and the rapidly multiplying wired McMansions. The larger units are in the right power range to be deployed, singly or in multiple-unit arrays, in the rapidly emerging CLEC hotels, network storage centers, Web caches, and server hotels. The big Powercosm hotel companies–Akamai, Exodus, Covad, Digex, Level 3, Equinix, Global Center, Qwest, Verio, WorldCom and others–are currently projected to build out over 20 million square feet of silicon floor space, with individual power requirements running 5 MW to 50 MW, for an aggregate total of at least 2,000 MW. There's plenty of opportunity to supply the high-9s, mission-critical power that such companies definitely need.

And as it happens, there's a real shortage of substitutes to serve this crucial stretch of the power curve. Almost all the high-duty-cycle commercial turbines are very big units reaching 200 MW. Very few manufacturers offer high-duty-cycle commercial turbines in the wide gap that separates the 30 kW to 75 kW microturbines (see July DPR) from the 50 plus MW utility units. (There are lots of mid-range aviation turbines, but they're short-duty-cycle and high-maintenance.) The venerable diesel spans the whole range–but again, most of the high-duty-cycle diesels (e.g. marine-type diesels) are very big.

And who wants a hot, noisy, smoky diesel running in the basement, or in the big power "closet" next door? Nobody–and that's the second major point in the fuel cell's favor. Its potential to save the planet is beside the point from the silicon's perspective–but its compact, clean, quiet operation is a cardinal virtue nonetheless. It lets you deploy the short-wire power plant in the heart of high-tech workplaces and wired homes, especially in the center of congested cities.

The fuel cell seems to fit the bill. The first serious commercial development was for the original cosmic workplace–the space capsule–where compact, clean, cool and quiet mattered a lot to the quality of the environment inside the capsule, however little anyone cared about environmental quality outside it. So far as compact goes, the alkali fuel cells used by NASA offer the highest power-to-weight ratio of any electrical generator ever devised, 50 times the energy density of the best batteries, lots of power without combustion or vibration in a technically magnificent–though tricky and dangerous–package. (United Technologies' International Fuel Cells division is the main NASA supplier.) Get past tricky and dangerous, and a fuel cell would seem to be an equally promising technology for providing short-wire power in the space capsule environs of techie cubicles, GHz server racks, and wireless base stations.

Finally, and most prominent in the public chatter, we have the green case for the fuel cell. Cool. Clean.

Editors	Peter W. Huber Mark P. Mills
Publisher	Richard Vigilante
Managing Editor	Debi Kennedy
Designer	Julie Ward
President	Mark T. Ziebarth
Chairman	George Gilder

The Digital Power Report is published monthly by Gilder Publishing, LLC. Editorial and Business address: P.O. Box 660, Monument Mills, Housatonic, MA 01236.

Editorial inquiries can be sent to: peterhuber@gildertech.com or markmills@gildertech.com. Single-issue price: $50. For subscription information, call 800.261.5309, e-mail us at help@gildertech.com, or visit our website at www.Powercosm.com

Hyper-efficient. These attributes matter for high-9s purposes too, because they win the fuel cell an essentially free pass–or even better, a subsidized pass–from zoning authorities and the green police. Deploying a big Caterpillar diesel in the parking lot in Manhattan, Kansas isn't too hard, but try installing one on the fourth floor at Four Times Square. Distributed generation, the key to short wire and high-9s, runs into lots of regulatory obstacles when it looks and smells like a truck. Fuel cells don't. That's a big difference, distant though it is from the engineering fundamentals of providing ultra-reliable power.

With so much to commend the fuel cell, it would seem hard to imagine a more promising Powercosm technology. Short wire, compact, clean, quiet, and inside the building–and greening the planet to boot. That's certainly how a great number of investors seem to have seen it. For a while, many fuel cell stocks were behaving like the dot-coms in their heyday.

We wish those recent high-flyers well. Lots of money and talent is certainly being thrown in their direction and, given time and resources, engineers can accomplish almost anything. But you won't find those companies on our Power Panel–not yet. For now, they're green plays, political plays, regulatory plays–anything but Powercosm plays. In this vast thicket of hype, the technologies we like are either too small and cool, or too big and hot, to have attracted much notice.

