CRAY CHANNELS

FALL 1991 · A CRAY RESEARCH, INC., PUBLICATION

Chemistry and Chemical Engineering

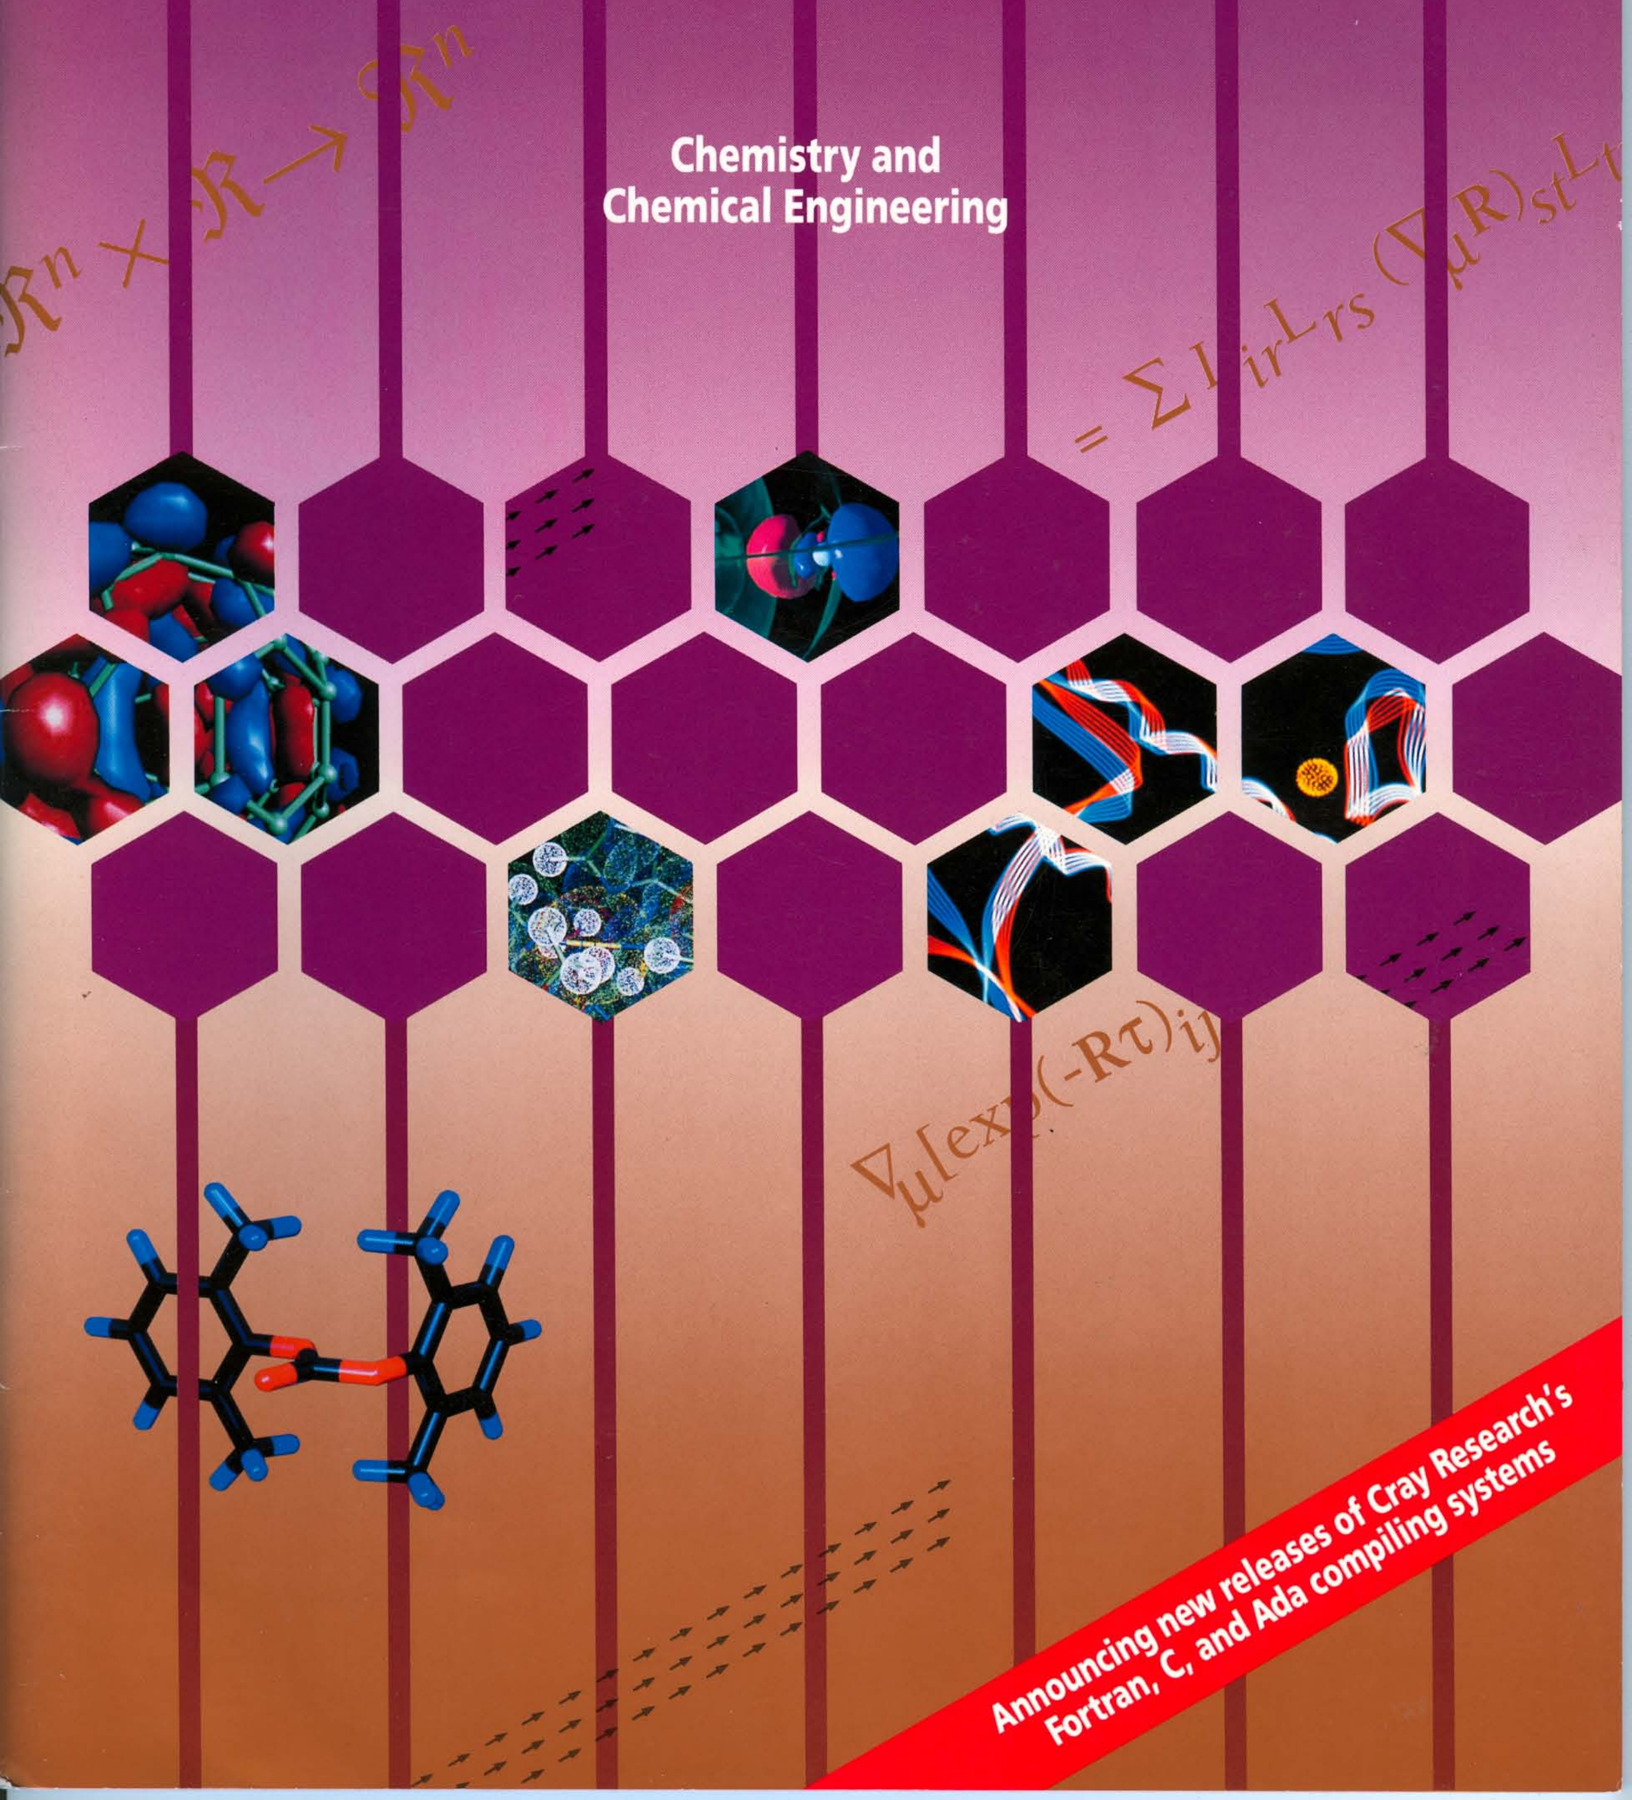

Announcing new releases of Cray Research's Fortran, C, and Ada compiling systems

CRAY CHANNELS

In this issue

Every attempt to employ mathematical methods in the study of chemical questions must be considered profoundly irrational and contrary to the spirit of chemistry. If mathematical analysis should ever hold a prominent place in chemistry — an aberration which is happily almost impossible — it would occasion a rapid and widespread degeneration of that science."

— A. Comte, *Philosophie Positive,* 1830

Contrary to the sentiments of the French mathematician and philosopher Auguste Comte, computational methods have acquired a prominent place in chemistry, vastly enriching the science. Chemists today use large-scale computational methods to reveal the properties and structures of atomic and molecular systems in detail. As a result, chemical research and chemistry-related product development are more efficient today than ever before.

In this issue of CRAY CHANNELS we focus on the practical benefits of supercomputing in chemistry and chemical engineering. We also describe enhancements included in the newest releases of Cray Research's Fortran, C, and Ada compiling systems. Our regular departments cover strategic partnerships between Cray Research and a number of research institutes and describe the computerized search for a better golf club.

To take full advantage of large-scale computational methods, chemists need easy-to-use and efficient software in addition to high-performance hardware. To meet this need, Cray Research developed the UniChem software environment. UniChem lets chemists import or build molecular structures, select appropriate methods, control computations, and visualize results all through a common interface. The speed and ease-of-use of Cray Research hardware and the UniChem software environment provide a significant competitive advantage to computational chemists in the pharmaceutical, chemical, petroleum, and other chemistry-related industries.

Features

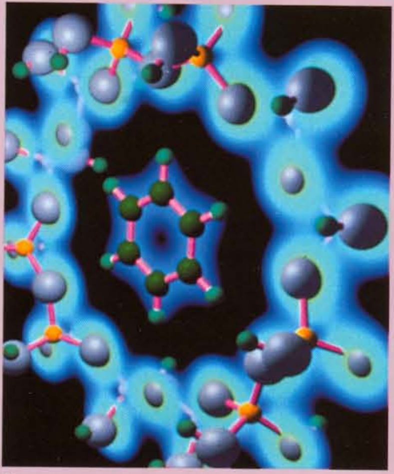

6

Departments

CRAY CHANNELS is a quarterly publication of the Cray Research, Inc., Marketing Communications Department, Tina M. Bonetti, Director. It is intended for users of Cray Research computer systems and others interested in the company and its products. Please mail feature story ideas, news items, and Gallery submissions to CRAY CHANNELS at Cray Research, Inc., 1440 Northland Drive, Mendota Heights, Minnesota 55120. Subscription inquiries and address changes should be sent Attention Dept. D.

Volume 13, Number 3

Editorial staff
Ken Jopp, editor
Mary Miller, associate editor
Vicky Frank, contributing editor

Design and production
Barbara Cahlander, typesetter, production artist
Eric Hanson, graphic designer
James Morgan, art director
Cynthia Rykken, graphic designer

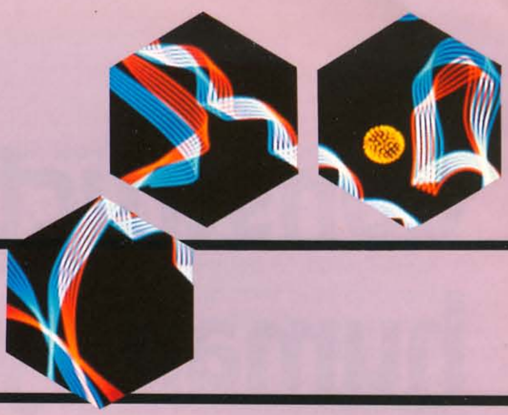

Construction of a human synovial phospholipase A_2 model

Robert B. Hermann and Ernst Dow, Lilly Research Laboratories, Eli Lilly and Company, Indianapolis, Indiana
Chemists use computational methods to determine an enzyme structure — a first step toward designing enzyme inhibitors to fight disease.

Computational quantum chemistry in catalysis research

Aileen E. Alvarado-Swaisgood, Exploratory & Catalysis Research, Amoco Oil Research & Development, Naperville, Illinois
Supercomputing helps petroleum industry researchers better understand the chemistry of catalysis.

Computational chemistry in parallel

John E. Mertz and Jan W. Andzelm, Cray Research, Inc.
New parallel algorithms improve the performance of computational chemistry software.

UniChem: a computational chemistry environment

Jeffrey F. Gaw, Cray Research, Inc.
Cray Research's UniChem software environment gives chemists several modeling options through a single, easy-to-use interface.

Improving chemical process design with supercomputers

Alan B. Coon, Union Carbide Technical Center, South Charleston, West Virginia
Christopher H. Goheen and Mark A. Stadtherr, University of Illinois, Urbana, Illinois
A new reordering algorithm streamlines equation-based flowsheeting simulation.

Supercomputing at Du Pont: a chemical engineering perspective

Richard J. Thomas, E. I. Du Pont de Nemours and Company, Wilmington, Delaware
After four years of experience, chemical engineers at Du Pont evaluate the effectiveness of the company's Cray Research system for process simulation.

