



“Have Method, Will Travel”

by Douglas L. Smith

Chevron's refinery in Pascagoula, Mississippi. The United States consumed some 123 billion gallons of gasoline in 1997, according to the American Petroleum Institute—at 20 miles per gallon, enough to drive the family minivan to Neptune and back 439 times. Photo courtesy of Chevron Corporation.

If it weren't for the oil in your engine, it would seize up in just a few miles of freeway driving. Trust me—I know this for a fact. But if it weren't for a dash of zinc dithiophosphate in that oil, the engine would grind its camshaft and valve train down to nubbins within a few thousand miles. This additive has been standard equipment in all high-quality motor oils since the late '40s, says William A. Goddard III (PhD '64), the Ferkel Professor of Chemistry and Applied Physics, but nobody knew why it prevented wear. “People weren't that interested in why, because it always worked. So why waste your money figuring out how? Problem solved.” The problem is rapidly coming unsolved, however, because federal law requires that by the year 2000, catalytic converters will have to last 200,000 miles. They don't last that long now, because the phosphorus in the additive slowly poisons the catalyst. Furthermore, the additive slowly gets burned along with the oil, and the oxidized zinc becomes part of the particulate ash that goes out of the tailpipe and into the smog. So the oil industry is scrambling to invent a replacement—a phosphorus- and zinc-free ashless inhibitor. Now, replacing something when you don't know how it works is a pretty good trick, so in 1995, the folks at Chevron Research and Technology Company in Richmond, California called on Goddard, the director of the Materials and Process Simulation Center (MSC) at Caltech's Beckman Institute, to help them out.


As the name implies, the center does simulations of materials, chemistry, and biology, making predictions before the experiment is actually done. It's the only chemistry lab at Caltech where the video tubes outnumber the test tubes. The center runs reactions in supercomputers, while collaborating experimentalists and engineers then perform the physical experiments and check the results. Now, you can't throw a brick in a chemistry department without hitting a theorist, and there are a goodish number of people doing com-

putational chemistry these days. (In fact, this year's Nobel Prize in chemistry went to Walter Kohn and John Pople for developing the mathematical and programming techniques that make computational chemistry possible.) What sets the Goddard group apart is that it applies chemical theory to real, industrial problems—real enough that industry underwrites half of Goddard's research.

Most university chemists prefer well-defined problems with sparse data (it's easier to get your own data published that way), whereas industrial chemists are often handed ill-defined problems with lots of data—sometimes several decades' worth. In the case of the chewed-up camshafts, Chevron's files bulged with some 50 years of engine-wear experiments. People would take an engine apart, measure all the clearances with a micrometer, put everything back together again, top off the oil, and run it on a test stand for 60 days. Then they'd tear the engine apart again, compare the clearances to the previous measurements, change the oil, and do it all over again. This isn't a very fast method, and it's not cheap, either—a complete set of tests on one inhibitor costs \$150,000. The relative wear-resistance of various additives had been ranked this way, providing performance trends among similar molecules.


But now Chevron needed to know if some unknown, radically different molecule might be better. In the search for a fast, cheap predictive method, Goddard recalls, “They tried pouring oil over an iron ball and rubbing it. They tried spectroscopy; they tried lots of different things, but nothing correlated with the engine experiments. So they turned to us, and for less per year than the cost of one engine experiment, we came up with a model they could use.”

The first thing the Goddard group did was step back and take a fresh look at all the data, and try to work out what might be going on “from the




Above: A top view of an engine on a test stand. The arrows point to the camshaft.

Below: A single cam, close up. The parallel scratches on what should be a nice, shiny, smooth surface are the wear marks.



viewpoint of what's plausible to us as chemists." This step back before plunging in is vital, because what the industrial collaborator thinks is the problem may not really be the problem. The theorists already had an inkling of the real problem on this one, having just finished designing some improved corrosion inhibitors for Chevron's Petroleum Technology Center at La Habra, just 30 miles from Caltech. Crude oil contains stuff that tends to adhere to and react with the inside surface of a pipe, rusting it away until it eventually begins to leak, causing costly oil-well shutdowns. And it turns out that keeping stuff from sticking to a surface has a lot more in common with keeping a surface from wearing away than you might think.