Small and Cool

Every fuel cell has two core parts: electrodes and electrolyte. Together, they sustain an electrochemical reaction in which (typically) hydrogen surrenders its electron at the anode, oxygen picks up an electron at the cathode, and the oppositely charged hydrogen ions migrate through the electrolyte, unite to form water, and exit the fuel cell**.** This is a 1.23 V reaction (for pure hydrogen-oxygen, at least), so for higher voltages you layer cells in series, much as cells (of a different kind) are layered inside a car battery. The total power generated depends on how fast you can effectively pump hydrogen and oxygen through the system, which depends in turn on the physical structure, surface area, and catalytic power of the electrodes.

Easy to describe; not at all easy to build into a practical, reliable system. If it were easy, fuel cells would already be ubiquitous. They aren't.

There are two basic fuel cell models. Call them big-and-hot versus small-and-cool. The hotter you operate, the easier it is to get the electro-chemistry running. (True of all chemistry; try starting your car when your battery sat overnight at 40° below zero.) The hot group demands much less of its electrolyte, and of its source fuel–the higher temperature itself does the heavy electrochemical lifting, so to speak. The cool crowd, by contrast, depends on a much more complex electrolyte and requires a very refined fuel–extremely pure hydrogen.

If you're going to run cool, the tough part is getting things to run at all. In the end, most of the small-and-cool players strike an uneasy compromise, they run at about 200°F, very cool compared to the hot technologies discussed in the companion issue, but not cool enough for very high reliability. What they do without in heat they make up for in expensive catalyst–platinum–and sophisticated electrolyte. The recent surge of interest in fuel cells can be traced back to breakthroughs in platinum chemistry and solid electrolytes achieved only a decade ago.

The remarkable thing about fuel cells is that there are so many of them, they are so widely hyped and the technology is so neat, but so few are practical outside the laboratory

Lead by John Appleby, one of the godfathers of fuel cell chemistry, scientists at Los Alamos National Laboratory, Texas A&M, and elsewhere found ways to deposit ten-atom-sized platinum particles on pure carbon, and bind them to a fuel cell's electrodes. This reduces by thirty-fold the amount of platinum required to operate a fuel cell, cutting catalyst costs from almost $200 of platinum per kW to about $7. DuPont then developed a semi-permeable, solid electrolyte–a sulfonated fluorocarbon teflon-like polymer, called a "proton exchange membrane" (PEM). The "Nafion" membrane comes in very thin (0.1 inch) sheets, and replaces unstable or volatile liquid electrolytes. Other solid-membrane electrolytes have been developed since, but DuPont's remains the most widely used. Fuel cell engineering has indeed come far. Far enough? Put it this way: Just give it a perfect PEM, and the fuel cell will indeed be cool, compact, perfectly green, and ultra-reliable too.

How perfect a PEM? First, a PEM that can run room-temperature, cool and still deliver serious amounts of power. At 200°F, most of the current PEM designs still run too hot for the long-term health of the PEM. At those temperatures, PEM's physically degrade, and much too fast and unpredictably for the Powercosm's objective of durable, trouble-free operation. So why not run cooler? Because lowering temperature lowers electrochemical performance.

The perfect PEM must also be immune to carbon monoxide. Ideally, it should process natural gas, gasoline, or other fossil fuels directly. None of the mainstream PEM-cell developers can. Their systems simply cannot tolerate any CO contamination above a few parts per million. Ironically, that superficial and entirely negative fact about the fuel cell is what makes it the darling of Kyoto greens. The PEM-based fuel cell emits no carbon into the air because it will admit no carbon monoxide into its own innards.