Computational brain dynamics: visualization of activity flow in the cerebral cortex

Robert E. Wyatt and Jesse W. Driver, University of Texas, Austin, Texas
Large-scale simulations of neural activity patterns contribute to the understanding of brain dynamics.

Announcing new releases of the Cray Research compiling systems

Building on its reputation in supercomputer compiling performance, Cray Research offers new releases of its compiling systems: CF77 5.0, Cray Standard C Compiler 3.0, and Cray Ada 2.0.

Corporate register

Applications update

User news

Gallery

Construction of a human synovial phospholipase A_2 model

Robert B. Hermann and Ernst Dow
Lilly Research Laboratories, Eli Lilly and Company, Indianapolis, Indiana

The design of inhibitors for enzymes that are known to play or suspected of playing a role in disease states is of therapeutic interest. It is expected that a small molecule designed to fit in a complementary fashion certain aspects of an enzyme structure should be bound strongly to the enzyme in solution, and, provided some other conditions are met, should affect the function of the enzyme in vivo. For this reason it is desirable to know the three-dimensional or tertiary structure of such enzymes. Recently, with the advent of powerful supercomputers such as Eli Lilly's CRAY-2 computer system, computational chemistry has been useful in determining these structures, which then are used in directing organic synthesis. In this manner, computational chemistry now can aid in the inhibitor design process.

Prediction of the tertiary structure of an enzyme from its amino acid residue sequence remains one of the outstanding unsolved problems of chemistry. However, one can, through homology modeling, sometimes determine the three-dimensional structure of an enzyme whose sequence is known, provided it belongs to a family of enzymes in which the tertiary structure of at least one member has been determined. It would be desirable to infer the structure of such an enzyme, through homology modeling, at least accurately enough to use in the design of specific enzyme inhibitors.

An example of such an interesting enzyme is human synovial phospholipase A_2^1, originally isolated from human rheumatoid synovial fluid. Phospholipases generally are ubiquitous, occurring intra- and extracellularly. Their substrates are the phospholipids, which include the building blocks of the cell wall. In general, all the PLA_2s produce fatty acids by hydrolysis of the sn-2 fatty acyl ester bond of the phosphoglycerides. Synovial PLA_2 is implicated in arthritis because it produces arachadonic acid directly in the synovial fluid by hydrolysis of the corresponding phosphoglyceride, and arachadonic acid is the first of several intermediates in the formation of mediators of arthritis. Here we build a workable model which we are using in our inhibitor design efforts. The model is constructed by combining the known features of bovine pancreatic PLA_2 and rattlesnake venom (*C. atrox*) PLA_2, whose structures have been determined by x-ray diffraction and the known sequence of human synovial PLA_2. We augment this structural information with available information pertaining to the binding of inhibitors[2,3] and calculations regarding both the nature of bound waters and the accepted mechanism of action. Lilly's CRAY-2 system was invaluable in completing the necessary molecular dynamics calculations. As a partial test of its validity, the model has been used successfully as the starting model for the x-ray determination of the molecular structure of human synovial PLA_2 by the molecular replacement method.[4]

Homology modeling

Building a protein by homology involves replacement of the amino acid residues of the old protein with the residues of the new one, thus keeping the backbone structure rather similar. To build the model, the *C. atrox* structure was chosen as a starting backbone configuration[5] because it already has 46 percent homology (residue identity) with the human enzyme, and like the human enzyme, it is a Type II PLA_2. The other available PLA_2, the bovine structure, is Type I. The main structural differences between the types (Figure 1) are the presence of a seven residue terminal flap in Type II and the presence of a cystine 11 to cystine 77

Figure 1. Bovine PLA_2 (Type I) in red and *C. atrox* PLA_2 (Type II) in blue. White region indicates overlap. Orange sphere indicates the position of the active site calcium in the bovine PLA_2.

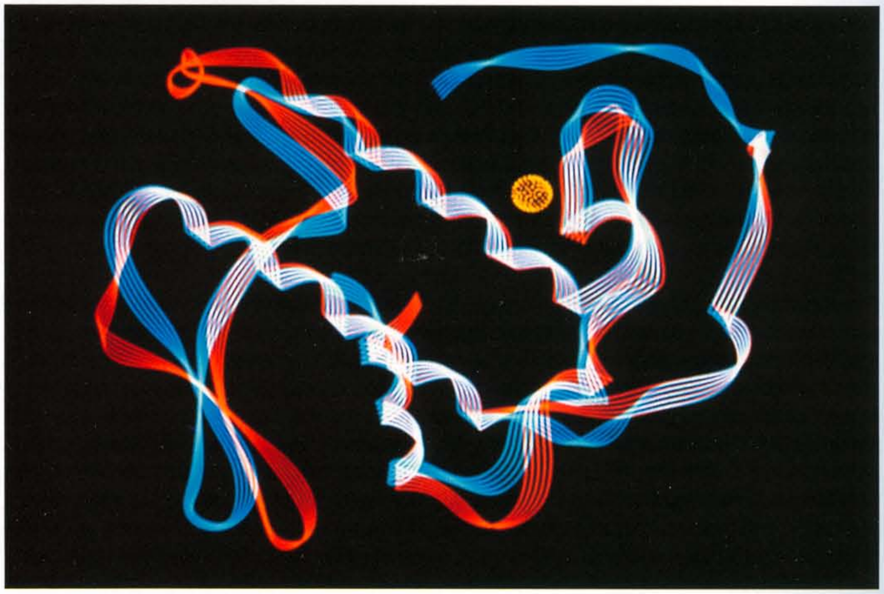

```
                5         10          15         20        25          30        35
HUM SYN   N L V N F H R M I  K - L T T G K E A A L S Y G F Y G C H C G V G G R G
C. ATROX  S L V Q F E T L I  M - K I  A G R S G L L WY S A Y G C Y C G WG G H G
BOVINE    A L WQ F N G M I  K C K I  P S S E P L L D F N N Y G C Y C G L G G S G

                40        45          50         55         60         65        70
HUM SYN   S P K D A T D R C C V T H D C C Y K R L E K - R G C G - - - - T K F
C. ATROX  L P Q D A T D R C C F V H D C C Y G K A T - - - D C N - - - - P K T
BOVINE    T P V D D L D R C C Q T H D N C Y K Q A K K L D S C K V L V D N P Y T

                75        80          85         90         95        100       105
HUM SYN   L S Y K F S N S G S R I  T C - A K Q D S C R S Q L C E C D K A A A T C
C. ATROX  V S Y T Y S E E N G E I  I C - G G D D P C G T Q I  C E C D K A A A I  C
BOVINE    N N Y S Y S C S N N E I  T C S S E N N A C E A F I  C N C D R N A A I  C

                110       115            120       125       130
HUM SYN   F A R N K T T Y N - K K Y Q Y Y S N K H C R G S T P R C
C. ATROX  F R D N I  P S Y D N K - Y WL F P P K D C R E E P E P C
BOVINE    F S K - - V P Y N K E H K N L D K K - N C
```

Figure 2. Sequences for PLA_2s discussed in the text, showing evolutionary correspondences. Blanks indicate deletions corresponding to deletions in the genetic code producing the sequences. The above proteins each consist of a single chain 125, 122, and 123 residues long, respectively. The numbering has been changed slightly from Ref. 6 to be consistent with the more familiar numbering.

crosslink in Type I. The one-letter sequences (Figure 2), show where these differences can be seen.[6]

It is believed that all PLA_2s contain calcium, which is necessary for activity. Since the *C. atrox* PLA_2 structure as experimentally determined by x-ray diffraction had no calcium (probably due to the particular acidity used for crystal growth), it is necessary to add this ion to the model. Furthermore, the environment of the ion in the model must be consistent with the accepted mechanism of enzyme hydrolysis.

The crystal structure of the *C. atrox* PLA_2 as given by Sigler et. al. (Brookhaven Data Bank entry 1PP2) was used as the starting structure for the modeling. Quanta/CHARMm software was used to make the residue substitutions as indicated in Figure 2. There was also one insertion, namely two residues after the *C. atrox* residue 56, on an outer loop of the protein.

The inserted residues together with three on either side were minimized with CHARMm. The structure was solvated within a 10-Angstrom sphere of water, and after 100 picosecond (psec) of dynamics using CHARMm there was no appreciable change in the backbone structure. The seven disulfide bonds impose a large degree of rigidity.

Modification of the calcium binding segment

The conformation of the segment 29-34 had to be modeled to look more like the bovine phospholipase structure, which has calcium, than like the published *C. atrox* structure (Brookhaven Data Bank entry 1PP2), which does not. Considerations of the mechanism of action and known modes of binding of certain inhibitors[2,3] lead us to this conclusion. According to the mechanism of Drenth et al., and the position of some inhibitors in the active site of several enzymes, there is an "anionic binding site," similar to the serine proteases, that binds the anionic oxygen formed when the carbonyl carbon of the substrate is attacked by the nucleophilic water molecule. This anionic binding site is a result of a special orientation of the glycine 30, a residue strictly conserved in all PLA_2s. This residue must be aligned so that the amide hydrogen atom bonds with the anion when present. Figure 3 shows a calculation of the position of the substrate in bovine PLA_2, having the substrate ester carbonyl oxygen in the anionic binding site. This orientation of residue 30 is different in *C. atrox* PLA_2, where there is no calcium. After placing the calcium ion in the structure, molecular dynamics with AMBER 3.0A in water was used to adjust this region of the model. The N-H of residue 30 remained directed toward the anionic binding site.

Adjustment of the side chains

Because the tertiary structure of the two available PLA_2s in certain regions are geometrically similar, it was felt that the same regions of our model might be similar to those corresponding regions of the available PLA_2s. For example, the sulfur-sulfur bonds are spatially related in each enzyme in a similar manner. Therefore, to remove the energetically unfavorable side chain orientations and preserve aspects of the structure deemed favorable, a limited form of dynamics was carried out whereby all cysteine, glycine, and alanine residues, and aspartic acid 99 and the calcium ion were kept fixed and everything else was varied. This was done on the solvated enzyme. Forty psec of dynamics at 300 K then yielded our structure.

Locating the binding sites

For the PLA_2s, the active site constitutes a large, chiefly hydrophobic binding region. To further improve the model as a working model for inhibitor design, it is desirable to locate within this region the nonhydrophobic binding sites. From the hypothesized mode of binding of substrates consistent with the accepted mechanism, and from examination of the positions of the water molecules in the bovine PLA_2,

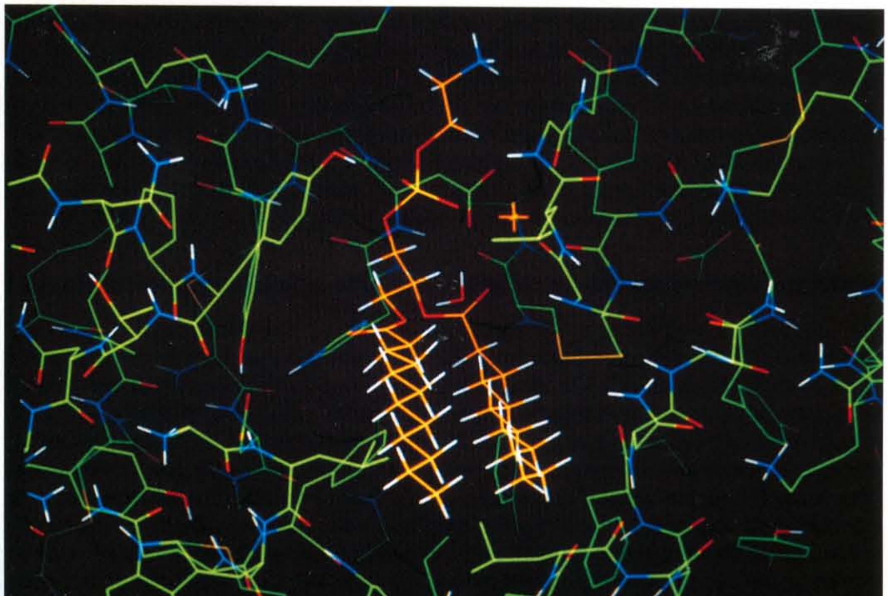

Figure 3. Position of substrate in bovine PLA_2 from energy minimizations and other considerations (see text). Calcium binding segment is shown in orange.

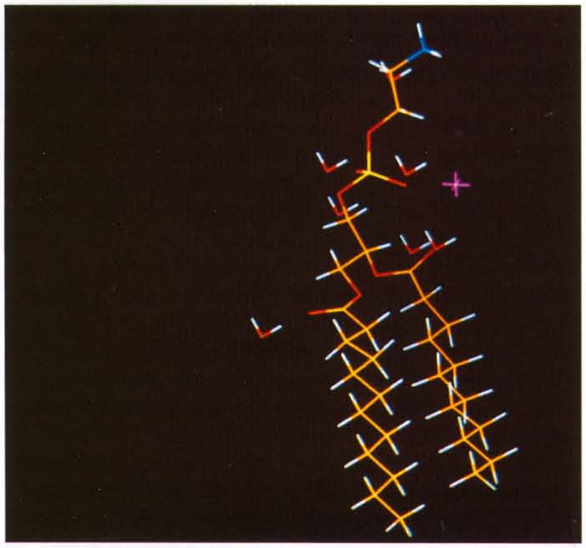

Figure 4. Overlap of the substrate as it sits in the bovine active-site (see Figure 3) with the x-ray-determined active site waters.

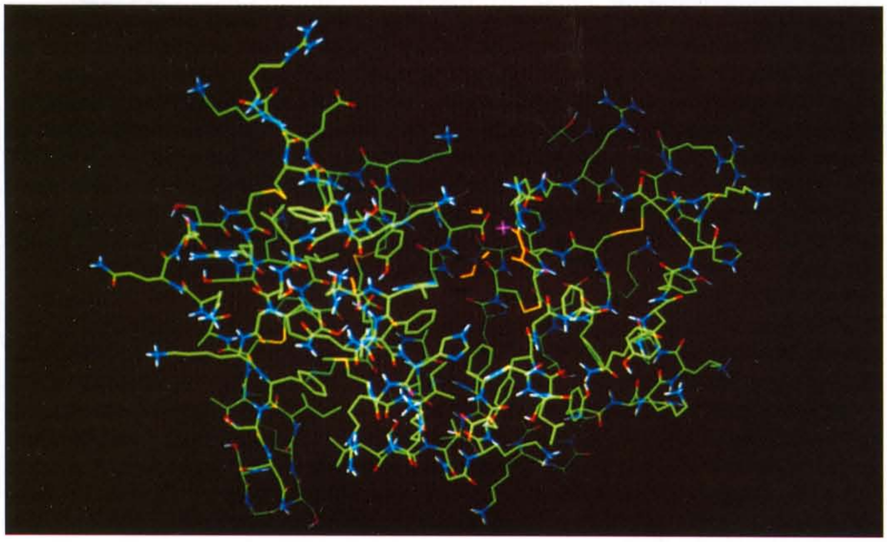

Figure 5. Final human synovial PLA2 model, including proposed binding sites.

it is possible to make a reasonable inference of these sites. The overlap of the substrate as it is presumed to sit in the bovine PLA$_2$ enzyme pocket with the x-ray determined water actually found in the bovine PLA$_2$ pocket is shown in Figure 4.