In the corrosion case, the Chevron chemists had assumed that the dissolved ions in the oil were the instigators, and tried to trap or somehow disable them. Nobody paid much attention to the properties of the pipe itself, and when they did look at it, says MSC project manager Mario Blanco, "People concentrated on modeling pure steel surfaces. We conjectured, and later found to be correct, that hematite was the mechanistically relevant surface." Hematite is one of iron oxide's several guises, and a thin skin of it forms naturally on the pipe's interior—and on the interior surfaces of an internal combustion engine—as soon as there's a trace of moisture. In other words, unless you're working in the vacuum of outer space, the hematite layer will form even as the pipeline (or engine) is being put together.



Above: The crystal structure of hematite predicted by the MSC's model. Iron atoms are rust-colored; oxygen atoms are red. Crystallographic data had shown that each iron atom is bound to six adjacent oxygen atoms, and that the bonds came in two lengths. It had been assumed that the bonds were all of roughly equal strength, but the model showed that the longer bonds (dashed lines) were fundamentally different in character, and significantly weaker, than the shorter bonds. This is why hematite flakes—the weak bonds between the layers break more easily than the strong bonds within a layer.


So postdoc Sunder Ramachandran, Bao-Liang Tsai (PhD '97), and Blanco developed a quantum-mechanical model of hematite from first principles, which, as faithful readers of *E&S* will know, means the dreaded Schrödinger equation, which provides a complete mathematical description of every particle in any atom. Undergraduates have been solving the Schrödinger equation for the one-proton, one-electron hydrogen atom since the 1920's, but as you add more particles to the system the number of terms in the calculation quickly spirals out of control. Recalls Goddard, "When I started in this business some 30 years ago, three atoms was difficult. We've made a lot of progress since then. Now it's hard to do more than about 50 atoms or so." Fortunately, hematite has a simple, repeating crystalline structure, so the investigators were able to model a unit cell, as the basic repeating structure is called, with a few dozen iron and oxygen atoms. The model hematite formed two-dimensional layers, explaining the flaky nature of the real crystal, and the crystallographic coordinates of the atoms in the model matched those derived from the real thing by X-ray diffraction studies—two reassuring reality checks. Once quantum mechanics had nailed down all the atoms' exact positions, the trio of scientists, aided by Member of the Beckman Institute Siddharth Dasgupta, derived a mathematical representation of a force field that numerically described every dimple and dome in the iron-oxide surface, and predicted the force with which that

The best inhibitors snapped into place like Lego blocks, paving the surface one molecule deep—forming a self-assembled monolayer, in the parlance of the trade.



Above: The best corrosion-inhibiting molecules look like tadpoles wearing party hats, with round heads and long, floppy tails. Carbon atoms are shown in gray, hydrogen atoms in white, and nitrogen atoms in blue. The nitrogen atom (arrowed) in the tadpole's chin is the part that actually sticks to the hematite surface.

Below: Atoms don't really have solid surfaces, but if you pretend that they do, it's easier to see how they pack. As the tadpoles line up with the "grain" of the surface, their tails overlap like shingles. The dashes are the nitrogen-iron bonds.




surface would attract or repel any other atom anywhere in space. Such a force-field model, which is much simpler computationally than the original quantum-mechanical model, can be scaled up to perhaps a million atoms more than enough for this project.