Finally, the perfect PEM will survive the twin assaults of heat and carbon for tens of thousands of hours–i.e. years–of operation, without requiring so much platinum that it has to be sold by Tiffany's. To our knowledge, nobody has yet come close to making such a PEM. The PEM crowd is desperately trying to balance temperature against catalyst to strike an attractive balance among power output, dollars, and durability. It hasn't found one yet. The devices built so far, around far-from-perfect PEMs, reflect a long series of compromises, in which reliability is undercut again and again to keep the cell itself affordable, and to raise its power and thermodynamic efficiency. For Powercosm purposes this is completely backward. The digital customer will pay sky-high prices for higher-9s electrons. Shedding 9s to cut emissions and save fuel may save the planet, but it will sink your chip fab, your dot.com, or your digital factory. Lots of people are green. Few are <u>that</u> green.

From a structural engineering perspective, the PEM fuel cell mirrors the grid's own frailties: The long, thin–and therefore unreliable–wires of the grid give way to the flat, thin–and therefore unreliable–acres of PEM. The industry's press releases are forever bragging about advances that make a membrane both thinner and stronger. But this only spotlights the dilemma. The membrane has to get thinner and hotter to work better. But it has to get stronger and cooler to last longer.

Bottom line: PEM fuel cells remain–for now at least–the quintessential low-9s technology, delicate, short-lived, and demanding exceedingly high levels of precision and purity, both in their own manufacture and in the fuels that they process. No PEM fuel cell today is reliable enough for even a car's low-duty cycles (i.e., 5,000 hours of operation over a vehicle's life). None comes even close to the reliability requirements of high-duty-cycle power application (30,000 plus hours). Yes, the fuel cell up in a space capsule is a high-9s device–but only because it begins with carbon-free hydrogen fuel on board, has a very low-duty cycle, and is meticulously maintained. You can keep a PEM-based fuel cell running here on earth in much the same way, burning copious amounts of technician time to nurse the thing along. But adding technicians and meticulous maintenance protocols isn't the right way to add reliability in most of the earthbound Powercosm.

Almost all the NASDAQ fuel cell favorites are pursuing the small-and-cool PEM. Sooner or later, one of them is going to come up with something approaching our "perfect" PEM. Possibly Ballard. The company has some 350 patents in the field, its own proprietary polymer PEM membrane, and ample capital. And if not Ballard, someone else. Polymer chemistry is a very fecund field of R&D, and at least 85 organizations around the world (four dozen in the U.S. alone) are engaged in PEM research. The company that builds a truly robust PEM is going to do very well indeed. We just don't know which company will. And we're inclined to doubt that a big breakthrough is imminent.

Scale Down

If membrane-based electrolyte fuel cells are going to find a place in the Powercosm, they will get there first by getting small. We're watching for the membrane-based micro- or nano-fuel cell–a 3 W to 30 W device–that can trickle-charge (or entirely replace) a laptop or cell phone battery.

The physical laws that govern a piece of engineering determine whether it will scale up well–or, conversely, whether its destiny (if it has one worth bothering with) lies in scaling down. With technologies that depend on surface effects and surface area forces, performance improves as size shrinks. Capacitors, silicon gates, and catalytic converters, are all surface-effect technologies. Smaller invariably lasts longer and works better here, because micro technologies are governed by diffusion and atomic-range forces, both of which become (relatively) faster and stronger as area increases and spacing shrinks. Large inertial and thermal systems, by contrast, get more robust and efficient as they get bigger. A giant power turbine is a lot more efficient than a tiny one. Scale such systems down too far, and they are over-powered by frictional effects and thermal losses. Surface area is their enemy; bulk, their friend.

As we discuss in Part 2 of this double issue, FuelCell Energy achieves reliability by building larger units propelled by higher-temperature chemistry. In this kind of heat-centered design, bigger works better. FuelCell has accordingly scaled up and up, toward the 250 kW to 1,000 kW level. ONSI has gone in the same direction. So too will Westinghouse and others. For now, all the reliable fuel cells are big and hot.