Of the sites marked by these waters, at least three possible significant sites should exist:

- ☐ The nucleophilic water in the accepted mechanism of PLA$_2$ action that is bound between the histidine 48 and the aspartic acid 49 oxygen.
- ☐ The anionic binding site — a site that holds the carbonyl oxygen of the substrate ester and is located as indicated by the water bound to the N-H group of glycine 30. This water is electrostatically bound to the calcium.
- ☐ A second water also electrostatically bound to the calcium. This last position is occupied by the phosphate oxygen in the bound substrate (Figure 3).

The determination of reasonable methods to find binding sites is an area of active research.$^{7-10}$ Here, in our work, molecular dynamics (at 300 K) was used to locate these binding sites as follows. Keeping the coordinates of the structure fixed and allowing only the waters to move gives rise to a distribution of water molecule migrations. After equilibrium is established, we look for the water molecules that move the least. We suggest that the positions of these water molecules are good candidates for significant binding sites. Water molecules that are immobilized are not only usually tightly bound, but also, if liberated through replacement by a ligand molecule, would be expected to contribute to the entropic part of the free energy of the binding process.

By subjecting the several enzyme structures in this manner to 40 psec of molecular dynamics while holding the structures fixed and allowing only the water to move, an average was taken over the last 4 psec and the root-mean-square positional deviation of the waters was recorded. A period of 4 psec was chosen for averaging rather arbitrarily because hydrogen bonds in water are made and broken on an average of 4 psec. Next, all waters that moved less than .2 Angstroms were selected from these.

Following this procedure for both the bovine phospholipase and the human synovial model, the three significant sites mentioned above were found in each case. A fourth immobilized water was found, namely the catalytic "network" water molecule, bound by the amino terminal NH_3^+ and the aspartic acid 99 oxygen. Several other waters not near the active site also were found. By changing the allowed motion of the waters to be .4 Angstroms, no increase in the number of water molecules in the active site was found. The final structure, together with the four waters in the binding site is shown in Figures 5 and 6.

Finally, the above technique was applied to the experimentally determined human synovial PLA$_2$ structure^4 (which contained no calcium) and the water molecule positions were compared. We did not attempt to determine the positions of all the waters observed in the crystal in this manner, only the waters in the active site cleft, which are the most immobile, and the network. Since calcium is lacking, the water molecule

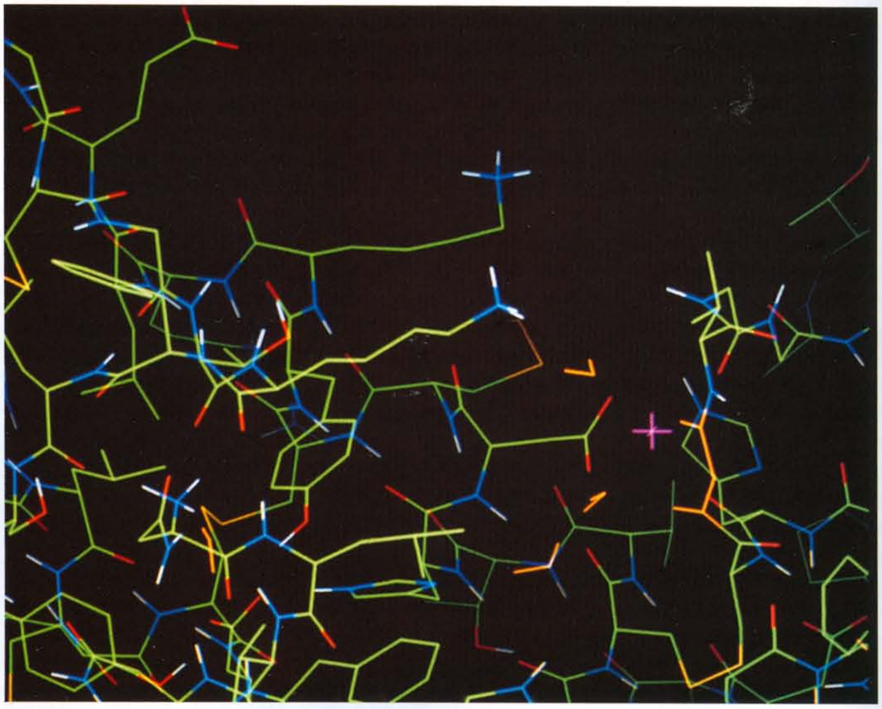

Figure 6. Final human synovial PLA$_2$ model, including proposed binding sites (closeup).

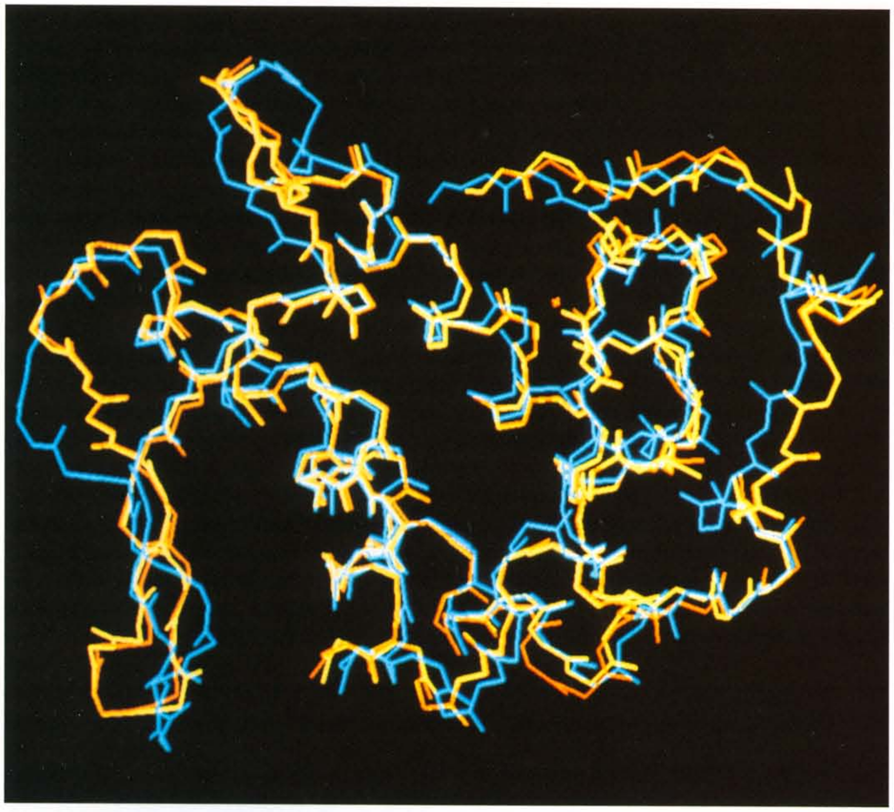

Figure 7. A comparison of the final human synovial model backbones (orange, with calcium; yellow, without calcium) with the experimental structure (blue) from Ref. 4.

positions are much different. The glycine 30 amide N-H bond is no longer directed toward the anionic binding site, and waters in the vicinity of the space occupied by the calcium in the bovine PLA$_2$ are bonded differently. A water molecule bound to the histidine 48 and the aspartic acid 49 presumably corresponds to the nucleophilic water in the PLA$_2$ containing calcium enzyme.

In this manner, several waters were found, the active site water and two others. In these experimentally determined human synovial PLA$_2$ coordinates, the network water molecule is absent. It is noteworthy that in agreement with this, we found no network water molecule during this simulation using these experimental PLA$_2$ coordinates; however, as mentioned above, using the model coordinates, a network water molecule was found.

Discussion

The model has been built on the basis of selected relevant information. It is necessarily a composite model. Its real test is a practical one: whether selective inhibitors can be designed on the basis of these results. The binding sites of the model as determined by looking for immobilized waters are similar to the binding sites of the bovine enzyme which contained calcium rather than the x-ray-determined synovial PLA$_2$ which contained no calcium, or the *C. atrox* structure.

A first test of the model was its usefulness in helping to solve the structure of human synovial PLA$_2$ by x-ray diffraction.[4] The model, after removal of the calcium, was used as a starting set of coordinates model for solving the x-ray structure. Figure 7 shows a comparison of the backbones of the models and the experimental structure.

Perhaps the model cannot be expected to have similarly configured loops as in the x-ray structure because the water environment is different than the crystal environment. More important is the comparison of the active site pockets. Except for several different residues, the model active site is most like the bovine PLA$_2$. This is basically due to the presence of calcium.

One feature not investigated in this work is the influence of the phospholipid membrane on the structural characteristics of our result. The PLA$_2$ molecule is associated with the cell membrane during its action. The efficiency and specificity of the inhibitors designed to fit the active site should be influenced by how those inhibitors interact with the membrane. The aqueous environment, the cell membrane, and the enzyme are all competing for the inhibitor, so that the binding free energy is a function of these effects. In addition, the membrane could change the configuration of the active site. Finally, the inhibition may be greatest when the enzyme, inhibitor, and membrane are bound together. The target may really be the enzyme-cell membrane combination. The limited goal of modeling the enzyme only is a necessary first step.

About the authors

Robert Hermann is a computational chemist at Eli Lilly and Company. He received his Ph.D. degree from Wayne State University in organic chemistry and was a National Science Foundation Postdoctoral Fellow at the Theoretical Chemistry Institute at the University of Wisconsin. His research interests include the calculation of hydrophobic interactions and the design of drugs through the calculation of ligand-receptor interactions.

Ernst Dow is an information analyst at Eli Lilly and Company and is currently working with artificial neural networks and their application to problems in chemistry. He obtained his B.S. degree in chemistry with a computer science option from the University of Pittsburgh.

References

1. Kramer, R. M., C. Hession, B. Johansen, G. Hayes, P. McGray, E. P. Chow, R. Tizard, and R. B. Pepinski, *Journal of Biological Chemistry,* Vol. 264, pp. 5768-5775, 1989.
2. Thunnissen, M. M. G. M., E. Ab, K. H. Kalk, J. Drenth, W. Dijkstra, O. P. Kuipers, R. Dijkman, G. H. de Haas, and H. M. Verheij, *NATURE,* Vol. 347, pp. 689-691, 1990.
3. Scott D.L., S. P. White, Z. Otwinowski, W. Yuan, M. H. Gelb, and P. B. Sigler, *Science,* Vol. 250, pp. 1541-1546, 1990.
4. Wery J. P., R. W. Schevitz, D. K. Clawson, J. L. Bobbitt, E. R. Dow, G. Gamboa, T. Goodson, R. B. Hermann, R. M. Kramer, D. B. McClure, E. D. Mihelich, J. E. Putnam, J. D. Sharp, D. H. Stark, C. E. Teater, M. W. Warrick, and N. D. Jones, *NATURE,* Vol. 352, pp. 79-82, 1991.
5. Ripka, W. C., 198th ACS National meeting, September 10-15, 1989. ACS Abstracts, MEDI 0045.
6. Davidson, F. F. and E. A. Dennis, *Journal of Molecular Evolution,* Vol. 31, pp. 228-238, 1990.
7. Subramanian, P. S., G. Ravishanker, and D. L. Beveridge, *Proceedings of the National Academy of Science,* USA Vol. 85, pp. 1836-1840, 1988.
8. P. J. Goodford, *Journal of Medical Chemistry,* Vol. 28, pp. 849, 1985.
9. Pitt W. R. and J. M. Goodfellow, *Protein Engineering,* Vol. 4, pp. 531-537, 1991.
10. Vedani, A. and D. W. Huhta, *Journal of the American Chemical Society,* Vol. 113, pp. 5860-5862, 1991.

Computational quantum chemistry in catalysis research

Aileen E. Alvarado-Swaisgood, Exploratory & Catalysis Research, Amoco Oil Research & Development, Naperville, Illinois

Analytical techniques in catalysis research continue to improve the level of detail at which we understand the chemical nature of the catalyst.[1] Ultimately, however, the chemistry of catalysis is the making and breaking of bonds, and analytical techniques do not yet provide sufficient information at this level of detail. Computational quantum chemistry in principle can provide such detailed information, but the computational challenge is enormous for even small tasks. Recently, however, high-speed, large-memory supercomputers, such as those from Cray Research, have made computational quantum chemistry a practical methodology for catalysis research. Today's supercomputers and computational techniques provide sufficient detail to supplement experimental catalysis research.

Computational approach

In choosing a computational technique to simulate a catalytic process, researchers need to decide specifically what they need to study. Catalytic processes can operate on various levels. Catalysis can be influenced by bulk phase properties, in which case the laws of thermodynamics apply. Catalysis also can be molecular, where molecules are attracted by electrostatic forces (physisorption), or an electronic interchange can occur among the molecules (chemisorption). Computational quantum chemistry allows for the study of the electronic structure of molecules, thereby providing the information necessary to distinguish between physical and chemical interactions.

