




Theoretical Chemistry Comes Alive

by William A. Goddard III



Top left. The optimized structure for CLT inhibitor (yellow) bound to the active site of the enzyme thermolysin. The tetrahedron represents zinc. The middle part of CLT is held in place by the hydrogen bonds to several residues of the enzyme (the clamp). In the lower region the two parallel hexagons (corresponding to phenylalanine residues) have a special stacking interaction (B. D. Olafson, W. A. Goddard).

Top right. The mayomycin molecule at the predicted bonding site of a CG tetramer of DNA (S. L. Mayo).
Bottom left. Beta silicon nitride, a new ceramic, showing the surface perpendicular to the c axis (M. M. Goodgame, W. A. Goddard).
Bottom right. Zeolite, A, showing the region accessible to solvent within one cavity (blue dots) (M. M. Goodgame, W. A. Goddard)

This article is dedicated to Linus Pauling, who blazed the trail for quantum mechanical contributions to the fundamentals of chemistry, and at 85 and professor emeritus of Caltech, is still active in unfolding the mysteries of nature.

CHEMISTRY IS THE SCIENCE of how electrons conspire to bond atoms into molecules and solids and of how to rearrange these bonds to transform one combination of molecules into a specific new combination. Dealing as it does with the making and breaking of bonds, chemistry is the fundamental science underlying nearly all aspects of modern technology — from drug design, to plastics, to dyes, to catalysts, to high-temperature ceramics and metal alloys, to electronic materials for microelectronics. Unfortunately, despite the amazing progress in all these areas, we often do not really understand why current chemical processes work and consequently cannot predict how to develop new materials with specific properties. Rather, these areas of chemistry continue to develop empirically with clever experimentalists using analogy and intuition to try new procedures with varying conditions until they find a satisfactory solution. However, a revolution is brewing in which this situation should change dramatically, and I would like to provide here some of the flavor of these changes.

The underlying physics governing the motions of electrons in atoms and molecules is quantum mechanics, and when quantum mechanics was developed in the roaring twenties, there was hope that all of chemistry could soon be explained. Indeed, a great deal of progress did occur with theorists such as Linus Pauling at Caltech developing from quantum mechanics simple concepts of bonding that revolutionized the concepts of chemistry. There is, however, an enormous gap between the equations of quantum mechanics and the details of how to transform one chemical into another, and chemistry remained a highly empirical science. (In the same way, discovering that cells are composed of a bunch of chemicals did not explain biology.)

In recent years theorists have learned how to reformulate the basic quantum mechanics into a form where, with the aid of high-speed computers, accurate answers can be obtained for molecules of chemical, biological, and

materials interest. This is most valuable, since the theorist can examine steps of reactions too ephemeral for experimental detection. Thus the theorist can examine the detailed trajectories of each atom during a reaction and can determine the properties of each reaction intermediate. However, this new-found ability to obtain such quantitative information will probably be less important than the ability of the theory to provide new *concepts* that collect together the quantitative results of theory and various experiments and that provide a *qualitative framework* useful for predicting how to modify the properties of a system.

By a conceptual framework I mean a simple picture that allows you to explain everything already known and allows you to predict how to change the system to do something new and neat. For example, one recent problem was to explain some puzzling tribological properties of the diamond surface. (Tribology is the science of friction and wear.)

The Diamond Surface

Diamonds are kind of expensive for bearings, but they do have very low friction (a coefficient of $\mu=0.1$ up to 800°). Recently some researchers at a NASA lab found that when they heated diamond to about 850° C, the coefficient of friction all of a sudden increased dramatically (to $\mu=0.7$). At the same time the observed surface changed irreversibly so that it decomposed easily and led to other special properties. New electron levels appeared at the surface leading to color and a Schottky barrier (surface diode). But then when they exposed it to hydrogen, the friction came back down, the color disappeared, and the surface no longer decomposed so easily. If they heated it up again, all the problems returned as it got above about 850° C.

How can we explain this? Let's think a little bit more about the general properties of diamonds. Diamond is, of course, made of carbon atoms, and carbon has four valence electrons that can be used to make four



Figure 1: Carbon in diamond is bonded to four atoms arranged in a tetrahedron.




Figure 2: Each bond involves two electrons localized on two adjacent atoms but overlapping.




Figure 3: Atoms at the surface can make three good bonds leaving one electron in a dangling bond orbital pointing into the vacuum.




Figure 4: Two pieces of unpolished diamond will have some surface dangling bond orbitals overlapping to bond the pieces together. This leads to high friction.