The other direction to go is down, way down, to the micro- or nano-fuel cell. PEM membranes perform better as they get smaller. Keeping them cool gets easier–small systems are self cooling–and cool membranes last longer. Lower temperature means you have to use more catalyst, which means that you have to build for markets that can afford platinum-coated power. Which, by happy coincidence, happen to be the markets that most value ultra-compact design. If carbon contamination is still a problem, and it is, you add still more catalyst (of another kind) to take care of it, and price be damned. Thin membranes are structurally weak, but the smaller the scale of your system, the stronger they effectively become. Just as pound for pound, insects are far stronger than we are, pound for pound, a square inch of membrane is stronger than a square foot.

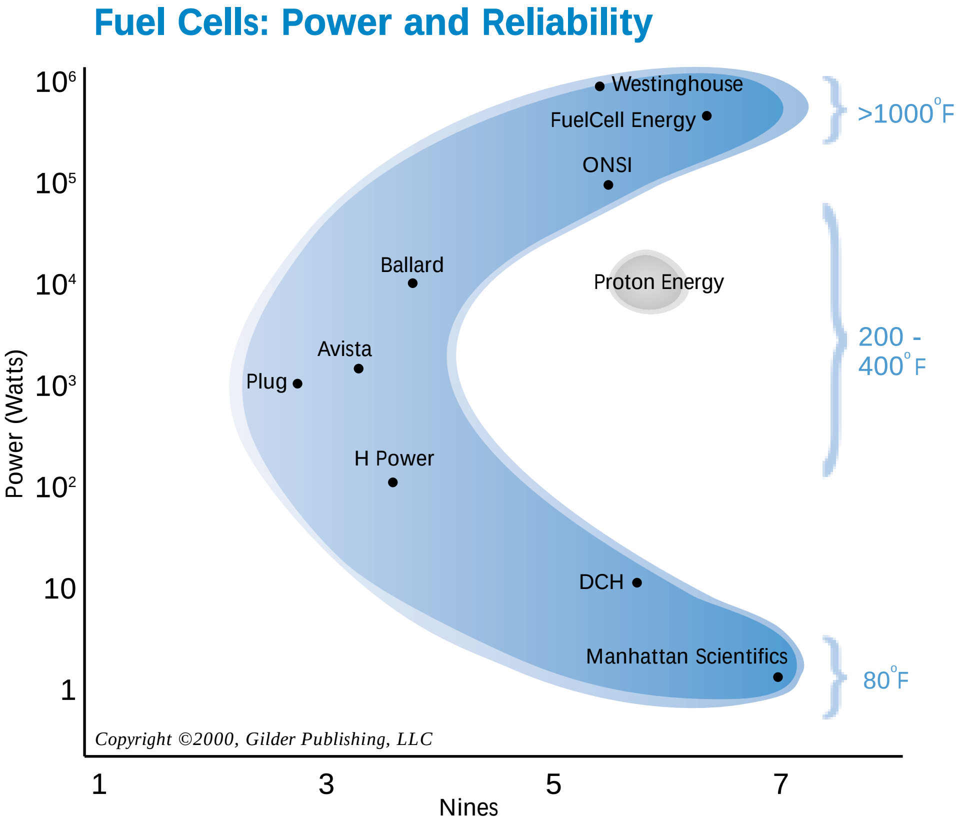

The fuel cell technologies with Powercosm promise, provide more reliability (higher 9s). They fall into two groups: low-power and cool, or high-power and hot. Most of the fuel cell hype, however, has focused on the less reliable technologies in the middle.

The one company we've come across that seems to have grasped this is such an insect itself that we hesitated to name it. We can only emphasize, again, that we're describing an elegant technology solution that makes sense to us, not a commercial product with customers, still less a set of corporate books.

The owner of the technology, Manhattan Scientifics (MHTX), is led by CEO Marvin Maslow. The company had three employees in 1997, went public in January 1998, and is up to a 30 today. Jack Harrod, the COO, had a thirty-two-year career at Texas Instruments, where he specialized in somewhat analogous technology niches, such as TI's Directed Light Processor, micro-mirrors on silicon. Manhattan Scientifics is an incubator, nothing more.