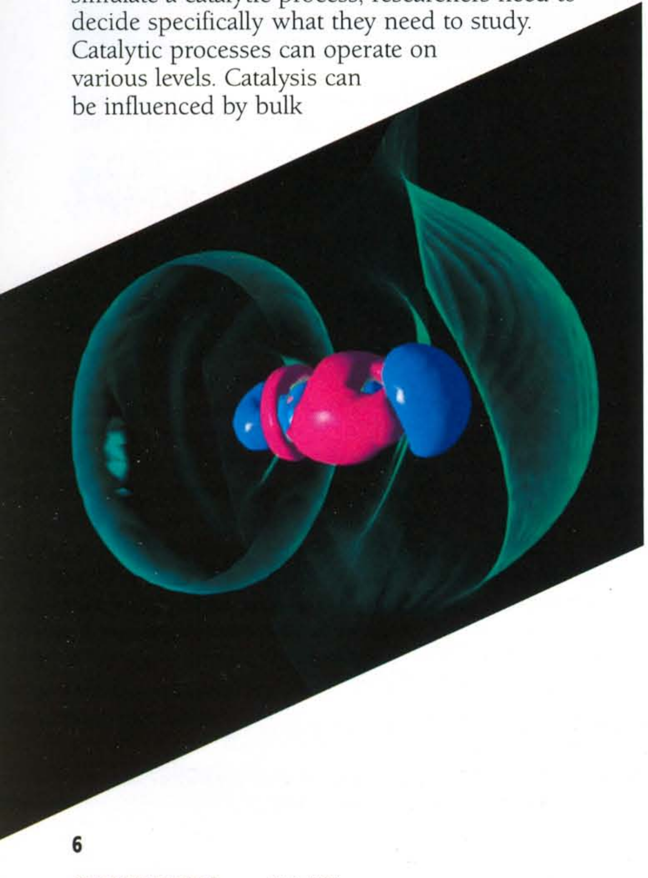

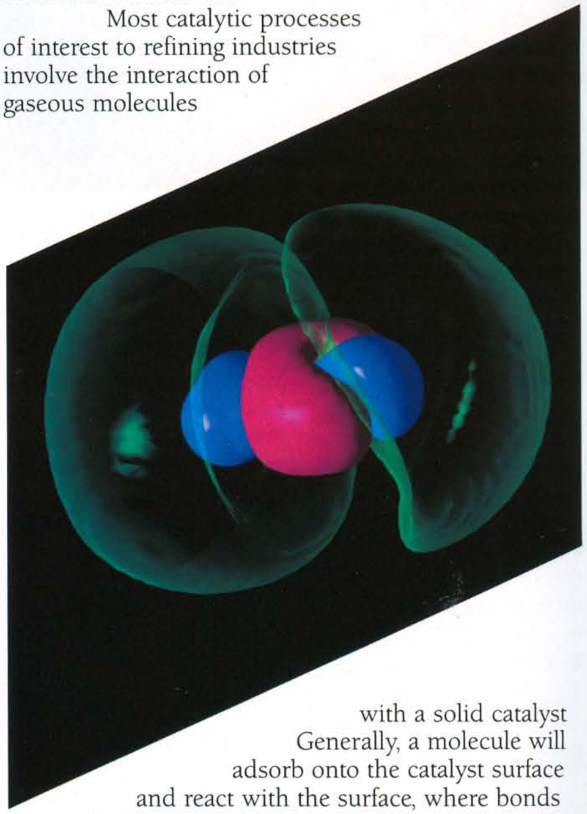

Figure 1. Three-dimensional volume renderings of the electron density difference for the $^5\Pi$ state of $CrCl^+$ as Cl approaches the Cr^+ ion. The red region is an isosurface for an electron density difference value of +0.001 e/au^3; the blue region is an isosurface for -0.001 e/au^3, and green shows regions of zero electron density difference.

Most catalytic processes of interest to refining industries involve the interaction of gaseous molecules with a solid catalyst Generally, a molecule will adsorb onto the catalyst surface and react with the surface, where bonds may be formed and broken, or molecules adsorbed on the surface may be more susceptible to reacting with each other. The resulting products then are desorbed from the catalyst's surface to complete the catalytic process.

Before we can study the interaction of a molecule with a solid catalyst, each needs to be studied independently. Once we understand them, we can understand their interactions. To represent the reactive process consistently from reactants to products, each system must be studied with the same computational technique. The interaction of CH_4 and Li/MgO provides an example.

A solid catalyst, like Li/MgO, is a large molecular system. To study solid materials as infinite systems, bulk techniques generally are used to provide information about the solid's electronic structure. These techniques take into account the periodicity of the solid structure. A gaseous molecule, like CH_4, is a small molecular system. To study gaseous molecules as independent systems, cluster techniques can be used.

Cluster techniques involve optimizing the electrons independently for a particular geometry. The calculations give an energy associated with the geometry and the electron density distribution. The density helps us define the chemical bonds. When groups of atoms are moved or rearranged, we can analyze which bonds are formed and which are broken from the electron redistribution. Stable molecular arrangements are those with the lowest relative energy, and the resulting electron density distribution will dictate the bonding structure.

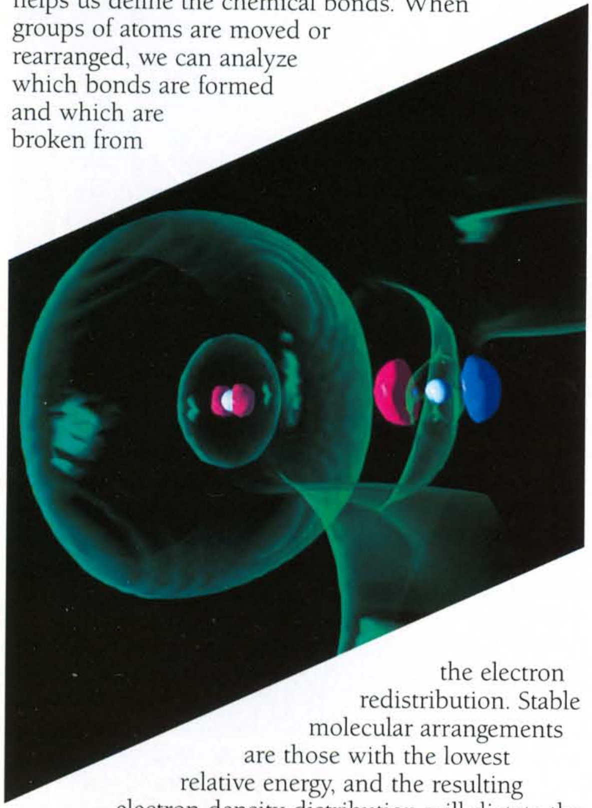

The interaction between CH_4 and Li/MgO entails the adsorption of methane on the catalyst surface, the activation of the methane C-H bond, and the subsequent formation of the resulting products (CH_3 + H). In comparison to the size of a solid catalyst, this interaction is localized at the molecular level. "Windowing in" on this interactive region will require a full description of the electrons of the gaseous, interacting molecules (reactants and products) in the vicinity of a representative section of the solid catalyst.

We cannot perform cluster calculations on an "infinite" solid catalyst structure. The amount of computer resources needed to attempt such a calculation is formidable in terms of memory, disk storage, and cost. We can, however, perform cluster calculations on smaller "cluster" (finite) models of the solid catalyst structure.

Applications

Following are three examples of catalytic mechanisms into which we have gained insights through computational research: electron transfer in $CrCl^+$, Li substitution into MgO, and benzene transport through faujasite pores.

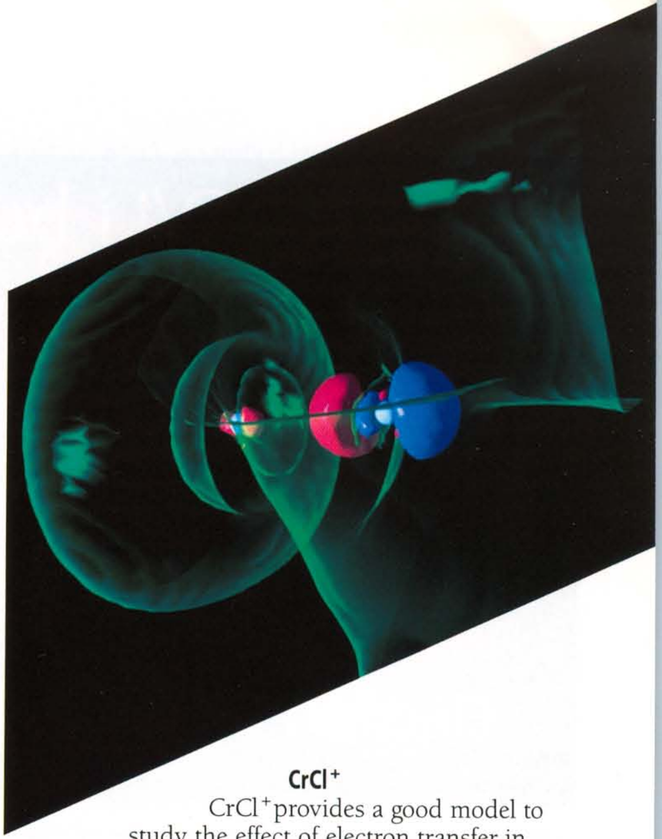

$CrCl^+$

$CrCl^+$ provides a good model to study the effect of electron transfer in enhancing the catalytic properties of a potential activant. Guided ion beam mass spectrometry experiments have shown that Fe^+ reacts exothermically with small alkanes, but Cr^+ does not.[2] On the other hand, $CrCl^+$ (Ref. 3) and CrO^+ (Ref. 2) are found to react exothermically with alkanes — $CrCl^+$ dehydrogenates and CrO^+ forms alcohols:

$$CrCl^+ + \text{n-butane} \rightarrow ClCr^+ - C_4H_8 + H_2$$
$$CrO^+ + RH \rightarrow Cr^+ + ROH$$

In these examples, chlorine and oxygen chemically activate an unreactive transition-metal ion.

To understand how $CrCl^+$ can activate alkanes, it is worth examining the electronic structure of each atom. The ground (^6S) state Cr^+ ion has a $3d^5$ electronic configuration and its first excited (^6D) state has a $4s3d^4$ configuration which is 1.5 eV (35.1 kcal/mol) higher in energy. For Cr^+ to be reactive, its stable half-filled valence shell of electrons will have to be excited to a state where the electrons (e.g., 4s electron) are more accessible for bonding. Fe^+ has a $4s3d^6$ ground state and does not require electron excitation for the reaction to proceed.

Cl in the ^2P ground state has only one unpaired electron. How this electron associates with the Cr^+ ion will define the electronic state of the $CrCl^+$ molecular ion. A $CrCl^+$ $^5\Sigma^+$ state is formed when the singly occupied Cl p_π orbital approaches Cr^+ and forms a σ bond. If, instead, a doubly occupied p_σ orbital from Cl approaches Cr^+, the resulting molecular state is a $^5\Pi$. The Cl unpaired electron (p_π orbital) then is available to π bond with the unpaired electron in the Cr^+ d_π orbital.

We performed ab initio calculations on $CrCl^+$ using a CRAY-2 supercomputer.[4] Figure 1 shows a three-dimensional volume rendering of the electron density difference for the $^5\Pi$ state of $CrCl^+$ as Cl approaches the Cr^+ ion. At longer distances, the electrons are polarized toward the σ bonding region between the two atoms. At distances below 5.2 Bohrs, a dramatic electron transfer occurs from the Cr$^+$ d_π orbital to the Cl p_π orbital. Electron density differences are obtained by adding the electron density distribution of the individual atomic fragments (Cr$^+$ and Cl) within the same three-dimensional region in space and substracting it from the electron density distribution of the molecule (CrCl$^+$). Each image shows the net electron density distribution at different Cr$^+$...Cl internuclear distances.

The CrCl$^+$ $^5\Sigma^+$ state is only 5 kcal/mol more stable than the $^5\Pi$ state. We believe that the $^5\Pi$ state has the electronic structure conducive to activate a C-H alkane bond.[3,4] The resulting $^5\Pi$ CrCl$^+$ takes on a [Cr^{++}...Cl$^-$] component in the π bonding region. "Cl$^-$" can attract the hydrogen atom. As H dissociates from C, the excess electrons on Cl can flow back easily to Cr, where the Cr d_π orbital then can bond to C. We have yet to determine whether this process occurs homolytically or heterolytically.

Another interesting aspect of CrCl$^+$ activating an alkane is that Cr$^+$ and Cl are not observed to dissociate during the reaction process. A calculation of Cr$^+$ interacting with HCl shows(4b) that there is a strong charge-dipole interaction of 20 kcal/mol — strong enough to withstand the next step involving dehydrogenation, where the H on Cl associates with a β-H from the alkane to form H_2.

Li substitution into MgO

The search for a catalyst that will convert natural gas to other profitable feedstocks is an ongoing research venture for many refining industries and universities. Natural gas is composed mainly of methane, which is generally burned as fuel for heating or cooking. Converting methane to, for example, higher hydrocarbons is a formidable task. Thermodynamics favor complete combustion (when oxygen is used) or coke formation (when pyrolyzed). The goal is to find a catalytic process that will activate the first C-H bond while providing the environment where the radicals formed can combine in a controlled fashion to the preferred products.

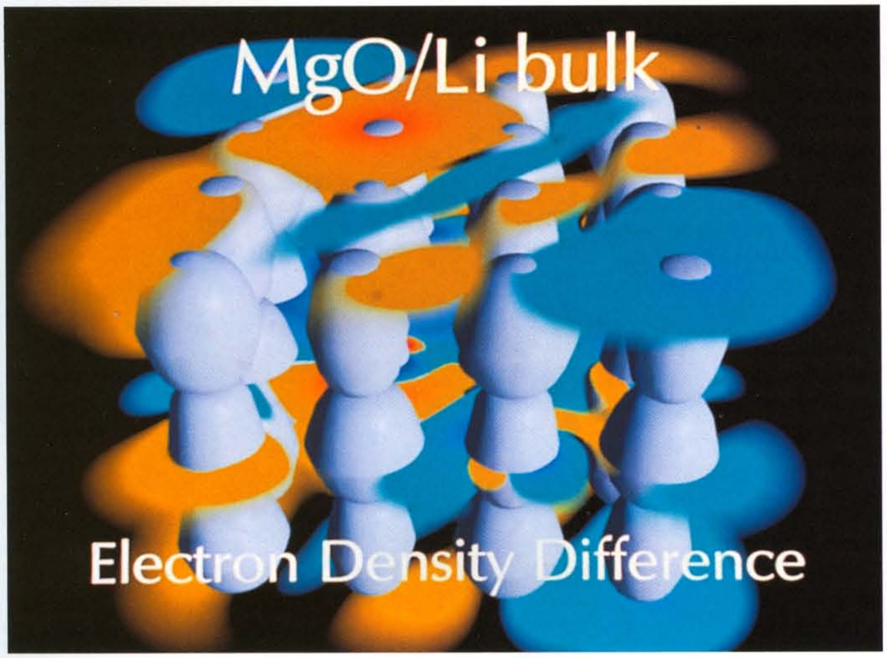

Figure 2. Net electron density distribution of MgO when Li is substituted into a bulk position. Yellow regions denote areas of electron gain due to Li substitution, and blue regions denote areas of electron loss.

Oxidative coupling has received considerable attention during the past five to ten years. This process involves the activation of CH_4 by O$^-$ species found on the catalyst's surface, from either catalytically activated O_2 or surface lattice oxygen atoms. One of the most widely studied oxidative coupling catalysts is Li/MgO. It is believed that when Li is substituted into the MgO framework, an O^{2-} atom is converted to an O$^-$ atom in the proximity of the Li$^+$ ion.[5] This makes a surface oxygen more reactive to activate methane.