Figure 5: The surface dangling bond orbital of unpolished diamond leads to an energy in the energy gap (a). This leads to a Schottky barrier (surface diode) as in (b).

strong bonds. Inside the crystal the four atoms bonded to any carbon form a tetrahedron (figure 1). Each such bond involves two electrons, where one electron is localized more on one atom and the second electron more on the other atom, as indicated in figure 2. Carbons at the surface can't possibly make four bonds because one of these atoms would have to be sticking out in a vacuum. So at the surface you know there's going to be some kind of change in the properties of the system. The best a surface carbon can do (as in figure 3) is to make three good bonds with the fourth electron just hanging around doing nothing; it's called a *dangling bond orbital*.

What are the properties of such a system? The surface has these dangling bond orbitals sticking out into the vacuum just aching to make a bond with someone. Now let's put two such pieces of diamond together (as in figure 4) and slide them with respect to each other. The surfaces will not match perfectly, however; at the high points (asperities) orbitals of one surface will overlap the orbitals of the adjacent surface to form a covalent bond. As we try to slide the one piece of diamond, it is necessary to break these covalent bonds, leading to high friction. (We have to push hard to provide energy for stretching the bond as the surfaces slide along, but as a bond is formed, the excess energy gets converted into heat.) Indeed, if you cleave diamond in a vacuum and are quick about it, the pieces will adhere when you put them back together again. You can't wait too long because the surface is quite reactive and even in a good vacuum will quickly react with residual gas molecules until the dangling bond orbitals are mostly used up.

Having these dangling bond orbitals on the surface leads directly to other properties, making the surface very easy to oxidize and also giving it special semiconducting properties. Since each dangling bond orbital has one electron, it's very happy to accept a second one (getting spin-paired), leading to surface charges and diode properties (a Schottky barrier), as in figure 5. This dangling orbital





Figure 6: The top layer of unpolished diamond is unstable.

electron also makes it easy to absorb light, making the surface colored.

Even worse, the cleaved diamond has a tendency to decompose, forming graphitic regions. Why is that? The problem is that the surface carbon is only bonded to three things; but carbon bonded to three atoms likes to be planar (for example, CH_3 is planar). Thus each atom in the top row (say C_1 of figure 6) is yearning to get down to the next plane (to become planar). If the atom (say C_2) in the second row now moves up toward the surface row, the electron on the second-row atom (C_2) previously used to bond to the third-row atom (C_7) can now bond to the dangling bond orbital. Thus the process is to break the bond between the second and third row of atoms (a sigma bond) in order to make bonds between the second and first rows (pi bonds). Doing the same thing for adjacent surface atoms leads eventually to a layer of graphite on the surface. But now the third row of atoms looks just like the original first row. If you keep on doing this, the diamond decomposes — not too good if it's supposed to be for your girlfriend; it probably wouldn't even last the time of a California marriage.

How do we solve this problem? It was actually solved about a millenium ago by ancient diamond cutters, long before theoretical chemists came on the scene. What they did was to cleave the diamond in special kinds of oil. If you put a hydrogen atom (with its own electron) on each surface carbon, it bonds up with the dangling bond orbital to form a two-electron covalent bond. Now the surface carbon is perfectly happy being tetrahedral, and the surface is now perfectly stable. The surface won't react with oxygen, because there are no longer dangling bond orbitals for the oxygen molecule to bond to (singly occupied orbitals are needed to bond to one oxygen and thereby weaken the O-O bond before combustion can occur). Getting back to tribology, we see from figure 7 that there is also very little friction, because with hydrogen on the surface there are no dangling bond orbitals on adjacent surfaces to overlap. There is no chemical bonding between the surfaces — just very weak forces called Van der Waals interactions — and the surface now has the properties of a heavy hydrocarbon (wax). In addition, now that the surface orbitals are paired, they are too stable to play a role in the electronic properties (no room for an extra electron, too stable for

removal of an electron), so now the surface won't lead to color or to diode properties. In the last five years surface scientists have shown that polished diamond has a hydrogen atom at each surface carbon (figure 7) leading



to all these neat properties while unpolished diamond (for example, diamond cleaved in a vacuum) has dangling bond orbitals and other crud on the surface.

Now that we understand all this we can go back and look at those experiments mentioned earlier. Polished diamond with a hydrogen on each surface carbon has low friction. At low temperatures it is way uphill for two hydrogen atoms to break their bonds to the surface and form an H_2 molecule. However at high temperatures the increased entropy of a free H_2 eventually favors desorption. This point occurs at about $850^\circ C$.