This million-atom level—called the nanoscale, because a cluster of a million atoms is about 50 nanometers in size—is where all the action is today, according to Goddard. "Integrated-circuit components are starting to get down to the nanometer scale, so the applied physicists and electrical engineers are approaching it from above. And chemists are starting to manipulate structures that assemble themselves into objects in that size range, so chemists and materials scientists are approaching it from below. And molecular biologists—a rhinovirus, the common cold virus, has half a million atoms in its outer coat. We've done simulations of that. So it's really the frontier, and the computational techniques and modeling methods that apply are not very well understood. Below us is quantum mechanics, and above us are finite-element analysis and other standard engineering methods. So it's open territory for theorists."

With the force-field model of the hematite surface up and running, the next step was to pour water molecules and a candidate corrosion inhibitor on it and see what would happen. The best inhibitors snapped into place like Lego blocks, paving the surface one molecule deep—forming a self-assembled monolayer, in the parlance of the trade. The molecules resembled tadpoles, having round heads and floppy, oily tails 12 to 20 carbon atoms long. Each tadpole would plant itself face-first onto an iron atom, shoving aside the water molecules already nuzzling there. And as rank upon rank of tadpoles lined up with the "grain" of the surface, their tails overlapped like shingles and created a waterproof roof that kept the corrosive ions at bay. "The self-assembled mono-

Zinc dithiophosphate is more avian than amphibian—the dithiophosphate portion looks like an eagle in flight, with the zinc atom clutched between its talons.




Zinc dithiophosphate looks like an eagle in flight, with the zinc atom (blue) clutched between its talons. Atoms of sulfur are shown in yellow; phosphorus is purple. Actually, because a zinc atom has four binding sites, there are always two eagles on opposite sides of the zinc, sort of like a face card in a poker deck. If you tore the card almost in half and twisted it 90 degrees, that is—the eagles are flying in perpendicular directions.

layers form an exquisite, geometrically regular pattern,” says Blanco. “Molecules with the wrong shape don’t fit the underlying structure and are not inhibitors at all.”

Could the same mechanism be involved in wear inhibitors? The two types of molecules are vastly different. Zinc dithiophosphate is more avian than amphibian—the dithiophosphate portion looks like an eagle in flight, with the zinc atom clutched between its talons. The wings are stubby, branched three-carbon hydrocarbon units called isopropyls. So postdoc Shaoyi Jiang released a flock of dithiophosphate eagles into the hematite model, and confirmed what had only been suspected—that once the swooping eagle struck, it let go of the zinc atom and sank its talons into an exposed iron atom instead. Says Goddard, “We calculated the monolayer’s cohesive energy, and found that it correlated inversely with the cam wear.” In other words, the eagle that gripped the tightest protected the best. But Jiang discovered that unlike the tadpoles’ overlapping tails, the eagles’ wings lay flat across the surface, wingtip touching wingtip. So he tried different wings on the same body, and behold! the isopropyl wings fit best. Smaller ones didn’t cover the surface completely, exposing some iron atoms to wear. And larger hydrocarbons got in each other’s way, like fledglings getting too big for the nest. And just as overgrown fledglings tend to fall out, so did these dithiophosphates have a lower cohesive energy. “So we did that the first year,” Goddard says. “And Chevron was very pleased, because it was the first time that they had anything that actually correlated with the engine-wear experiments.”

In the contract’s second year, Chevron suggested 13 new families of potential wear inhibitors, and asked the Goddard group to predict how these untried additives would work. (The MSC works in three-year contracts, cancelable at the end of each year if the client doesn’t feel sufficient progress has been made—to date, not one of




These graceful trefoils, like a window in a Gothic cathedral, appear when the hematite model is seen from above. In the left view, only the very top layer—the surface atoms to which corrosive molecules (or corrosion inhibitors) can bind—is shown. The dashed lines mark the “unit cells,” or repeating units, of the crystal structure. Each successive layer is offset by half a unit cell, so when a second layer is added (right), the trefoils turn to arbesques and the cathedral becomes a mosque. A third layer (not shown) converts the arbesques to hexagons.

Sarah Sanders (BS '98), working as a Summer Undergraduate Research Fellow (SURF), obliged, making the model much easier to use in the process.