We were interested in studying the effect Li$^+$ has on the electron density distribution of MgO. To study this effect, we used a 48-atom MgO cluster model. We substituted a Li atom for a Mg atom and used a CRAY-2 supercomputer to determine the optimum electron density distribution. Figure 2 shows the net electron density distribution of MgO when Li is substituted into a bulk position. The effect Li has on MgO is not localized to a specific oxygen but to the entire oxygen lattice framework. As methane approaches a "perturbed" surface environment where Li is substituted into the MgO structure, it will encounter these surface electron density differentials.

Benzene transport through faujasite pores

Zeolites provide another example of an oxide system used in a number of petroleum refining processes. Zeolites are large three-dimensional molecular sieves noted for their internal pores and cavity structures. Reactants are converted to the desired products as they flow through these pores and cavities. Knowledge of the transport properties of zeolites is of critical importance to ensure that the reactants reach the reactivity zones and products are able to exit before unnecessary conversions take place.

One type of zeolite widely used in the refining business is called *faujasite*. To study transport in the faujasite pores at the molecular and electronic levels, we prepared a 72-atom cluster model taken from the peripheral atoms that form a pore. The peripheral atoms comprise 12 silicon and 12 oxygen atoms connected in alternating order to form a pore 7.4 Angstroms in diameter. Because the silicon atoms are tetrahedrally coordinated, two more oxygen atoms are bonded to the silicon atoms to form part of the rest of the three-dimensional faujasite molecular framework. To limit the size of the model, these "framework-connecting" oxygen atoms are capped with hydrogen atoms, thus giving us the 72-atom cluster model.

Benzene is used as the probing molecule. Benzene can be found in the catalytic cracking process as either a by-product, where aromatic fragments are inherent in the reactant structures, or as a product through dehydrocyclization of $\{C_{6+}\}^+$ carbocation intermediates. Figure 3 shows benzene migrating through the pore "head-on" starting with two hydrogens, then two carbons, and so on (2-H entrance). Figure 4 shows all atoms in the benzene molecule entering at the same time (face entrance).

There is a stabilization energy of 5 kcal/mol when benzene enters the pore head-on. The stabilization energy gained by a face entrance is 18 kcal/mol — more than three times the stabilization energy of a perpendicular head-on (2-H) entrance. In a face entrance, all six benzene hydrogen atoms are attracted to the oxygens in the pore structure.

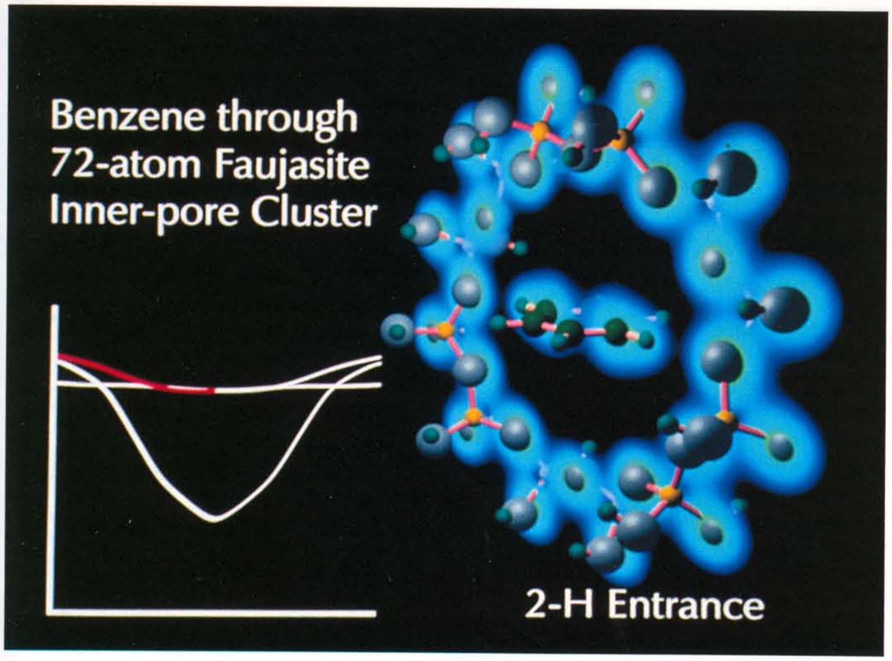

Figure 3. Model of a benzene molecule migrating through a faujasite pore perpendicular to the plane of the pore.

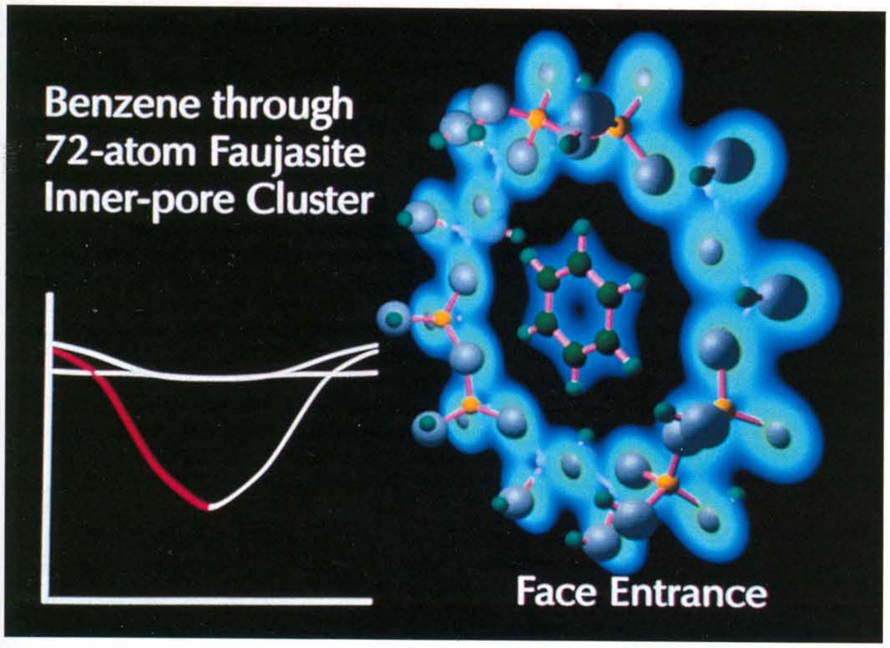

Figure 4. Model of a benzene molecule entering a faujasite pore parallel to the plane of the pore.

Experimental researchers have shown that only one benzene molecule is found to occupy the pore site.[6] We have shown that it would take more than three benzene molecules to enter head-on to be as energetically stable as one benzene molecule entering with all its atoms at the same time. Fitting more than three benzene molecules head-on is sterically impossible, thus verifying experimental observations.

Catalysis research in the future

As mentioned in the previous section, zeolites are used in a number of petroleum refining processes, not only for their transport properties through the three-dimensional porous structures, but also for their potential catalytic activity. When Si^{4+} atoms are substituted for a number of trivalent cations (such as B^{3+}, Al^{3+}, Ga^{3+}, etc.), there is a charge imbalance which is satisfied by including monovalent cations (such as H^+, Li^+, Na^+, K^+, etc.). Transition metals also are added to zeolites to enhance their catalytic activity; in such cases catalysis often occurs on the transition metal surface, and the zeolite serves as a molecular support system. Exchanges and additions of this nature are commonplace in the development of new catalysts, providing a reactive environment "custom-made" to obtain the desired products.

Computational quantum chemistry will continue to play an active role in defining our needs for the future. Electronic effects during reaction processes are only beginning to be understood. Computational quantum chemistry also is used to develop parameters for popular empirical techniques, such as molecular mechanics and molecular dynamics. With experimental techniques continually improving, and computational chemistry providing new insights, the eventual formulation of methodologies for catalyst design seems assured. ▬

Acknowledgments

The author thanks her collaborators in each of the applications discussed.

$CrCl^+$: *Jeffrey Tilson (Amoco graduate summer intern) and James F. Harrison of Michigan State University.*

Li/MgO: *Mark Barr, Keith Baumruck, and Don Washecheck of Amoco Oil R&D; Gerry Zajac of Amoco Corporation Analytical Services Division.*

C_6H_6/faujasite: *April Swanson, Don Hopkins, Mark Barr, Chris Marshall, and Keith Baumruck of Amoco Oil R&D.*

Scientific visualization: *Jeffrey Thingvold, Steve Pruitt-Lane, and Matthew Arrott, Scientific Visualization Group, National Center for Supercomputing Applications.*

Computational resources: Amoco Oil R&D, Amoco Corporation SFD and I&CS Divisions (project funding); National Center for Supercomputing Applications (access to CRAY X-MP and CRAY-2 systems); Arthur Freeman and Ye Ling of Northwestern University, George Fitzgerald of BioSym Technologies, Inc. (now with Cray Research, Inc.); and Harrell Sellers, NCSA Chemistry Group.

About the author

Aileen E. Alvarado-Swaisgood received her B.S. degree in chemistry from the University of Puerto Rico-Rio Piedras and her Ph.D. degree in theoretical physical chemistry from Michigan State University. Since leaving Michigan State, she has been a research scientist in the Exploratory & Catalysis Research Division of Amoco Oil Research & Development.

References

1. Hegedus, L. L. (ed.), *"Catalyst Design: Progress and Perspectives,"* John Wiley & Sons, Inc., New York, New York, 1987.
2. Armentrout, P. B., L. F. Halle, and J. L. Beachamp, *Journal of the American Chemical Society,* Vol. 103, p. 6501, 1981.
3. Mandich, M. L., M. L. Steigerwald, and W. D. Reents, *Journal of the American Chemical Society,* Vol. 108, p. 6197, 1986.
4. (a) Alvarado-Swaisgood, A. E. and J. F. Harrison, *Journal of Physical Chemistry,* Vol. 92, p. 5896, 1988; (b) Tilson, J., J. F. Harrison, and A. E. Alvarado-Swaisgood, in preparation.
5. Ito, T., J.-X. Wang, C.-H Lin, and J. H. Lunsford, *Journal of the American Chemical Society,* Vol. 107, p. 5062, 1985.
6. De Mallmann, A. and D. Barthomeuf, *Journal of Physical Chemistry,* Vol. 93, p. 5636, 1989.

Computational chemistry in parallel

John E. Mertz and Jan W. Andzelm, Cray Research, Inc.

Researchers in the competitive fields of drug and materials design know the value of timely solutions to their research problems. Unfortunately, answering the routine question "What does my molecular system look like?" is often a lengthy process involving sophisticated and time-consuming experiments. In the race to design novel compounds, knowing the molecular structure of a target system has obvious advantages. Because computational methods now play a vital role in the prediction and refinement of molecular structures, many industrial researchers involved in molecular design use Cray Research supercomputers to increase their competitiveness. With the need for fast turnaround in mind, the chemistry group at Cray Research, along with several academic research groups, has developed parallel algorithms that have been incorporated into many popular application programs.[1] In each of the four major methods used in computational chemistry — force field, Hartree-Fock, semiempirical, and local density functional — the parallelized application programs can solve problems at a rate of well over one GFLOPS (one billion floating-point operations per second) on Cray Research supercomputers.

Force field methods

Force field methods can be used to investigate the structure and dynamics of large molecules such as proteins and polymers. Most force fields are of the functional form, $U = \Sigma_q U(q) + \Sigma_{i,j} U(r_{ij})$, where $U(q)$ is the part of the potential that is described in the internal coordinates (q) of the molecules, and $U(r_{ij})$ accounts for the interactions between pairs of atoms. This potential is based on geometric and energetic data obtained from computational and experimental sources. Powerful techniques have been developed that combine x-ray and nuclear magnetic resonance (NMR) data with this potential.[2] These techniques allow researchers to refine their molecular structures faster and in a more automated fashion. In these application programs most of the computational effort is spent in subroutines that are involved in the evaluation of the potential energy, *U*, and its derivatives. In these subroutines the typical DO-loop structure is like that shown in Figure 1a. In some cases, vectorization of this loop is possible because the values of the index array are sorted in ascending order. In other cases, where it is not possible (or convenient) to organize the elements of *index*, the "update loop" feature included in version 5.0 of Cray Research's CF77 Fortran compiling system allows this loop to be executed in vector mode. If this loop, as written in Figure 1a, is executed in parallel, incorrect results will occur because different processors can update the array elements of *f(j)* simultaneously. (In the jargon of parallel processing this constitutes a *critical region.*) This problem is easily overcome by allowing each processor to have its own copy of the *f(j)* array (Figure 1b). The command *tskvalue* returns the logical CPU number of the processor executing the task. This guarantees that each processor works with its own copy of the *f(j, icpu)* array. After the entire set of subroutines in the program has computed *f(j, icpu)*, these private copies are combined to yield the final *f(j)* array (Figure 1c). Parallel speedup factors (the ratio of single-processor wallclock time to multiple-processor wallclock time) of 7.9 on eight processors are not difficult to achieve with this loop restructuring.

Hartree-Fock method

One of the most popular methods used to predict the equilibrium geometries of small organic molecules is the Hartree-Fock (HF) method. Organic and medicinal chemists routinely use this technique to study their compounds and drugs. This method is an ab initio method because molecular properties are

Figure 1. Transformation of the force calculation from sequential/vector mode to parallel/vector mode.

```
                            cmic$ parallel
do i = 1, n                 icpu = tskvalue + 1                        do icpu = 1, ncpus
  j = index (i)             cmic$ do parallel vector                   cmic$ do all vector
  f(j) = f(j) + (∂U/∂rj)  = do i = 1, n                              +   do j = 1, m
end do                        j = index(i)                                 f(j) + f(j) + f _p(j, icpu)
                              f _p(j, icpu) = f _p(j, icpu) + (∂U/∂rj)   end do
                            end do                                     end do.

        (a)                               (b)                                  (c)
```

```
do i = 1, p_basis
  do j = 1, q_basis
    do k = 1, r_basis
      do l = 1, s_basis
        call twoel
        call fock
end do
```

(a)

=

```
do i = 1, p_basis
  do j = 1, q_basis
    do k = 1, r_basis
      do l = 1, s_basis
        ntask = ntask + 1
        call collect
end do
```

(b)

+

```
cmic$ do all
do i = 1,ntask
  call p_twoel
  call p_fock
end do
```

(c)

+

```
do icpu = 1, ncpus
fockmat = fockmat + fockmat_p(icpu)
end do.
```

(d)

Figure 2. Transformation of the building of the Fock matrix from a sequential to a parallel algorithm.

calculated using only the basic information about the molecule: the number and type of atoms and the total charge on the molecule.

The HF method usually is applicable only to systems containing less than 50 atoms. This is because the computational effort scales as N^4, where N is the number of atoms in the molecule. Although algorithmic developments such as the direct scheme[3] allow for treatment of larger systems (the scaling is slightly less than N^3), the HF method does not correctly address the issue of electron correlation. This effect is important when investigating inorganic and metallic compounds, but including electron correlation in the model requires algorithms that scale as N^5 or more.