Currently nestling under its corporate wings are four little techno-eggs. One is holographic storage, sounds like neat stuff, but not our field. The second is a nano-membrane for water purification. Same reaction. The third, an investment in NovArs, a German developer of a 70 W, PEM-based, micro-fuel cell. We're getting warmer here: NovArs has developed a system prototype to replace batteries for portable communications for the U.S. Army. It delivers four times the performance from half the weight of the Army specs. And the fourth is a nano-fuel cell, a credit card-sized power plant.

Manhattan Scientifics acquired the nano-cell IP in 1997 from Bob Hockaday, a physicist at Los Alamos National Laboratory. In 1990, well before *Wired* or NASDAQ got on the bandwagon, Hockaday set about making tiny fuel cells in his basement. In due course he took an entrepreneurial leave of absence from the Lab, then turned around and signed a cooperative research agreement with the Lab for additional technical help, and set to work. His three core patents (a fourth pending) now cover the critical structure and concepts surrounding a heretofore ignored architecture for nano-fuel cells. In the fall of 1997, ready for seed capital and the first steps towards the commercial world, Hockaday made several dozen VC presentations, earning tepid responses, until Manhattan Scientifics' Maslow heard the story. Now, with nearly a decade of diligent work into his project, Hockaday is on the verge of producing the first, pre-commercial prototype. If all goes well, commercial production could begin in 2001.

So how is this different from any of the dozens of other companies that keep promising something commercial "real soon?" Perhaps it isn't. But if it is, it's different because Hockaday has followed the PEM membrane to where the technology's underlying physics inexorably leads. And if it works–commercially that is, because it already works in the lab–the Hockaday cell could end up providing a whole lot of portable electrons to the wireless Telecosm. No battery anywhere on the horizon can run a high-bandwidth, color-screen, wireless, portable Web terminal–the inevitable, convergent destiny of cell phone and "personal digital assistant"–for weeks, or even days. The Hockaday cell could.

Hockaday builds a very small system, which makes it inherently sturdy. Inertial and thermal stresses are lower in smaller systems, so membranes last longer. The cell runs at room temperature, which is where delicate membranes want to operate if they're going to last. And it runs on a readily available carbon fuel–methanol. Methanol isn't cheap as fuels go, but there's plenty for sale–it's in your windshield-washer fluid. As it happens, the hydrogen in methanol is also less tightly bound to the carbon atoms than it is in natural gas, making it easier to coax it free.

Per Watt generated it takes a lot of platinum to metabolize hydrogen out of methanol at room temperature. So

Hockaday uses it. And carbon monoxide poisoning is still a problem–Hockaday uses the Dupont PEM, too. But Hockaday deals with the carbon poisoning by adding still more catalyst inside the cell itself–ruthenium oxide, which immediately converts the CO to harmless CO_2, before it can corrupt the platinum alongside.

The Hockaday cell is built on a simple polymer plastic substrate, about three square centimeters, irradiated to make it porous. The DuPont PEM electrolyte is dissolved and infused into the substrate. Standard vacuum deposition is then used to layer the platinum ruthenium catalysts on the surface. The cell's anode is a bottom layer, the cathode the top layer, in a multi-layer "sandwich" (around the substrate and electrolyte); the kind of construction that's standard in the chip-fab, circuit board and even lithium batteries industry. Using straightforward photolithography, a very simple channel is etched in the polymer to create a fuel path and links to the anode and cathode.

A one-ounce methanol container supplies the fuel. The methanol reaction clocks in at 0.5 V; ten to twenty cells are linked in series. Use it to trickle-charge a cell phone battery (for example), and it will run the phone for about 20 weeks on an ounce of methanol.

Whether the Hockaday cell emerges as a winner in the nano-fuel cell space remains to be seen but the engineering fundamentals are on the mark

None of this comes cheap on a dollars per kW basis. It doesn't win you anything in the way of efficiency or environmental protection either. What it can potentially do is significantly outperform other technologies currently used to power small-scale, off-grid applications. The electron competition here isn't a Honda engine or a utility's giant turbine–it's the battery. Portable battery power is very expensive from the get go, and users already pay huge premiums for higher energy densities. The price paid for a kW in a Palm is 10,000 times what the auto industry will pay for a kW to electrically propel a Pontiac. The right fuel cell can almost certainly deliver what a wireless Palm needs, and soon. The fuel cell for a Pontiac will take quite a while longer.