Unfortunately, many molecules that work well in the model may be too expensive to be practical, so now, in the project's final year, the theorists are collaborating with synthetic chemists at Chevron to try to pick compromise molecules. "We'll show them a molecule that we think will have really good properties," says Blanco. "And they'll say, 'No, we can't make that. But we can make this, which is kind of close—do you think it will work?'" In the meantime, postdoc Yanhua Zhou and Tahir Çagin, the MSC's director of simulation technology, are adding more detail to the model. Says Goddard, "We're applying pressure and looking at the dynamics of shear as one hematite surface moves past another, just as in a real camshaft. We are learning how shear resistance works. This is the first time that such dynamical studies have ever been done, thanks to some computational breakthroughs that Tahir has made."

The key to the group's creativity, says Blanco, "is the presence of other researchers doing unrelated work. Self-assembled monolayers—thioalkanes on gold—were a curiosity back when we started looking at the corrosion problem, but Jim Gerdy (PhD '96) was doing some work on them. A hallway meeting combined my knowledge of how surfactants work and Jim's ideas on monolayers in a way that applied to these projects. To someone on the outside, things would appear a lot messier (and to some extent they still are), but we could focus on the essential features of these compounds and learn how to improve on them. We call this process 'finessing the problem.'" To that end, the group is home to about 20 grad students and as many postdocs—roughly equal numbers of physicists, chemists, and materials scientists, with a sprinkling of applied physicists, biologists, and biochemists thrown in for spice.

And if ever a project called for finesse, it was another assignment from Chevron that the group took on at about the same time: helping to wring the most gasoline from a barrel of crude oil. Whether you're filling up your Ferrari or topping off your Toyota, you expect what's going into your tank to be pretty consistent. So does your engine, which is designed to burn gasoline with a specific octane rating. But no two samples of crude oil are completely alike. Even crude from a single well changes as the well gets pumped down, says Anil Patel, lead planning engineer at Chevron Research and Technology Company—the dregs are always thicker, as anyone who has ever drunk homemade chocolate milk through a straw knows. To try to limit the variability, most refineries were originally designed specifically for the crude from a given region, or even a particular oil field. But in today's global economy, the crude can come




How the isopropyl eagles stick to the hematite. In the figure at left, the bonds between the sulfur talons and the iron atoms are shown as dashed lines.

Note how the top two layers of the hematite surface are actually distorted by the presence of the inhibitor molecule.

The figure at right shows how the isopropyl wings cover the surface and shield the iron atoms from harm.

two dozen or so contracts has ever been canceled.) Jiang ran all 13 families through the model, and one family emerged as the leading contender. Instead of zinc and phosphorus, these birds were built around an environmentally friendly carbon-nitrogen body. However, a little more testing revealed that the isopropyl wings weren't going to fly any longer. Because the new body contained more atoms, it sat up higher on the hematite, so longer, floppier chains of six or seven carbon atoms were needed to drape down and completely cover the surface.

Jiang rated the 13 body types in his second end-of-the-year presentation to the Chevron folks, and explained how the wings altered the performance. After the talk, he and Goddard learned that Chevron had gone ahead and done some experiments already, based on preliminary results Jiang had given them earlier. The best-performing body in these experiments proved to be the one that the model had predicted, and the wings behaved in real life the same way that they had in the model. "The boss of the whole operation was a couple of feet off the ground," Goddard recalls. "He took a risk to fund us, and we came back with a model that actually made useful predictions. The Chevron group included several very good organic chemists, but they hadn't been thinking about wear inhibitors as chemists because they had no model for what properties of the molecule were important. So they'd try things similar to what they had already studied, instead of reasoning, 'Well, on this surface, this molecule is likely to do this, and it's likely to fit in this way.' So they asked us to put together a graphic interface that would let them look at our results on their own computers, see how the molecules fit, and start thinking about them as chemists." Blanco and