Almost all the computational work in HF application programs is done in the evaluation of the "two-electron integrals," I_{ijkl} and the Fock matrix, F. The two-electron integrals are used to construct the Fock matrix through formulas such as $F_{ij} = F_{ij} + D_{kl} * I_{ijkl}$. In the program DISCO/2 these tasks are performed in the subroutines *twoel* and *fock*[3] (Figure 2a). These subroutines are called repeatedly from inside nested loops that run over the number of basis functions on sets of atoms. At first glance this algorithm may seem trivial to parallelize, but several technical points complicate a straightforward transformation.

First, *twoel* calls many subroutines that in turn call more subroutines. Within this calling tree, hundreds of variables are shared among the procedures. In addition, hundreds of variables are local to one or a few subroutines. To ensure proper execution, each of these variables must be classified properly, or scoped, as private or shared among the processors. Cray Research software tools such as Atscope and Atexpert, along with parallel processing features such as task common, make this job manageable.

Second, the outermost loop typically completes less that 50 iterations, which is not enough to distribute work equally among the processors. To remedy this situation the loops are unrolled by collecting all the necessary information (Figure 2b) and executed in parallel as one loop (Figure 2c). The subroutines *p_twoel* and *p_fock* are versions that have been modified for the correct data scoping. The critical region that exists when the Fock matrix is collected is circumvented by allowing each processor to work on its local copy of *fockmat_p(icpu)*, each of which is combined subsequently into the final, shared values of *fockmat* (Figure 2d).

Because such a large portion of the overall computational work is done in this parallel region, speedups of 7.9 on eight processors are obtained routinely. Figure 3 displays a molecule that is used to model a section of the plastic, polycarbonate. The DISCO/2 program running on a Cray Research supercomputer solves the HF equations for this molecule at a rate of 1.5 GFLOPS.

Semiempirical methods

Semiempirical methods have been developed to facilitate calculations on larger molecular systems. Their name derives from the fact that experimental data, such as molecular properties, are used to parameterize certain terms in the HF equations. This reduces the number of two-electron integrals necessary to solve the HF equations. When used on classes of compounds where the parameterization has been performed, these methods provide reasonable energetic and structural results.

Computationally, semiempirical methods are very efficient; the scaling can be less than N^3. This is because the work is dominated by matrix

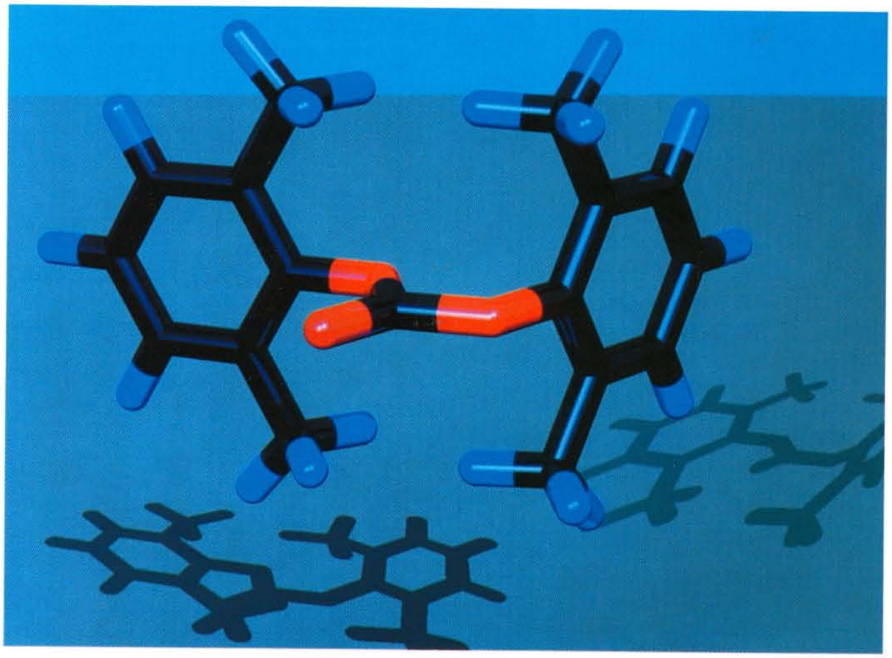

Figure 3. Molecular model representing polycarbonate. The calculations were done at the Hartree-Fock level of theory using 316 basis functions and no symmetry.

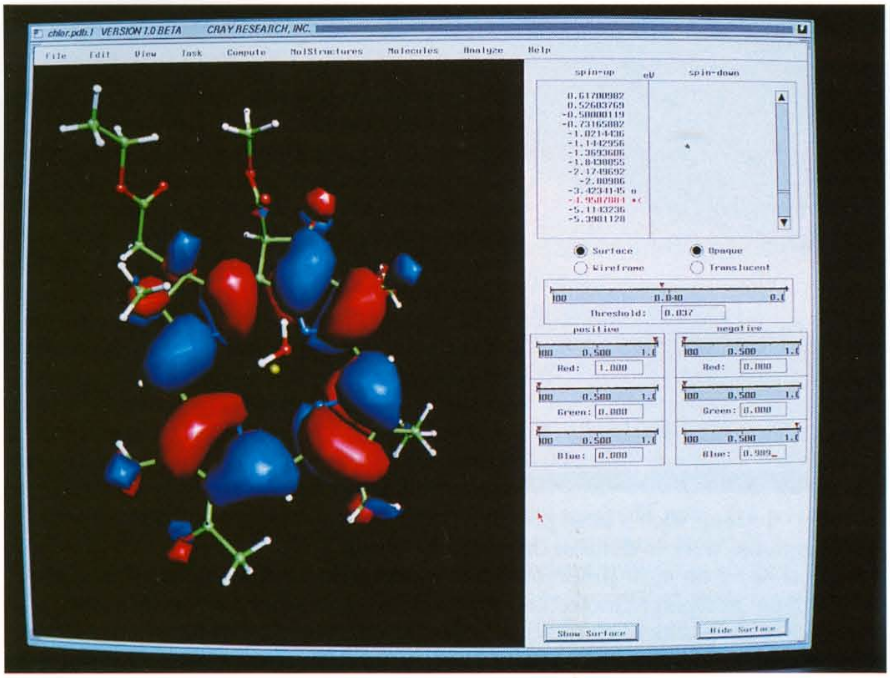

```
spin-up        eV        spin-down

 0.61700982
 0.52603769
-0.50000119
-0.73165082
-1.0214436
-1.1442956
-1.3693606
-1.8430055
-2.1749692
  -2.00986
-3.4234145 o
-4.9587884 *(
-5.1143236
-5.3981128
```

● Surface ○ Wireframe

● Opaque ○ Translucent

Threshold: 0.037

positive

Red: 1.000

Green: 0.000

Blue: 0.000

negative

Red: 0.000

Green: 0.000

Blue: 0.989_

Model of a section of bacteriochlorophyll *b* created with Cray Research's DGauss quantum chemistry program.

operations, some of which can be multitasked automatically by the Autotasking feature of the CF77 5.0 compiling system, and others can be parallelized via compiler directives. Through careful vectorization and calls to Cray Research library routines, Professor Walter Thiel of the University of Wuppertal, Germany, has achieved over 2.2 GFLOPS on a CRAY Y-MP8 system using his semiempirical program MNDO.[4]

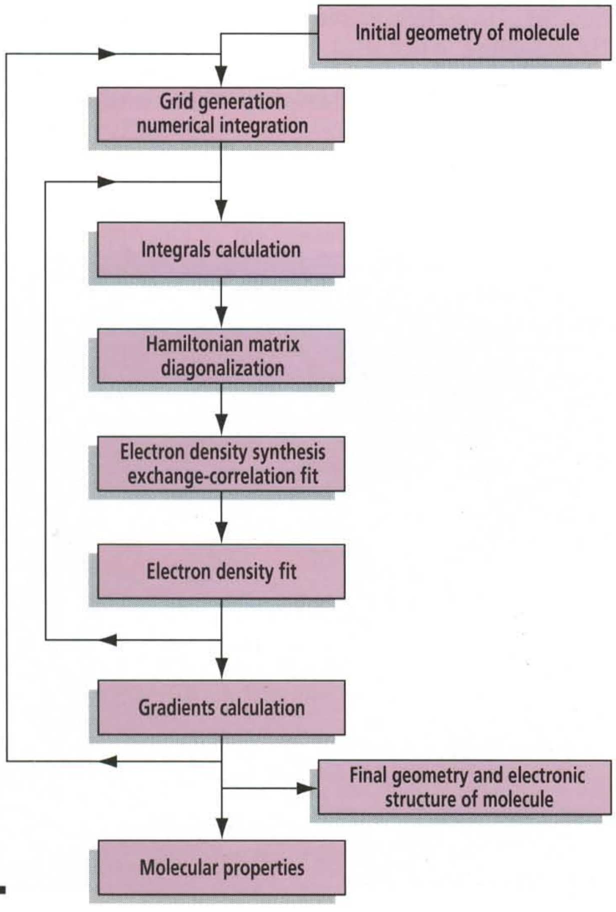

Figure 4. Organization of the DGauss program.

Density functional methods

Although the density functional (DF) method has a rich history in physics, it has entered chemistry only recently. This is because for a long time it could not predict molecular structures and energies to the accuracy needed by chemists. Therefore, the majority of chemists relied on ab initio and semiempirical techniques instead. The advent of automatic (gradient) algorithms to optimize molecular structures and the accurate (nonlocal) functionals have made the DF method a valuable tool for chemical research. Like the HF method, the DF method is an ab initio method; unlike the HF method, it takes into account electron correlation and therefore can be used to study a variety of systems, including metals and inorganic systems.

Various techniques are being used to solve DF equations for molecular systems. For the most part these methods differ in the form of functions used to represent molecular orbitals. One can use numerical functions, Slater-type orbitals, plane-wave functions, or Gaussian-type functions. Here we focus on the Gaussian-type functions.

The DGauss program (for Density-Gaussian) is an analytical implementation of the DF method that uses Gaussian-type orbitals.[5] This constitutes a bridge for comparison between HF and DF Gaussian-based methodologies and allows DF methodologies to build on the wealth of experience gained from HF methods. The Gaussian-based DF method is potentially the fastest approach to solving DF equations for molecular systems. The Gaussian DF method provides important functionality by allowing the calculation of forces and molecular force constants through first and second derivatives of the energy. This functionality leads to automatic geometry optimization of molecular structures, calculations of molecular vibrations, and the study of molecular reactivity. The DGauss method uses Gaussian-type functions to approximate molecular orbitals, electron density, and the DF exchange-correlation potentials. This is done by expanding the orbitals, density, and potentials in terms of Gaussian functions where the goal is to find linear coefficients (vectors) of those expansions by solving DF equations.

The scheme of calculations is presented in Figure 4. The structure of the DGauss program includes two cycles: the outer geometry cycle and the inner SCF cycle. Four computational steps are involved in the SCF cycle. For a particular starting geometry, the initial coefficients are assumed, and the multicenter integrals involving Gaussian-type orbitals are calculated. Then the Hamiltonian matrix is built and diagonalized. This step provides vectors representing molecular orbitals. Next, the molecular orbitals are used to generate DF exchange-correlation potential and total density coefficients, respectively. After this, we go back to the first step, calculate integrals, and build the Hamiltonian matrix. This process continues until molecular orbital vectors do not change from iteration to iteration. This process is called the self-consistent method (SCF) because the output of one iteration is being used to generate input for the next. At the end of the SCF cycle the converged vectors are used to calculate gradients of the total energy with respect to the nuclear displacements. This information is used to calculate a new molecular geometry. Iterations in the geometry cycle typically continue until the minimum energy structure of the molecule is reached.

Most of the computation is in the SCF cycle. Depending on the specifics of the calculation, one of two algorithms is used in this cycle: the direct algorithm, which is designed to minimize input/output (I/O) by calculating the integrals at every iteration of the SCF cycle, or the I/O algorithm, which calculates the integrals once at the beginning of SCF cycle and stores them on a disk from which they are retrieved during every SCF cycle. For a molecular system with 1000 basis functions, approximately 350 million integrals will need to be generated; they must be processed 20 to 30 times during the SCF procedure. On a system of this size it is obvious that the I/O approach saves considerable processor work (up to 40 percent), but because this work is sequential it is a bottleneck for parallel processing. Therefore, in this case the direct scheme is the best approach for reducing wallclock time. In other cases, optimal performance is reached by balancing these factors.

The integral calculation step is a major portion of the overall computational work, and it has been efficiently vectorized and automatically multitasked through compiler directives. Speedups greater than a factor of 7 typically are attainable in this section of the program (Figure 5). BLAS3 matrix multiplication routines are used to set up the Hamiltonian matrix, and the Eispack procedure is used for diagonalization. The "exchange" step involves the synthesis of the density from molecular orbitals, which can be accomplished very effectively with a sparse matrix algorithm. Like the integral step, the fitting of density ("density step") is dominated by vector algorithms and can be parallelized by compiler directives. There is also an additional startup time that involves matrix algebra operations for the direct scheme and the vector work for the I/O approach. Before the SCF procedure begins, the grid generation and numerical integration ("grid step") are performed using a sparse matrix algorithm. To achieve significant performance improvements, each of these sections of the program must be parallelized.

Summary

We have shown several examples of how parallel processing is implemented in existing application programs on Cray Research computer systems. Although the results demonstrate that a moderate number of processors can yield significant performance improvements, the enormous speedups promised by massive numbers of processors are difficult to achieve. Consider, for example, that in the force field methods we typically see that 85 to 95 percent of the computation can be executed in parallel. The serial section of the program is indispensable because it is collecting data that the program is generating. For example, if a job is 95 percent parallel and requires 1000 hours to complete, it executes 50 hours in serial mode and 950 hours in parallel mode. The wallclock time required to complete the job would be (50 hours + 950/8 hours) = 169 hours: a speedup factor of 6. But the same job executed on a huge number of processors can achieve a speedup factor of no better than 20; (50 hours + 0 hours) = 50 hours. In this "real-world case" thousands of processors are at best a disappointing factor of 4 times faster than eight processors. Unless the level and quality of parallelism are near perfect, as in the Linpack Benchmarks, the full power of a highly parallel system cannot be delivered. Achieving this perfection is and will continue to be a grand challenge to software developers.

	2 CPUs	4 CPUs	8 CPUs
Integrals	2	3.9	7.7
Exchange	2	3.9	7.4
Density	2	3.8	7.0
Grid	2	3.7	6.2
Total	2	3.8	7.1

Figure 5. Performance improvements obtained from vectorized and parallelized DGauss calculations for a model of zinc-insulin using 1031 basis functions. The integral step is the most parallelized section of the code. The overall speedup of 7.1 corresponds to a processing rate of 1.1 GFLOPS (1.1 billion floating-point operations per second).