Others are on the same trail. Researchers at Sandia Labs are pursuing a similar design based on a silicon substrate. Fraunhofer Institute for Solar Energy Systems (Freiburg, Germany) with Siemens, has built a 20 W fuel cell prototype for notebook computers, using a solid-metal-hydride fuel source. Last May, Case Western researchers announced a miniature fuel cell the size of a pencil eraser, again built with chip-fab technology. Christopher Dyer (formerly Bell Labs, now at Motorola) has developed an ultra-thin film platinum-based fuel cell. DCH (DCH) has obtained intellectual property from Los Alamos lab to make a small circular fuel cell that can be packaged as a D-cell battery-sized stack; fueled from a metal hydride canister, it lasts three times as long as a comparably sized pack of nickel-cadmium batteries.

Whether the Hockaday cell will emerge as a big winner in the nano-fuel cell space remains to be seen–it's still way too early to tell. But the engineering fundamentals are all on the mark. Membrane technologies should lead to nano and micro, to small structures, low temperatures, smart combinations of catalysts, and the cleverest micro-manufacturing. Which is exactly where Hockaday, and now Maslow, have followed them.

Deliver Hydrogen

So why can't the manufacturers of larger fuel cells for stationary applications use ruthenium oxide or bottled hydrogen to beat the carbon problem, and more platinum to beat the temperature problem? As membranes grow, the more fragile and unreliable they become; catalysts get prohibitively expensive; the bottles get big, clumsy and expensive, compared to stationary technology alternatives available at higher power levels. The portable markets targeted by nano-fuel cells are already paying huge price premiums for inferior power. The stationary markets aren't.

If small-and-cool fuel cells are going to move up into the higher power ranges, whether in offices or cars, they will most likely do so not by piling on more catalyst or perfecting the refractory PEM, but by perfecting its fuel. The challenge is to deliver perfect fuel–hydrogen–to the highly imperfect PEM fuel cell. Much as their developers hate to dwell on the fact, all the small-and-cool PEM-based fuel cells at hand require a very pure hydrogen feed. And all still have a world of trouble getting it where they want it.

Hydrogen is a plentiful fuel, if you happen to be in a space capsule. Pounds being the hardest thing to get into orbit, NASA propels its rockets with liquid hydrogen plus liquid oxygen, the highest energy-density fuel combination to be found outside the nucleus of an atom. For NASA it makes perfect sense to use these perfect fuels to generate electricity too, which is just what the alkali fuel cell does. Small wonder that it runs clean and compact. All the dirty work is done back on earth.

The greens' fondest fuel cell dreams revolve around the magical emergence of a hydrogen-fuel economy. Their hydrogen is to come from solar powered electrolysis of water. Perhaps some day. But a hydrogen-centered energy economy is about as close to current reality as a hydrogen-fusion economy, which is to say, way too distant and speculative to matter.

The main problem with hydrogen isn't production–production is easy, though not cheap.

Distribution is the problem. To move a lot of energy, you either have to liquefy the hydrogen or pump the gas under high pressure. Under high pressure, however, it's quite viscous. But because of its low molecular weight, it's terribly leaky and it sneaks through even the tiniest cracks, which makes for serious safety problems. Hydrogen can be adsorbed on to metal hydrides and transported that way which works fine for smaller, portable applications, but isn't yet practical for moving large quantities of the gas.

The companies that are solving the fuel cell's fuel problem fall into two main groups: Those that strip hydrogen out of carbon fuels on site, and those that generate hydrogen from water on site.