Thus at $850^\circ C$ the surface hydrogens come off as H_2 , leaving dangling bond orbitals on the surface and all the properties of the unpolished diamond surface — high friction, color, and a tendency to decompose. Adding hydrogen again forms surface C-H bonds and all is hunky-dory (low friction, stable surfaces).

So now the question is, if we really understand things, how could we modify the system to get *different* properties? A good conceptual understanding allows us to make predictions on what we should do new to change the *properties of the system*. So let's say we want a diamond that would have low friction at $1,500^\circ C$ (maybe diamond turbines for a new high-performance DeLorean). Is there some way to modify this system to make the surface more stable? We have to find something that will bond to each surface carbon, keeping it tetrahedral, but we want to make the bonds much stronger, so that the surface atoms won't leave the surface until higher temperature than hydrogen does. One possibility is to replace each hydrogen with a fluorine atom. The C-F bond is about 0.5 eV stronger than the C-H bond and, more importantly, the F_2 bond is far weaker (about 3eV) than the H_2 bond. Consequently we expect that the fluorinated diamond surface would be much more stable than the hydrogenated one. Simple estimates suggest that the fluorinated sur-

face might be stable up to about $1,500^\circ C$. The fluorinated surface would also lead to low friction (like Teflon) and to other properties like those of polished diamond. (At Seminar Day I said that the experiment had not yet been done; since then I have learned of experiments at Rice University showing that the fluorinated surface retains low friction up to higher temperatures than the normal polished diamond.)

The important point here is to illustrate how a simple concept (about the nature of the surface bonding) can be used to predict a number of different properties of diamond (both polished and unpolished) and how this same idea can be used to design and predict the properties of modified systems. One concept can tie together experiments in completely different fields — from tribology to oxidation resistance, to the semiconducting properties, to the optical properties. They all tie back to this one concept of the tetrahedral bonding of carbon.

Catalysis

Another prime area of chemistry where theoretical chemists are getting into the act is making catalytic reactions selective. As an example, consider the molecule methanol (H_3C-O-H). Adding oxygen to methanol can make carbon dioxide and water, but you can't sell them for very much. What you'd like to do sometimes is remove only two hydrogens and form formaldehyde, as in figure 8. But it's much easier to make car-




Figure 8: Catalysis involves making the less favorable reaction go faster than the more favorable one by appropriately adjusting the barriers. MoO_3 does this to get nearly 100 percent H_2CO , but how does it work?

bon dioxide (making formaldehyde is about one-fourth as favorable). The challenge to the catalytic chemist is to design a system so that the reaction has a small hill to climb for going to the desired product and a much higher hill for going into the deeper valley of the detested reaction.

How does a theorist get ideas about what's going on in such a system? Let's consider molybdenum trioxide powder, which is




Figure 9: Two stable configurations for surface sites in MoO_3 . Mo has six valence electrons and therefore makes six bonds. The left figure (mon-oxo) has one $\text{Mo}=\text{O}$ double bond and four $\text{Mo}-\text{O}$ single bonds (a total of six). The right figure (dual-oxo) has two $\text{Mo}=\text{O}$ double bonds and two $\text{Mo}-\text{O}$ single bonds for a total of six. In both cases the double bonds point into a vacuum so that no broken bonds are necessary.

Figure 11: We are halfway finished but can't finish the reaction if only one Mo center is involved.

Figure 12: If a second dual-oxo site is close by, the second step of the reaction is favorable.

known to catalyze conversion of methanol selectively to the high-energy product (H_2CO) rather than the most stable product (CO_2). What we'd like to understand is what's happening on the surface. What are all the atoms doing, and why does it work that way?

Molybdenum has six valence electrons, and hence the most stable states of molybdenum have six bonds. This leads to two likely configurations on the surface, as indicated in figure 9. These species should be quite stable on the surface since there are no broken bonds.

What chemistry is expected for these surfaces? Janet Allison (PhD 1985) did the quantum mechanical calculations and found that the mon-oxo site (10a) is not very reactive. It's uphill to react with methanol. (Remember, our catalyst has to rip two hydrogen atoms off the methanol.) On the other hand, the reaction is favorable for the dual-oxo site (10b). The reason for such a dramatic difference has to do with the special properties that occur when an oxygen makes a double bond to a metal that already has another double bond. The second $\text{Mo}-\text{O}$ double bond may not seem to be involved; however, this *spectator oxo* bond stabilizes the products resulting from reactions at the other double bond, thereby promoting the reaction. [The real story here is that Tony Rappé (PhD 1981), now professor at Colorado State, and I had discovered this spectator oxo stabilization effect while examining some reactions that occur in solution (homogeneous catalysis), and we guessed that such spectator oxo effects might play a role in reactions in molybdenum





Figure 10: First steps: reacting H_3COH with the surface. The mon-oxo site is not favorable, but for the dual-oxo site the reaction to break the $\text{O}-\text{H}$ bond is favorable because of the spectator oxo group.


crystals. We then looked into the chemistry of molybdates and got interested in the methanol-to-formaldehyde reaction. Of course, in the published research papers, we

start off as if we had started off just to explain reactions on this neat surface.]