Acknowledgments

The work on the DISCO/2 program was performed in collaboration with Hans-Peter Luthi of the Interdisciplinary Project Center for Supercomputing at the Eidgenossische Technische Hochschule (ETH) in Zurich. John Mertz thanks Dr. Lüthi for his hospitality and support during sabattical at ETH.

About the authors

John E. Mertz joined Cray Research in June 1987, where he now researches and develops applications of statistical mechanics to solve problems in the pharmaceutical and chemical industries. Mertz received his Ph.D. degree in chemical physics from the University of Illinois in 1987.

Jan W. Andzelm joined Cray Research in April 1987 as an expert in ab initio quantum chemistry. He is the developer of the DGauss computational chemistry program, which allows the quantum-mechanical simulation of large and complex systems of molecules and materials. Prior to joining Cray Research, Andzelm was a research associate at the University of Montreal. He received his Ph.D. degree from the University of Warsaw, Warsaw, Poland.

References

1. See *Directory of Applications Software on Cray Research, Inc., Supercomputers,* Cray Research, Inc., 1991.
2. Brünger, A., "Crystallographic Refinement by Simulated Annealing on Supercomputers," *CRAY CHANNELS,* Vol. 10, Number 3, pp. 16-19, fall 1988; Yip, P. and D. A. Case, *Journal of Magnetic Resonance,* Vol. 83, pp. 643-648, 1989.
3. Almlöf, J., K. Faegri, Jr., and K. J. Korsell, *Journal of Computational Chemistry,* Vol. 3, pp. 385-399, 1982; Lüthi, H. P., J. E. Mertz, M. W. Feyereisen, and J. Almlöf, University of Minnesota Supercomputer Institute Research Report UMSI 91/141, 1991.
4. Thiel, W., program MNDO90, Version 3.1, July 1990.
5. Andzelm, J. and E. Wimmer, *Journal of Chemical Physics,* submitted.

UniChem

A supercomputing environment for computational chemistry

Jeffrey F. Gaw, Cray Research, Inc.

Scientific research as it is practiced today comprises three interacting components: experiment, theory, and computation. Computation is theory in action; it allows researchers and engineers to model complex real-world behavior numerically and make meaningful predictions. In chemistry, ab initio computations — those based on quantum-mechanical first principles — are even challenging many experiments. Cray Research's UniChem software environment provides the chemist with a window into the world of computational quantum chemistry.

The UniChem environment is Cray Research's response to the application software needs of customers in the chemical and pharmaceutical industries. Industrial computational chemists were searching for ways to maximize their computational performance; increase their efficiency, productivity, and effectiveness; and access the latest computational methods without taking time out from their primary job of "doing chemistry." To ensure that its software solution met the needs of industrial chemists, Cray Research organized a consortium of five major chemical and pharmaceutical companies to help develop and test UniChem: 3M, Du Pont, Eli Lilly and Company, Exxon Research and Engineering, and Monsanto and its subsidiaries G. D. Searle and NutraSweet. This group together with Cray Research established the software requirements and participated in the testing of the alpha and beta versions of the product.

Network supercomputing for chemists

Version 1.0 of the UniChem software environment was announced in February 1991. The UniChem environment eliminates barriers that many industrial chemists face in running complicated calculations. It couples Cray Research supercomputer power with

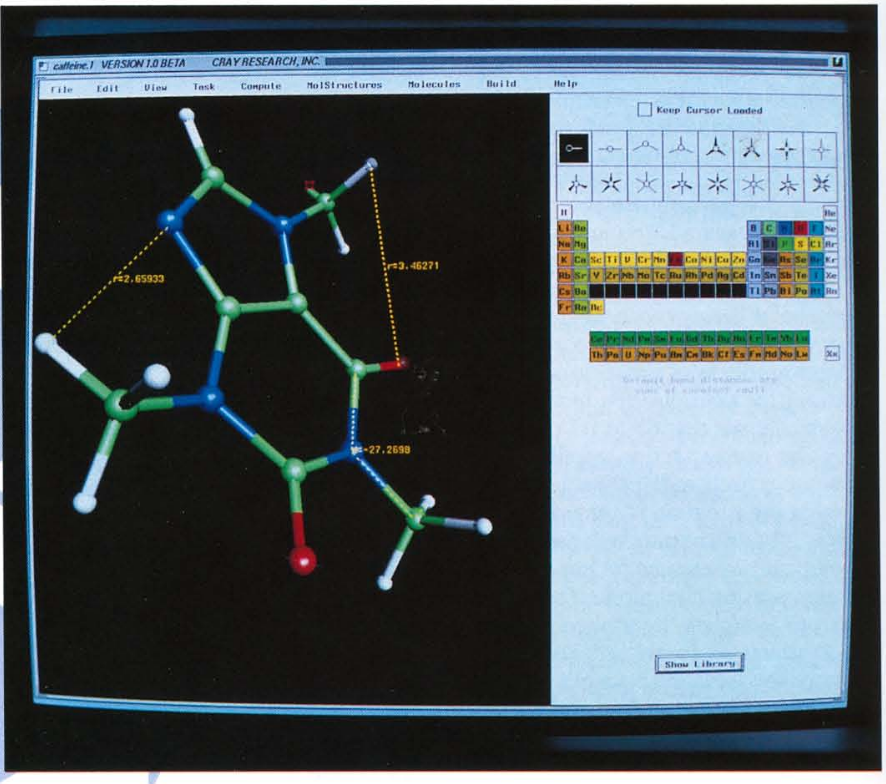

UniChem screen in "build" mode showing a model of the caffeine molecule.

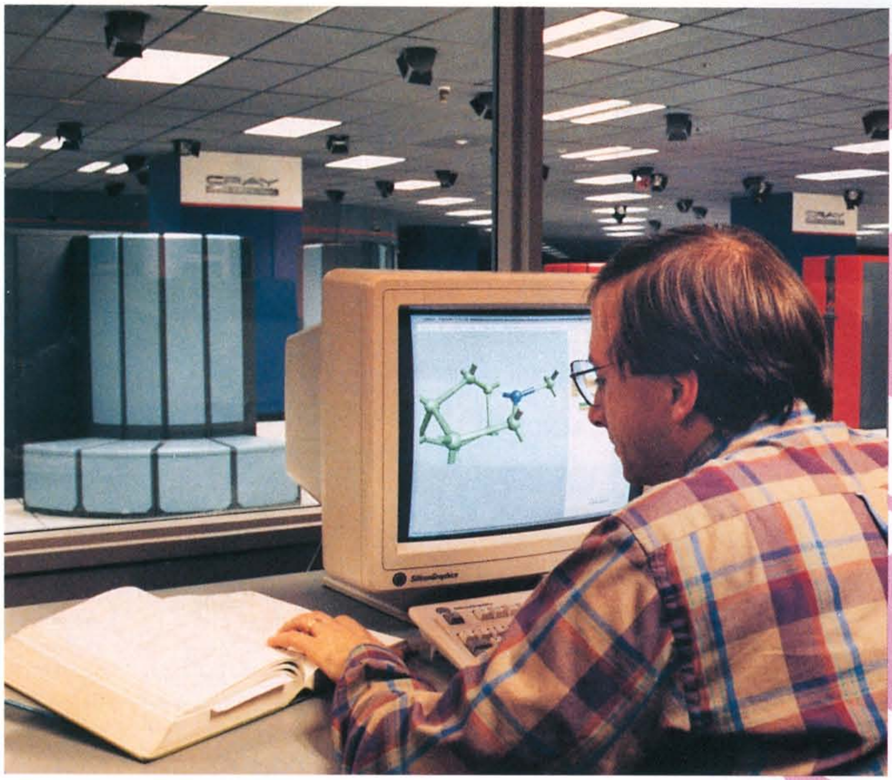

the graphics display capabilities of Silicon Graphics workstations.

The workstation occupies one side of the UniChem window, through which the chemist

☐ Builds molecular structures, either with single atoms or through a molecular fragment library
☐ Edits structures
☐ Sets up the input to the calculation through a graphical point-and-click display
☐ Launches the calculation on the Cray Research supercomputer through simple menus
☐ Monitors the progress of the calculation running on the Cray Research supercomputer
☐ Graphically analyzes the computed data from the calculation

The quantum chemical programs that reside on the Cray Research supercomputer occupy the other side of the UniChem window. These are

☐ MNDO90, a semiempirical program developed by W. Thiel at the University of Wuppertal, Germany
☐ DGauss, a density functional program using Cartesian Gaussians developed by Cray Research
☐ CADPAC, an ab initio Hartree-Fock program developed by R. Amos and coworkers at the University of Cambridge, England

These quantum chemical codes provide an unprecedented range of scientific coverage through a single interface. The UniChem environment provides both the simplicity and speed of semiempirical methods and the complexity of density functional and Hartree-Fock methods. UniChem can perform calculations on molecules that range in size from a few atoms to 500 or more. By integrating academic third-party codes, such as CADPAC and MNDO90, into the UniChem architecture, Cray Research can make the latest computational tools available to the largest possible audience.

The chemist at the workstation requires seamless two-way connectivity to the Cray Research supercomputer, because the input data are generated on the workstation; the calculations are sent from the workstation to the Cray Research supercomputer; and the results and data files generated by the calculations are transmitted back to the workstation (Figure 1).

Fail-safe mechanisms are needed to prevent loss of data in case of a network disruption. In addition, all network traffic must flow securely. UNIX operating system facilities such as *rsh* and *rcp* therefore were insufficient for the UniChem environment. The UniChem development team developed an Application Integration Toolkit (AIT) for use instead, within which communication between the workstation and the Cray Research supercomputer is handled by a passive daemon, called RCAgent, running under *root* mode. All UniChem jobs communicate with RCAgent. The quantum chemistry codes have routines that send data to the agent, which in turn sends the data to the appropriate workstation. The converse is also true; the user can request that the agent perform certain functions in relation to a specific

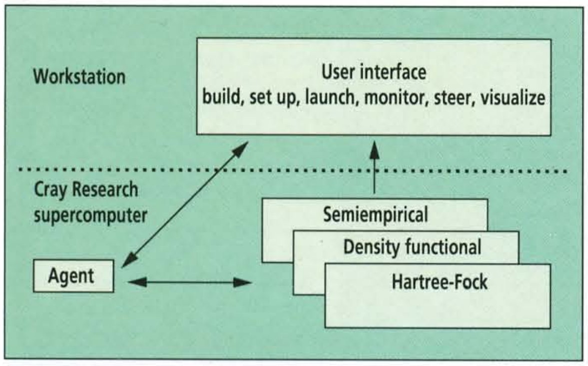

Figure 1. Outline of the UniChem architecture.

The UniChem interface includes point-and-click menus such as the density-functional menu shown below.

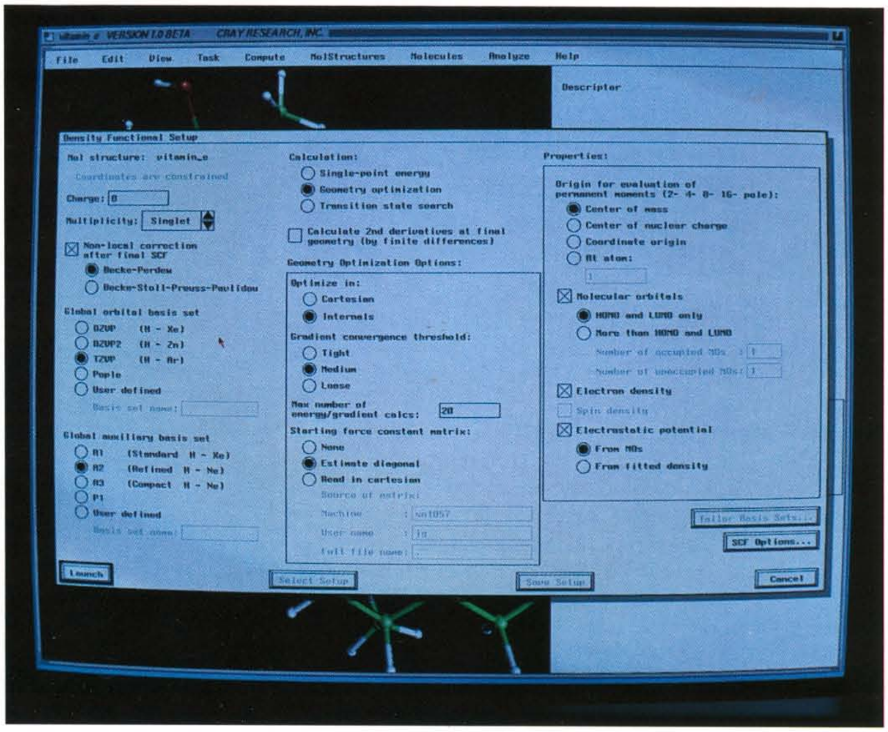

Density Functional Setup

Mol structure: vitamin_e
Coordinates are constrained
Charge: 0
Multiplicity: Singlet
☒ Non-local correction after final SCF
● Becke-Perdew
○ Becke-Stoll-Preuss-Pavlidou
Global orbital basis set
○ DZVP (H - Xe)
○ DZVP2 (H - Zn)
● TZVP (H - Ar)
○ Pople
○ User defined
Basis set name:
Global auxiliary basis set
○ A1 (Standard H - Xe)
● A2 (Refined H - Ne)
○ A3 (Compact H - Ne)
○ P1
○ User defined
Basis set name:

Calculation:
○ Single-point energy
● Geometry optimization
○ Transition state search
☐ Calculate 2nd derivatives at final geometry (by finite differences)
Geometry Optimization Options:
Optimize in:
○ Cartesian
● Internals
Gradient convergence threshold:
○ Tight
● Medium
○ Loose
Max number of energy/gradient calcs: 20
Starting force constant matrix:
○ None
● Estimate diagonal
○ Read in cartesian
Source of matrix:
Machine : sn1057
User name : lg
Full file name: .