The first and cheapest alternative is to strip hydrogen out of a fossil fuel. Natural gas is the main candidate, because that's something our energy economy does already distribute and store. You can also pull hydrogen from methanol and coal, but both are richer still in PEM-poisoning carbon. So natural gas is what just about every real, terrestrial fuel cell configuration depends on–not as a fuel, but as a chemical feed stock from which pure hydrogen can be extracted. But it takes quite some extracting. The "energy system of the future" turns out to depend a whole lot on the "chemical industry of the past," a chemical refinery, basically, and a very good one. All the PEM fuel cells require a reformer that gets CO levels down below 2 ppm.

Very few companies, if any, know how to make an ultra-reliable reformer of the right grade to feed a PEM the ultra pure hydrogen needed, at a cost that doesn't sink the whole venture before you even get to the fuel cell. Plug Power's stock rose sky high the first of this year after GE signed a contract to buy 485 of its units. The subsequent collapse occurred after Plug Power announced design changes to the reformer stage of its unit that added five years to the already 10-year pre-commercial and commercial delivery schedule. This change allowed GE to bail out of the original contract, which it promptly did, staying in for a diluted, very long-run (face saving) play. Plug has since signed on with Engelhard Corporation, a venerable old-world company, to pursue a better reformer.

United Technology's ONSI division has focused a lot of its recent effort on next-generation reformers, an area where they clearly have deep experience. Northwest Power Systems (NPS) is developing units to convert methanol, ethanol, gasoline, diesel, kerosene, methane, and propane into hydrogen. Nuvera (Cambridge, MA) is focusing on commercialization of a Multi-Fuel Processor, the first ever small gasoline- and ethanol-powered reformer targeted at vehicles (they've also shipped a unit to Plug). Other independent reformer makers include Haldor-Topsoe (Denmark) and Hydrogen Burner Technology (CA). IdaTech, a subsidiary of IdaCorp (OR), makes a spark-ignition internal combustor to pre-heat the fuel whether gas or methanol; it is combined with a steam reformer to yield the hydrogen, and in a third step, a hydrogen purifier. The big oil companies are getting into reformers, too. Exxon, for example, has developed a slick prototype of a small reformer to produce hydrogen from gasoline, and has a joint venture going with GM.

Proton's HOGEN could be what it takes to someday make sense out of the most over-hyped but least reliable fuels on the landscape, the sun and the wind

All reformers remain centered on old-world chemical engineering–lots of catalysts, multiple stages, heat, filters, and so forth. This is mostly your grandfather's chemistry here, and if it's going to take something pretty remarkable and unexpected to change that. That's also why much of the reformer development is coming from big oil and big chemical companies–they understand refineries and the chemistry of carbon fuels. If the internal combustion engine is going to be reinvented as a hydrogen machine, big oil will be happy enough to provide gasoline just one tank higher up the energy food chain. The green love affair with fuel cells will cool fast if that turns out to be the most practical way to produce on-board hydrogen in your Pontiac.

There are also companies who will make hydrogen for you on site by running a PEM-based fuel cell in reverse, and renaming it an "electrolyzer." GE originally developed this PEM application for splitting water to make oxygen for Navy submarines; the hydrogen was a waste gas. But the reverse fuel cell also offers a way to store grid electricity, not as a bucket of raw electrons (which is physically impossible), but as the closest practical thing to it, a tank of pure hydrogen, which can then be converted back into electrons as the need arises.

The most interesting technology we've come across in this line of work has been developed by Proton Energy Systems (Rocky Hill, CT)–still private, but with an IPO likely in September. They're already selling viable hydrogen generation systems for the industrial market.

Chip Schroeder, Proton's CEO, is yet another MIT grad (the Powercosm seems to be littered with them) and one of the company's founding four. Bob Shaw, VC of Arete Capital, was the one who almost single-handedly spurred them into the business just four years ago. Schroeder had been at AES, one of the most successful independent electric power companies, with earlier stops in the natural gas industry and on Wall Street.

Proton's HOGEN looks much like a dishwasher on steroids; white, some buttons and gauges, a few connections. Water and electricity go in, 99.999% pure hydrogen and oxygen come out, and without a trace of membrane killing carbon or carbon monoxide. No noise, no chemicals, no pressure system. Ironically, the unit once again houses a PEM membrane.