The naphtha fraction of gasoline is a wild assortment of hydrocarbons containing between five and eight carbon atoms per molecule, some of which are shown at left. In these diagrams, each line segment represents a carbon-carbon bond, with a carbon atom and a full complement of hydrogen atoms sitting at each vertex and endpoint.

from anywhere. So refineries are always resetting their process controls to maximize the amount of consistent product that can be made from this inconstant raw material. Chevron wanted a model that would use an analysis of the feedstock to predict how the controls should be set, which is not as easy as it sounds. To see why, we need to take a quick look at how gasoline is made.


Crude oil is a mishmash of molecules containing up to 60-something carbons, so the refinery distills the crude into fractions containing successively larger molecules with successively higher boiling points. The so-called "naphtha" fraction, with five to eight carbons per molecule, becomes a major component of gasoline. Actually, gasoline includes hydrocarbons containing up to 12 carbon atoms, because the distillation process is not 100 percent efficient and a few heavier molecules sneak in. In fact, because of the voracious demand, a lot of the even heavier stuff eventually becomes gasoline, too, through a process called "cracking," which breaks big molecules down into smaller ones. More on that later.

The naphtha fraction includes 265 distinct species of hydrocarbons. There are linear molecules, branched molecules, and even circular molecules. Some of these are desirable; some aren't. Highly branched molecules burn smoothly—they have a high octane rating—while linear ones make your engine knock. Molecules with double bonds have higher octane ratings than the equivalent molecules with no double bonds. And the so-called "aromatic" molecules—benzene, toluene, and the xylenes—have the highest octane ratings of all. Naphtha might have an octane rating of, say, 65, but the unleaded regular that comes out of the pump is 87 octane. And, in fact, says Patel, "We want an octane rating of 95 to 100, because some other components in gasoline tend to be very low octane. The naphtha is processed further to bring the overall octane rating up when we blend the various streams to make


gasoline." And the higher the naphtha's octane rating, the more of the low-octane stuff you can blend with it, and the more gasoline you get.

The naphtha goes to the reformer, which is not a school for wayward youths but a reactor that heats the naphtha till it vaporizes. The reformer then blows the vapor through a catalyst that rearranges low-octane molecules into high-octane ones. While in the catalyst, those 265 species participate in half a dozen different families of competing reactions. Not all of the outcomes are good—a reaction that improves one molecule's octane rating may lessen another's.


Each species undergoes each reaction at a different rate, so the operator adjusts the reformer's settings to maximize the good reactions and minimize the bad ones, based on a quick-and-dirty analysis of that batch of naphtha. Until recently, the lab gave the operator four numbers: the percentage of paraffins (linear and branched molecules with no double bonds), olefins (the same molecules, but with double bonds), naphthenes (circular molecules built around a five- or six-membered ring of carbon atoms), and aromatics; plus a boiling-point curve and a specific-gravity measurement that together provided a rough idea of the size distribution of the molecules. The operator then set the reactor's temperature, hydrogen pressure, and feed rate based on gut instinct and the accumulated wisdom of the plant, often enshrined in hand-drawn graphs taped up in the control-room window. These performance curves were worked out through trial and error, says Patel. "It was strictly an empirical model, with no chemistry behind it. If the new feed was radically different, they'd do a pilot-plant run on it first. But if it was just a little different, say maybe two percent more paraffins than the last batch, they'd say, 'Hey, I can handle this,' and just go ahead and




The reaction families of catalytic reforming. Reactions shown as double arrows can go either way, while single-arrowed reactions are irreversible. In this case, the series of four reactions at left lead to benzene, which is desirable, but the two reactions at right led to smaller hydrocarbons with fewer branches, which is not so desirable.



This sample paraffin isomerization reaction shows how the reforming process works. The catalyst actually uses two different components (the zigzag surface), mixed together in a heterogeneous mass. In the first step (above), a rhenium/platinum component strips a pair of hydrogen atoms off of a paraffin, converting it into an olefin, which then takes wing again.