This is a start, but we've only done half the reaction; there is still another hydrogen to rip off, and the reaction won't be any good unless we can finish the job. Unfortunately with only one molybdenum center the second step of the reaction is unfavorable (figure 11). However, we find that the second step would be favorable if there were a second dioxo site close to the first one to pull off the second



hydrogen and make the final product (figure 12). Thus we concluded that a properly configured site with two dual-oxo units could be the catalytically active surface site for giving CH_3OH this one-two punch. The next question is whether the real surface can have such a configuration.



At this point in the research, Janet and I examined the bulk structure of molybdenum trioxide from x-ray structure studies and found that one surface, (010), of the crystal has exactly the configuration and the properties that we wanted (figure 13). Looking sideways at the surface there's a whole row of dual-oxo molybdenum sites. This is the most stable surface plane (since no chemical bonds are broken in making the surface), and we found that this surface leads to a plausible catalytic cycle. At this point we ran across an experimental paper that had just come from France. These researchers had examined the catalytic reactions on itty-bitty molybdenum trioxide surfaces and found that one surface — precisely the (010) surface we had deduced — was responsible for the selective catalysis of H_3COH to H_2CO . We were really elated and rushed to finish the paper; we submitted



it a month later, saying that the French experiments proved our theory.

We then waited patiently for three months to get comments back from the referees. One referee said, Hey, this is really great work. It's good to see the theorists are finally doing something useful. Publish it. The other guy said, Well, it's interesting, but it can't be right. Some new experiments in an American industrial lab show that this surface is unreactive. We immediately located this new experimental paper, and it turned out that they hadn't actually done the chemistry. In a high vacuum system they had exposed the (010) surface of a single crystal of molybdenum trioxide to the CH_3COH , and since they didn't see any change in the spectroscopic properties of the surface, concluded that it's probably not reactive. But what they neglected to do is to expose the surface to oxygen. The French researchers had shown that if you don't have an atmosphere of oxygen on the (010) surface, you get no chemistry whatsoever. Presumably, without oxygen the catalyst loses some of its surface oxygen and hence loses the dioxo units required for the chemistry. So now my friends at this industrial lab are busy doing experiments adding oxygen to the surface. I should emphasize here that the theoretical results do not *prove* that the dual-oxo sites on the (010) surface do the chemistry. We have considered only the chemistry expected for these stable surface species. There could be an unstable species on this or other surfaces which is special to catalytic conditions and which would do the observed chemistry more rapidly. Now that there is a specific model (with a new principle — spectator oxo stabilization) for the surface configuration responsible for the catalysis, new special experiments will be designed to make specific tests of this model. As these tests proceed, the theorists will learn from the experiments and will examine various details more carefully. The result from

the theorists and experimentalists working together and separately will be a new level of understanding which will eventually lead to the knowledge needed to design new catalytic processes.

Simulation

The above examples focused on concepts about what the electrons are doing at surfaces and how to understand specific properties of various surface electronic states. There will be a flourishing of such theoretical activities in coming years; however, I believe an even more dramatic impact will come out of a related area of theory — *computer simulation of materials processes*. The idea here is to take the results of quantum chemical calculations on clusters of atoms (up to 10 or 20 atoms), as was used to get the numbers in the above sections, and to extract from these results an analytic description of energies and forces in terms of two-body, three-body, and four-body interactions. With such a description computer programs have been developed that allow the forces and dynamics of, say, 5,000 atoms representing the surface plus reacting atoms to be rapidly calculated on minicomputers (like the DEC VAX). With proper software the results of these calculations can be displayed in real time on a graphics system so that the scientist/engineer can actually "see" the reaction as it proceeds. With proper graphics equipment (such as the Evans and Sutherland PS 300), the user can interactively rotate and zoom the system to see the details in specific regions of the system and can pop the system into stereo to see the dynamics in 3-D. This overall process — *Computer Aided Materials Simulation (CAMS)* — will, I believe, have a significant impact upon the areas of drug design, chemistry (catalysis, synthesis, polymers), and materials research (ceramics, semiconductors, metallic alloys) even more dramatic than what CAD/CAM has done for the engineering design and processing communities.