Run this way the PEM is enormously reliable. Proton uses a PEM that dates back to 1953. The first PEM electrolyzer was developed in 1973, and it now has millions of operating hours behind it. Its feedstock is very pure water, which is easy to obtain. Proton cleans and deionizes its water using standard, readily available purification technology. There's no carbon monoxide in sight, because there's no carbon. The system can run at relatively low power levels in the background, and therefore cool, and most critically, well hydrated–which means less stress on the membrane. The only apparent catch: It takes about 4 kWh of grid electricity to make a quantity of hydrogen that (if run back through a fuel cell) ultimately returns about 1 kWh or so of electricity–a 75 percent loss in raw energy content.

Why then would anyone do it? First, for reasons that have nothing to do with the Powercosm. There are plenty of higher-value uses of hydrogen than generating electricity. Industrial, chemical, and laboratory users spend $1.5 billion a year on hydrogen, and the HOGEN can deliver it at one-thirtieth the cost of hydrogen in cylinders, and from a much more convenient, toaster-sized unit. For Powercosm purposes, the point is to turn the hydrogen back into electricity at times when the grid price is a lot higher than when the hydrogen was made–when the grid is down and won't supply electrons at any price at all.

The point of adding extra hardware on premises isn't to save energy or to make electrons cheaper, it's to make them more reliable, which an electrolyzer plus fuel cell system can indeed do. Proton has built a unit for NASA that can make hydrogen from solar panel electrons on the bright side of an orbit, and make electricity on the dark side. Along similar lines, Regenesys, a subsidiary of Britain's National Power utility, is currently building a hectare-sized sodium bromide/sodium polysulfide reversible fuel cell to serve as a 5 MW to 500 MW regenerative system that will consume power off-peak and produce on-peak.

Proton already has a contract to supply 1,000 benchtop toaster-sized units to Matheson Gas, a supplier of hydrogen to the laboratory market. The dishwasher-sized unit will be commercialized next, targeted at industrial (e.g., powdered metals), and chemical markets, as well as microprocessor fabs, whose appetites for hydrogen (used for cleaning residual coatings and other processes) are growing very fast. The "dishwasher" unit will use an ultra-thin palladium film, vacuum-deposited on stainless steel, to physically filter all traces of water out of the hydrogen, to achieve the purity required in many industrial applications. This will be hydrogen even purer than PEM fuel cells require. Indeed, wet hydrogen is kinder to fuel cell PEMs than dry; for PEM membranes, proper hydration is essential to survival.

While it is (wisely) targeting the industrial hydrogen market first, Proton does grasp how well positioned it also is to fuel the fuel cell. Its NASA unit is a fully integrated file cabinet-sized 200 W reversible electrolyzer fuel cell, built around a single PEM membrane, that can flip from hydrogen generator to electricity generator in 250 microseconds. The commercial target is a 20 kW system using the HOGEN PEM to make hydrogen and oxygen out of electrons and water; and the UNIGEN (a sister PEM fuel cell) to make electrons and water out of hydrogen and oxygen.

Efficient? No. Green? Not at all. The UNIGEN consumes a lot more electricity than the HOGEN returns. But it is a technology that can add reliability. Even the worst utility grids function most of the time, so this kind of configuration can make good sense for adding backup capacity. It's inherently simple and robust. Once it gets to the end-user's premises it depends on the cleanest and least-regulated fuel, grid electricity.

In the end, the hydrogen electrolyzer's green virtues might emerge, too. Proton's HOGEN, or a unit much like it, could be what it takes to someday make sense out of the most over-hyped but least reliable fuels on the landscape, the sun and the wind.

Peter Huber & Mark Mills
August 28, 2000

Announcing...

Beginning with the October issue, the *Digital Power Report* will be available on the web for subscribers only.

Log on to www.powercosm.com and click on the subscriber login button located on the left-hand side of the home page.

Your password will be your email address. If you have already provided us with your email address, you are set to go! If not, simply follow the prompts and fill out the form and your password will be activated within 24 hours.

Richard Vigilante
Publisher