Next, the olefin slams into the hydrogen-ion-studded surface of a second, "acid" component. The olefin's double bond picks up a proton, or hydrogen ion, converting the hydrogen into a neutral atom and leaving the now-ionized olefin stuck to the catalyst by what used to be the other end of the double bond. The ion flops around like a break dancer as other parts of the molecule try to fill the ionic hole. Eventually, a new double bond establishes itself and the rearranged olefin departs.



And finally, the newly branched olefin can have another run-in with the metallic component, emerging as a paraffin once again but with a higher octane rating.


try it."

But today's instrumental methods can give you tons of data instantly—for example, that this morning's naphtha has 6.3 percent unbranched paraffins with five carbon atoms, 4.7 percent single-branched olefins with 12 carbons, 3.2 percent benzene, and so on. (The breakdown isn't so detailed above 12 carbons, but that's another story.) There aren't enough windows in the control room to post 265 graphs, and even if you had an infinite sheet of glass, there's precious little data at that level of detail to plot. Pilot-plant studies are expensive and time-consuming, so gathering process-control data hasn't kept pace with analytical advancements. Chevron estimated that it would cost several million dollars to do all the experiments needed to find just 20 different rate constants over the three reactor parameters, says Goddard, "so they decided, for a small fraction of that cost, to see if we could predict these things from theory. And I told them there's about a 50-50 chance that we could do it."

The 265 species mentioned above don't begin to reflect the full complexity of the problem. The catalyst rearranges a molecule by converting it to an ion, which you can think of as having a hole where something ought to be. Other parts of the molecule slide into the hole to try to fill it, but this just moves the hole around. Eventually, the ion spits out a hydrogen ion, becoming a stable molecule with a double bond and detaching itself from the catalyst. The original ionized site can be anywhere along the carbon backbone, so the number of different ions increases with the length of the backbone. How the molecule morphs depends on where the hole is, adding another layer of calculations. If you have up to eight carbon atoms in your feedstock, you have 540 different hydrocarbon species—including the intermediate ions and the byproducts with fewer than five carbon atoms—that you have to keep track of in your calculations. (Even though the naphtha fraction starts at five carbons, you can't neglect the smaller molecules—some of the reaction paths cause the parent molecule to break off a carbon atom. If an eight-carbon molecule becomes a seven-carbon molecule, that's fine, but if a five-carbon molecule loses a carbon, it becomes butane—a gas, not gasoline.) These 540 species can undergo 883 possible reactions. And the first set of products can hit the catalyst again downstream to undergo a second set of reactions, and so on.

So the first simplification was to limit the problem to molecules with six or fewer carbon atoms. This brought the number of species down to 85, and the number of possible reactions down to 105—much more manageable numbers.

Second, "We had to assume that the catalysts were perfect," says Goddard. "That they didn't get plugged up, and that they could hydrogenate and dehydrogenate, and protonate and deprotonate when they wanted to. Catalysis is our busi-



Above: The pared-down reaction network for six or less carbon atoms.

Again, reactions with double arrows can proceed in either direction. The color coding matches the graph above right, which shows how the model evolves over time at a temperature of 950° F and a hydrogen pressure of 350 psig, or pounds per square inch (gauge)—in other words, pounds per square inch above normal atmospheric pressure.

ness, and normally we would have looked in detail at what goes on on the catalytic surface, but that would not have been practical here. We could not have dealt with this unbelievably complicated set of reactions if we'd had to worry about the surface too." In the end, the catalyst was represented by seven parameters: one for each of five types of ion rearrangement—the "reaction families" mentioned earlier—that measured how much of a kick it took to get that rearrangement started, and two that reflected the catalyst's ability to ionize the hydrocarbon at two different temperatures.

