

Hill, proved that result. Giray Okten of the University of Alaska-Fairbanks then spoke on some new constructions of quasirandom sequences that combine traditional quasirandom sequences with pseudorandom sequences. This series of talks addressed some of the mathematical difficulties behind providing parallel streams of quasirandom numbers. During the extensive discussion on this topic, new results on how to use full-period pseudorandom numbers to provide many related quasirandom streams were presented. This seems like a promising approach that needs more study.

Next, Greg Astfalk, chief scientist of Hewlett Packard's High-Performance System's Division, shared his views on the future of high-performance computing. He made two especially telling points:

- Vector architectural features (such as flat memory, low latency, extremely high sustainable-memory bandwidth, multiple pipes to memory,

hardware support for gather/scatter, efficient single-word access, vector instructions—all the other parts of the vector architecture that matter more than the vector instructions themselves) are dead.

- Almost all vendors seem to be converging on clusters of symmetric multiprocessing systems for their high-end products. These clusters will be based on the same SMPs that make up their server and high-end workstation markets: success will ride on maximally leveraging commodity technology.

The SPRNG tool

The final two speakers spoke about SPRNG and its suitability for this community. NCSA's Ashok Srinivasan described SPRNG in detail, while I presented SPRNG in terms of the design decisions made in its implementation. This led into a broad discussion on whether SPRNG was the right tool for the present and what things an improved

SPRNG should offer in the future. Ashok and I concluded that SPRNG is well-designed for current applications, but that new generators should be designed to better meet the demands of ASCI-class applications. In addition, applications-based testing, coupled with a Web-accessible test site, would serve both the theoretical and applications community. Quasirandom numbers should be made available for parallel and distributed computing, but first a better mathematical understanding of key issues will clearly be required.

Workshop sponsors included the University of Southern Mississippi's Program in Scientific Computing, the Center of Higher Learning and the Programming Environment and Training Program at Stennis Space Center, the NCSA, the University of Illinois at Urbana-Champaign, and Darpa.

Michael Mascagni is an associate professor of mathematics at the University of Southern Mississippi; michael.mascagni@usm.edu.

QUANTUM CATALYSIS

Donald G. Truhlar

The recent National Meeting of the American Chemical Society in Dallas included a three-day (March 30–April 1) international symposium on Transition State Modeling in Computational Catalysis, sponsored by the Division of Computers in Chemistry. Forty papers by computational chem-

ists and physicists from 13 countries covered the whole gamut of computational issues involved in current models of catalytic reactions. All kinds of catalysis were represented: acid-base, homogeneous organic and organometallic, metal surfaces, metal oxides and zeolites, and enzymes.

Catalytic reactions are widely used for polymerization and for producing fine chemicals and pharmaceutical molecules. The computational methodologies applied to the wide variety of catalytic processes share a number of common features, providing yet another example of the phenomenon

whereby computational science provides a unifying element that promotes cross-disciplinary interactions. Catalytic reactions will become increasingly important because of the new demands placed on the chemical industry for enantiospecific syntheses and environmentally friendly reaction schemes.

Understanding catalytic reactions

Computational chemistry has made great strides in modeling chemical reactions, due partly to faster computers but even more to improved algorithms. Most reactions of industrial or biological importance are promoted by catalysts, which makes them intrinsically more complex than the simple reactions on which most theoretical progress has been made so far. Because of this complexity, catalytic reactions are often poorly understood. Consequently, chemical reaction conditions for reactions that underlie multibillion-dollar industries are often adjusted on a trial-and-error basis rather than on fundamental understanding. Furthermore, drug design and biomedical engineering endeavors often proceed under equally clueless circumstances. But this situation is changing rapidly due to advances in computational science.

Catalysts are molecular-scale engines, and—like other atomic and molecular phenomena—their ultimate understanding rests on quantum mechanics. Thus, the symposium focused on computational quantum mechanics and, in particular, on the quantum mechanics of transition states. Transition states are the dynamical bottlenecks gating the flow of quantum probability density from the configurations we label as reactants to those that we label as products. We can identify them as the configurations along the reaction path that have the highest free energy of ac-

tivation. Although the molecular transit time through the transition state is short, $O(10)$ femtoseconds, its properties determine the rate of reaction as well as its selectivity and specificity. But a transition state by its very nature has partial bonds, which are much harder to treat in a computationally reliable fashion than the full bonds of stable molecules. This is sometimes called the open-shell problem, referring to incomplete occupancy of one or more shells of electronic orbits in the molecular complex.

Transition states are high-energy species, and one of the most critical pa-

◆

***Computational
chemistry has made
great strides in modeling
chemical reactions, due
partly to faster
computers but even
more to improved
algorithms.***

◆

rameters to calculate is their internal energy relative to reactants. Because the transition states have open-shell electronic structures, whereas the reactants often have closed-shell structures, it is hard to treat the two types of systems in a balanced way, and this leads to systematic errors. Traditional methods for overcoming this problem have a computational effort that scales as a high power, for example, N^5 or N^7 , of the number N of electrons in the system. Because a practical catalyst for

polyethylene production might contain 150 electrons, and even a small protein (a protein that serves as a catalyst is called an enzyme and is the most common form of biochemical catalyst) might contain many thousands of electrons, the problem is clear.