At the moment biological systems are the only ones for which we have a good enough understanding of the forces so that we can reliably represent the quantum mechanics in terms of force field functions. The system illustrated on the cover (and at top left on page 2) is an enzyme, thermolysin, which is a model for an enzyme (angiotensin-converting enzyme) that you want to block in fighting hypertension. It's also resistant to high temperatures and so is a good model for research

Figure 13: The (010) surface of MoO_3 . This surface shows the adjacent dual-oxo sites needed for the steps in 10 and 12.

on commercial biocatalysis. Thermolysin has more than 3,000 atoms, and it's selective for breaking peptide bonds connecting hydrophobic (water-hating) amino acids.

Including a layer of water and some salt to mimic thermolysin in solution, postdoctoral fellow Barry Olafson and I used CAMS to calculate the forces on all 3,500 atoms as a function of time for simulations of the dynamics at various temperatures and to predict the optimum structure. In some of these systems experimental crystal structures were available for comparisons, and the excellent agreement confirmed the overall adequacy of the force fields. In the cover illustration the blue dots show the surface of the enzyme as seen by the solvent (water). The tetrahedron represents a zinc atom, and the entire cleft region represents the active site for this enzyme to bind to the substrate (the molecule it will cleave). A drug molecule to block the enzyme might be a molecule that would bind so strongly that it could not be displaced. Such an inhibitor (shown in yellow) was known from studies at Merck to be effective, but the structure was predicted theoretically without knowledge of the experimental structure. (Later comparison showed good agreement.)

The point is that theory is now in a position to give a credible prediction of structure even when there is no substantiating experimental evidence. Theory can also predict the interaction energies, which allows one to analyze why something (say, an inhibitor) works the way it does and then modify it on the computer to design an even better drug.

Probably the real payoff for theory in drug design over the next few years will be in understanding how various kinds of molecules bind to nucleic acids. This process plays a critical role in determining which genes are expressed and how fast, determining, for example, why the DNA in your earlobe makes earlobe cells and not brain cells. Olafson and I are now working on a regulatory protein binding to DNA but haven't yet optimized the structure. Using theory we're trying to calculate how the interactions work and predict how and where the protein wraps around the DNA. In a way we're in a race with experimentalists who are trying to cocrystallize this system so that they can do x-ray diffraction studies to get structure. (The experimentalists never believe that the theory can provide real predictions unless the theory was done before the experiment.)

As an alternative attack on this overall problem, grad student Steve Mayo has designed a molecule (top right on page 2 — we call it mayomycin) to be selective for binding to sections of DNA that are rich in CG base pairs. The idea here is that if we can design a drug to recognize only a very specific sequence of base pairs, then we might be able to control the expression of a particular gene of a particular organism.

We're also working on understanding the forces for semiconductor surfaces so we can examine new synthetic techniques for making semiconductor devices (for example, the use of molecular beam epitaxy to make heterojunctions and superlattices). We are also examining catalysis (such as in zeolites, bottom right on page 2) and ceramic surfaces to understand what happens when one ceramic rubs against another (tribology). The figure at bottom left on page 2, made by postdoctoral fellow Marv Goodgame, shows a surface for $\beta\text{-Si}_3\text{N}_4$, a new ceramic that is being developed for high-temperature gas turbines and adiabatic diesel engines.

A new project funded by the National Science Foundation will involve five Caltech faculty members interested in materials (Bill Johnson and Thad Vreeland from materials science, Marc Nicolet from applied physics, Tom Tombrello from physics, and myself from chemistry) who will carry out joint experimental and simulation studies to examine kinetic processes in materials synthesis. A unique aspect of this *Caltech Materials Center* (CMC) will be the integral use of simulation techniques in the various experimental programs. Because it is important for such a materials development program to have strong coupling with industrial research and development organizations, the group will probably involve a few industrial sponsors committed to a strong interaction with the CMC.

Various portions of the research reported were supported by the Department of Energy (Energy Conversion and Utilization Technology); the National Science Foundation (Division of Materials Research); the National Science Foundation (Chemistry Program); the Petroleum Research Fund of the American Chemical Society; the Office of Naval Research; and the Army Research Office. The software used for the biological simulation is BIOGRAF, written by Mayo, Olafson, and Goddard, and the hardware is an Evans and Sutherland PS-300/DEC VAX 11/780. □