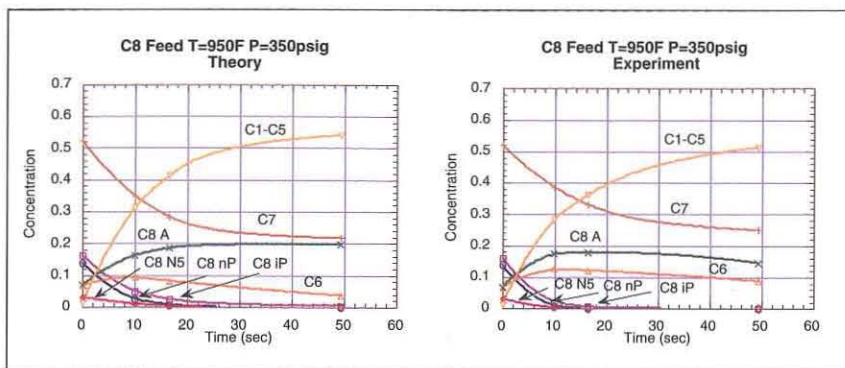
Jason Perry (PhD '94) mapped the pared-down network of possible reactions and used quantum mechanics and the idealized catalyst to calculate the rate constant for each link in the network, creating a model that would predict the product distribution for any six-carbon feedstock at any combination of reactor settings. (Perry, a freelance quantum chemist—yes, there are such things!—and president of First Principles Research, was a grad student of Goddard's.) The researchers checked these predictions against real data taken at two temperatures, two pressures, and three elapsed times (the equivalent of flow rate—higher flows mean the feedstock is in contact with the catalyst for less time). Once the catalyst parameters had been adjusted to correctly predict the six-carbon products, the researchers successfully expanded the model to predict product distributions for up to eight carbon atoms.

As with the wear-inhibitor project, the researchers developed a graphic interface before giving the model to the engineers. "It's the simplest model to use they've ever had," says Perry. "Yet it gives the most detailed results. And Chevron uses lots of different catalysts, but this model could apply to all of them with the appropriate tweaking—the seven parameters all have physical meaning, so it's easy to adjust them." Patel is now using the model for process design on the pilot-plant level. "This was a real success," says Goddard. "We said

we'd need three years, and it took two." "It would be cool to push the model further," says Perry. "But we've now moved on to modeling the cracking process, which is where the real action is, in collaboration with Michael Shippey [PhD '78], a senior research chemist at Chevron. Instead of doing a few hundred species going up to eight carbon atoms, we're modeling the reactions of millions of species going all the way up to 60 carbon atoms. The accounting gets a little tricky, but the principles are the same."

These projects are typical of how the group works. When an industrial partner signs on, Goddard hires a postdoc to work on the project full-time, and assigns a senior staff scientist at the MSC to direct the effort. The Caltech people and the industry people meet monthly to update each other on what's going on. Information flows in both directions, and a casual comment by one person can spark a revelation elsewhere (or dredge up a forgotten tidbit of information). "But you've got to be eyeball to eyeball," says Goddard. "If you don't have this sort of direct, continuous interaction, you don't get the right information." And sometimes the only way to get that information is to talk around the problem in a way that e-mails and phone calls just don't do. Furthermore, it's always the same industry people at each meeting—an important point for project continuity and efficient transfer of knowledge; a collaboration is only as good as the institutional memory of the collaborators.

Some face-to-face interactions go much farther than monthly meetings. The MSC issues a standing invitation to its industrial partners to come to Caltech and work in its labs for a while. Yongchun Tang, a senior research chemist at Chevron's La Habra facility and leader of their molecular modeling group, spent several months at the MSC back around the turn of the decade, when the collaboration was first getting off the ground. "We've been so successful in our model-



The model's predictions for an eight-carbon-atom feedstock matched experimental reality pretty well, tracking the rise and fall in concentration of the various hydrocarbon families. The abbreviations are as follows: Cx is everything with x carbon atoms, A is aromatics, nP is unbranched paraffins, IP is one-branched paraffins, and N5 is five-carbon rings.

ing efforts because our scientists have worked at the MSC," Tang says. "It's a critical link. Being in the MSC allowed me to see how its methods could be applied to other problems we are working on, and then to bring the relevant people in. We are way ahead of the other oil companies as a result." (The projects mentioned in this article are just a sampling of the ones Chevron and the MSC have collaborated on over the years.)