Divide-and-conquer methods

A recent development in computational chemistry is the use of divide-and-conquer molecular orbital methods that decompose a molecule into small overlapping subsystems. This technique allows us to devise algorithms for molecular orbital calculations whose computational resource requirements scale linearly with the size of the system. A few years ago, chemists first performed full self-consistent-field molecular orbital calculations for whole proteins. At the Dallas symposium, Kennie Merz of Pennsylvania State reported a calculation in which he included not only an entire protein with 69 amino acid residues but also a large number of surrounding solvent molecules, for a total of 5,000 atoms and over 10^4 electrons. These calculations led to the interesting result that 10 of the amino acids each transfer about 0.2 electrons to the solvent. These protein-water charge-transfer interactions have been uniformly neglected in the force fields for molecular dynamics and Monte Carlo simulations of proteins. Thus, if Merz's findings are substantiated, they should profoundly influence force fields for condensed-phase molecular-dynamics simulations.

An example of a computational science problem that challenges workers in the field of transition-state modeling is nonlinear optimization. To the first approximation, transition states are stationary points (points where the gradient vanishes) of a multidimensional (50 to thousands of dimensions) potential energy function. However, they are not

extrema (local minima or maxima) but rather first-order saddle points. At a first-order saddle point, the Hessian has one negative eigenvalue with all the others positive. A special twist on the problem is that the Hessian is often far too expensive to calculate. Thus, several papers were concerned with numerical methods for finding saddle points of a multidimensional function using only its value and gradient, without any knowledge of Hessians. Especially stimulating discussions of this problem were presented by Eric Sandre of Cambridge, England, and Jens Nrskov of Lyngby, Denmark.

In related work, Hans Martin Senn of the Swiss Federal Institute of Technology in Zurich presented a scheme based on driving a system over a barrier by molecular dynamics with a continuously moving, one-dimensional constraint. In a variation on the theme, Joachim Sauer of Berlin discussed a dual-level approach in which low-level Hessians (which are affordable) are combined with Hessian updates based on high-level gradients.

Zeolites, quantum enzymology, computer-aided catalyst design

Scott Auerbach of the University of Massachusetts presented the fascinating challenge of modeling diffusion transition states in zeolites. Zeolites are microporous, crystalline silicate solids with regularly arranged molecule-sized cavities containing catalytic cation sites. Under certain circumstances, the rate-limiting step of the overall catalytic reaction is diffusion of molecules through the windows into and out of the cavities. In solving this problem, Auerbach combined fast-multipole and Ewald methods for electrostatics, the kinetic Monte Carlo method with infrequent event sampling for discrete jumps between binding sites on the lattice, and

the transition-state theory of activated diffusion. For future progress, he plans to consider the transport equations for diffusive flows. The reason for this is the increasing interest in zeolite thin films, for applications where a high gas pressure is imposed on one side of the film and a vacuum is pulled on the other side. This sets up a steady-state flow, which he will model by a kinetic Monte Carlo method. The pressure-dependent rates of penetrating and exiting the film on both sides will be obtained from the intracrystalline transport rates discussed above. Interestingly, anisotropic intracrystalline dif-

fusionally promoted reaction steps with and without isotopic substitution. Then you use computational transition-state theory to extract the structure of the transition-state complex between the enzyme and the reactant. Next, you calculate the electrostatic potential of the reactant in its transition-state configuration. This is the configuration that is most stabilized by the enzyme. Finally, you attempt to design drugs that have a similar electrostatic potential to the transition state complex. Such transition-state analogs should show potent binding to the active site of the enzyme.

Quantum enzymology in this form is a case of computer-aided design of catalytic inhibitors. An even more general approach is computer-aided catalyst design. This subject is blossoming especially rapidly in the area of organometallic catalysts, as discussed by Senn, Karsten Krogh-Jespersen of Rutgers, Don Lauffer of Phillips Petroleum, and Tom Ziegler of the University of Calgary. As computational chemists use the techniques of quantum mechanics to understand the most intimate details of the fleeting bottleneck configurations of the molecular-scale paths taken by molecules in complex catalytic reactions, the use of computers to design and fine-tune catalysts in all branches of industry becomes inevitable. The age of quantum catalysis has begun, and it is entering a period where we can expect exponential growth for years to come. ♦

◆

The age of quantum catalysis has begun, and it is entering a period where we can expect exponential growth for years to come.

◆

fusion can couple to the flow to create anomalous diffusion, which breaks spatial symmetry. This project provides an excellent example of combining a variety of scientific computation techniques for the solution of a single practical problem.

Vern Schramm of the Albert Einstein College of Medicine of Yeshiva University presented a vision of quantum enzymology that he believes will become increasingly visible in the very near future. In his approach, you begin by measuring the relative rates of en-

Donald T. Trublar is a professor of chemistry at the University of Minnesota; trublar@umn.edu; <http://www.chem.umn.edu/~trublar>.