Goddard's original notion was to have his theorists talking to the theorists in industry. But there aren't that many out there any more—changes in management strategies have caused most of the companies that maintained showcase research facilities as corporate crown jewels back in the '60s and '70s to abandon or severely curtail them. Like the Cold War-driven space race, they appear to have been a prestige item whose time has passed. It's become much better for the bottom line to farm your research out. "So what we do now is interact with engineers, who are interested in solutions."

Goddard is a theorist, but he's also a pragmatist. He uses these corporate contracts—which now account for about 50 percent of his budget—as a vehicle to advance his theoretical work on computational methods and verify the results. Says Goddard, "We continue to push the frontiers of the methods, but by and large the methods themselves aren't being funded. That's just part of our overhead. If you say in a grant proposal that your aim is to develop a new method, it's very easy for the reviewers to poke holes in you. Someone will say, 'Oh, the assumptions are bad,' or, 'It's not going to be practical,' or, 'You won't get the answers well enough...'" Adds Blanco, "The trouble is, when you propose method development alone, there are so many paths you can take. But our work is all based on using that method to solve a real-life problem. Then the critical issues come out naturally as you work on the problem, and we focus on solving those issues." And if those issues call for a new method, they assign a postdoc to develop it.

"I think it's fair to say that 20 years ago, you could have characterized what we did as 'have gun,

will travel,'" Goddard resumes. "Have method, will simulate.' We looked for whatever problems our method would be useful for. These days, it's 'what's the problem? We'll solve it.'

"I could spend hours just talking about applications, even though that's only about 20 to 30 percent of what we do, because that's the part that most non-theorists are going to see as interesting and important. And perhaps 25 to 30 percent of our theory work I wouldn't even tell to my closest [chemist] friends, because they wouldn't understand why we're doing it. They'd say, 'It's naive to think you can do that problem,' or, 'It's already been done, so why would you want to do it again?' or, 'You've made the wrong assumptions,' or something. Like any kind of research enterprise, some parts of it underwrite the rest of it at any point in time. You just have to keep working on things that will generate the new knowledge that will be useful years from now.

"It's actually evolved into a really good training ground for graduate students and postdocs, because today's world is very different from what it was 20 years ago. Typically then, a person got a PhD in basic science, and joined a company like Exxon or Chevron, and gradually learned the company's technology, gradually learned what was interesting to the company, and eventually got into a position to contribute. These days, if you're not useful on a project from Day One, they're not going to hire you. So our grad students and postdocs have to understand what the technological issues are in order to know how to use their knowledge of basic science. They have to know how to ask the right questions. There's no prescription; you have to learn how to think about problems. And you only learn to think about them by being exposed to them. It's good for them to see how we struggle with defining the problem. Mostly, in science, the important advances are in figuring out what the questions are. And the graduate students get the pleasure of seeing how interested somebody other than their advisor is in their research. Because the advisor is always interested. But to have other people who *really* think it's neat stuff, and to have students see how it's directly related to some important technology, that's good for them to see. So we think it's developed into a good model for education in addition to one for doing research."

Besides the lab's corporate clients, says Siddharth Dasgupta, "We're really a resource for the whole Caltech community. We do a lot of custom software development, and people in a lot of different fields are finding us useful. For example, if some geophysicists want to find out what happens to some kind of rock under tremendous pressure, but they don't have the software to simulate it, they can come here and work with us. They just have to give us a half-page summary of what they want to do, and we give them the resources and training to do it." □