Properties:
Origin for evaluation of permanent moments (2- 4- 8- 16- pole):
● Center of mass
○ Center of nuclear charge
○ Coordinate origin
○ At atom:
1
☒ Molecular orbitals
● HOMO and LUMO only
○ More than HOMO and LUMO
Number of occupied MOs : 1
Number of unoccupied MOs: 1
☒ Electron density
☐ Spin density
☒ Electrostatic potential
● From MOs
○ From fitted density

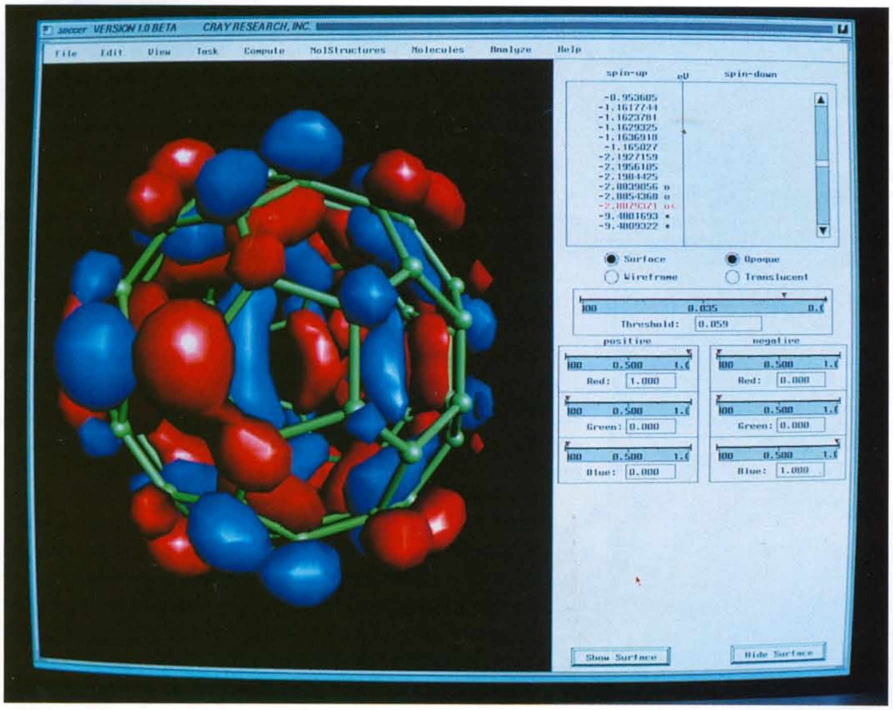

```
-0.953685
-1.1617744
-1.1623781
-1.1629325
-1.1636918
-1.165027
-2.1927159
-2.1956185
-2.1984425
-2.8039056 o
-2.8054368 o
-2.8079321 o<
-9.4001693 •
-9.4009322 •
```

Threshold: 0.059
positive: Red: 1.000, Green: 0.000, Blue: 0.000
negative: Red: 0.000, Green: 0.000, Blue: 1.000

Lowest unoccupied molecular orbital (LUMO) for the C-60 molecule (Buckminsterfullerene) modeled with UniChem.

calculation. Two such functions have been designed for the agent: the monitor function, which sends the current copy of the monitor buffers of a specific calculation back to a requesting workstation (these buffers contain information about the progress of the calculation and are updated continually by the application itself) and the steering function, which sends requests (Connect, Kill, Stop after next checkpoint, Pause, and Resume) directly to the application. Monitoring information also includes the calculation's current CPU time, SSD storage-device usage, and memory size, in addition to the chemistry information (for example, the current energy, cycle of the geometry optimization, or gradient components). Each of the quantum chemical codes has routines that check for incoming agent requests and respond accordingly. Data files are transferred from the application via *ftp* for TCP/IP connections or Cray Research station software. All file transfers are transparent to the user and use either destinations and file names supplied by the user or the UniChem system defaults based on the job name.

Methods

The chemistry applications, based on quantum mechanics, reside on the Cray Research supercomputer. Each application solves the nonrelativistic electronic Schrödinger equation, but with a different approximation. The codes work under the Born-Oppenheimer approximation, which allows the separation of nuclear and electronic motions. Conceptually, the simplest of these codes is the semiempirical code MNDO90, which treats only the valence electrons and uses various approximations to reduce the complexity of the theoretical problem. At its heart is the neglect (setting to zero) of terms in the equation that describe the energy of the system. The neglected terms involve a subset of the electron-electron repulsion terms. When these terms are set to zero, the equation becomes unbalanced and favors attractive terms. Experimental data are introduced to balance the remaining terms in the energy expression. This is the empiricism of the method. Ionization potentials are one type of experimental data often used in the parameterization, or fitting, of the energy equation. MNDO90 offers a series of semiempirical methods: MINDO, CNDO, MNDO, MNDO/H, AM1, and PM3. Semiempirical methods are fast because the energy equation has been simplified by neglecting some of the computationally intensive electron-electron repulsion terms. Because the methods scale as N^3, where N is the number of atoms, quantum mechanical calculations on very large systems are possible. A 540-atom carbon cluster recently ran on a CRAY Y-MP system at over 2.2 GFLOPS using a parallelized version of MNDO90.

The introduction of experimental data into a theoretical setting has its drawbacks, however. If the molecule in question has an unusual electronic structure and therefore does not fall within the domain of the experimental data used in the parameterization of the method, then a more rigorous approach is necessary. UniChem includes two ab initio methods: density functional and Hartree-Fock. Although each of these methods is based on first principles, each solves the electronic Schrödinger equation differently. Density functional methods produce exact solutions to an approximate form of the energy equation, whereas Hartree-Fock methods produce approximate solutions to an exact form of the energy equation. As the approximations are removed, the results of the methods converge.

In density functional theory, the energy equation is approximate due to the presence of an exchange-correlation potential. This exchange-correlation potential is a quantum effect and describes the instantaneous electron-electron repulsion; it differs from the Coulomb, or electrostatic, repulsion. Unlike Hartree-

The UniChem environment also can be used for nonpharmaceutical application. Shown here is a materials science application: a model of the lowest unoccupied molecular orbital of a titanium complex from catalysis research.

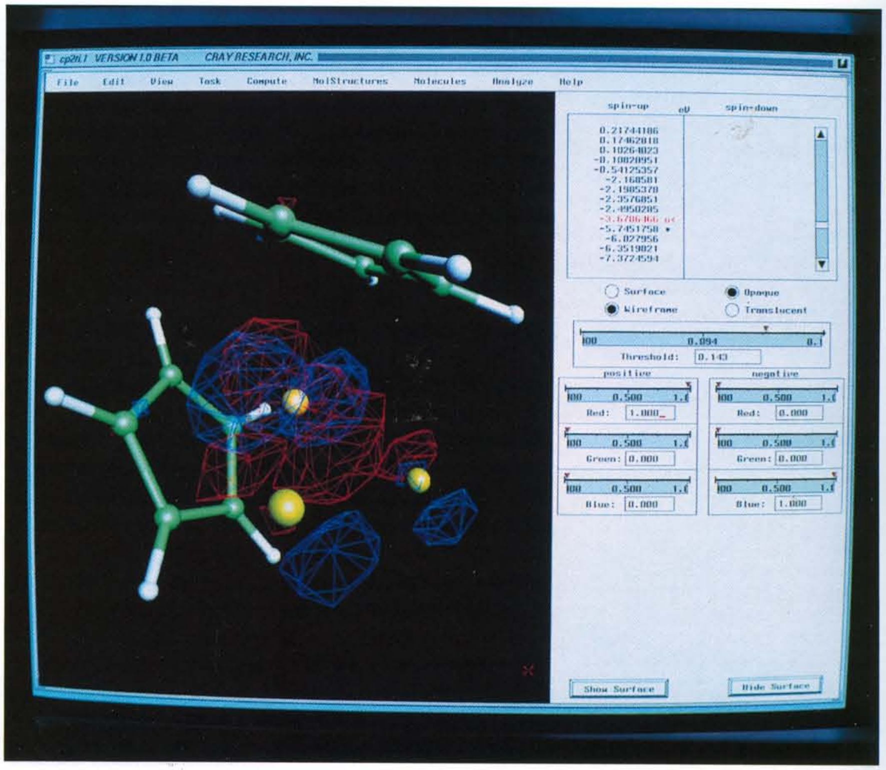

```
0.21744186
0.17462818
0.10264023
-0.10828951
-0.54125357
-2.168581
-2.1985378
-2.3576851
-2.4950285
-3.6786466 o<
-5.7451758 •
-6.827956
-6.3519821
-7.3724594
```

Threshold: 0.143
positive: Red: 1.000, Green: 0.000, Blue: 0.000
negative: Red: 0.000, Green: 0.000, Blue: 1.000

Fock theory, described below, density functional methods such as DGauss have the electron density, not the molecular orbitals, as their conceptual centerpiece. DGauss solves for the electron density using a local density approximation that sometimes overemphasizes electron correlation, yielding an electronic structure that has the electron density too delocalized over the molecular framework. By applying a nonlocal correction, DGauss can compensate for this error, leading to results typical for correlated Hartree-Fock methods. Systems with many electrons that are known to be sensitive to electron correlation, such as organometallic compounds and transition metal complexes, are now open to successful computation. Density functional methods scale as N^3, where N is the number of basis functions, thus making it possible to run first-principle calculations on large systems, where the number of basis functions exceeds one thousand, which translates to over 150 atoms. With density functional theory, chemists have, for the first time, a first-principles method that allows them to treat entire molecules, not just portions or models of molecules.

In UniChem, ab initio Hartree-Fock theory is accessed through the CADPAC program. Within Hartree-Fock theory, the electronic energy equation is exact, but the wavefunctions used to solve it are not. These wavefunctions are a Slater determinant (an antisymmetrized product) of linear combinations of one-electron hydrogen-like atomic orbitals represented by Cartesian Gaussians (called the basis set). We are using a collection of one-electron functions to describe a many-electron problem — an approximation. Solving the Hartree-Fock equations yields the energy and wavefunctions — the molecular orbitals. But because Hartree-Fock theory does not treat electron correlation correctly, it yields molecular orbitals that are less delocalized than those from density functional methods. However, Hartree-Fock theory can be extended easily to obtain accurate electron-correlation treatments. With CADPAC, users can choose from a wide range of correlated methods. These include Moller-Plesset perturbation theory, coupled cluster doubles, and configuration interaction with single and double excitations from the reference state. Of the three computational codes present in UniChem, these correlated methods offer the greatest accuracy, but at a cost. The computational demands of many of these correlated methods scale as N^5 or greater, where N is the number of basis functions. Such a scaling law prevents correlated methods from being applied to large systems.

In addition to its collection of correlated methods, CADPAC offers a rich set of standard Hartree-Fock-based properties. These range from permanent moments to dipole-dipole polarizability to IR and Raman intensities, and harmonic frequencies. The latest version of CADPAC uses a direct self-consistent-field (SCF) approach for large molecular systems that eliminates the need to write very large two-electron integral files. For moderately sized systems, these Hartree-Fock methods scale as N^4, where N is the number of basis functions, but reduce to $N^{2.5}$ or N^3 for larger systems.

The UniChem environment not only provides the chemist with a common interface for MNDO90, DGauss, and CADPAC, but also offers many of the same types of calculations within each method. All of the chemistry codes can perform geometry optimizations (searching for minima or transition states) and calculate second derivatives and dipole moment derivatives. In the UniChem environment, a computational chemist can perform a calculation in the semiempirical framework, watch for the convergence of the results, then transfer smoothly to the more rigorous first-principles methods. The chemist can take the geometry and second derivatives from any level of theory and use them as input to another level. UniChem also provides common analysis and display tools such as the visualization of electrostatic potentials, molecular orbitals, total electron density, electronic populations, permanent moments, polarizabilities and vibrational spectra.

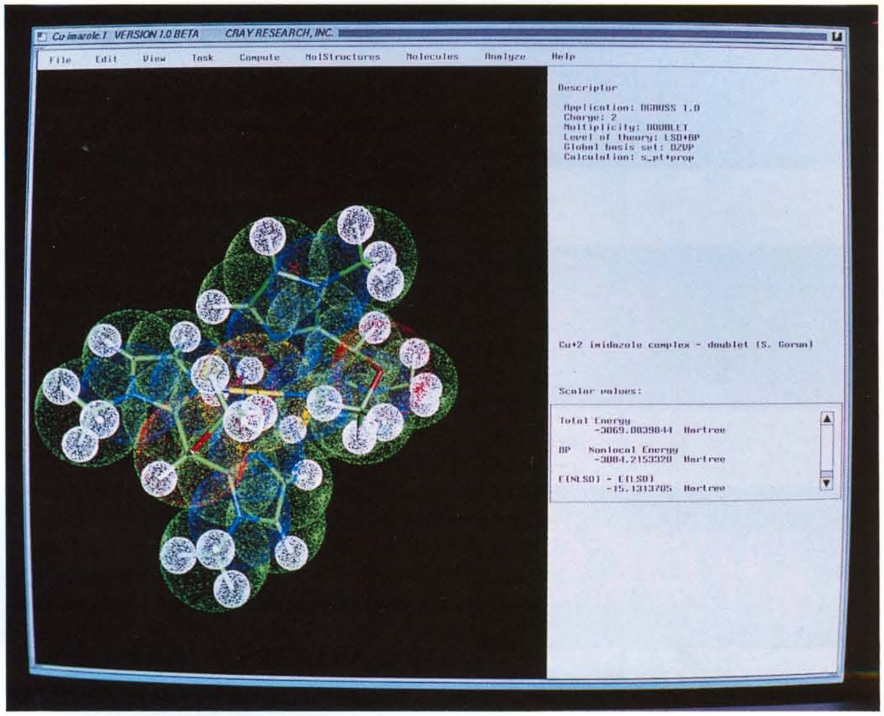

```
Descriptor

 Application: DGAUSS 1.0
 Charge: 2
 Multiplicity: DOUBLET
 Level of theory: LSD+BP
 Global basis set: DZVP
 Calculation: s_pt+prop

Cu+2 imidazole complex - doublet (S. Gorun)

Scalar values:

Total Energy
      -3069.0839844  Hartree

BP   Nonlocal Energy
      -3084.2153320  Hartree

E(NLSD) - E(LSD)
        -15.1313705  Hartree
```

Density functional model of a copper imidazole complex displayed with a Van der Waals surface.

UniChem II

Cray Research is continuing to respond to the needs of its chemistry customers with the creation of the UniChem II consortium. As much as the UniChem I consortium concentrated on lifting the barriers to computational chemistry, UniChem II will develop tools for correlating multiple molecules and properties. The consortium will help set the direction of UniChem II and help formulate its requirements. By integrating third-party tools and programs into UniChem II, Cray Research will provide industrial computational chemists with leading-edge technology. Cray Research intends to provide computational chemists with not only the best hardware, but also the most productive software environment possible. ▬

About the author

Jeffrey F. Gaw is the chemistry architect of the UniChem development group. He received his Ph.D. degree in chemistry from the University of California, Berkeley, and postdoctoral training in the Department of Theoretical Chemistry at Cambridge University.

Improving chemical process design with supercomputers

Alan B. Coon, Union Carbide Technical Center, South Charleston, West Virginia
Christopher H. Goheen and Mark A. Stadtherr, University of Illinois, Urbana, Illinois

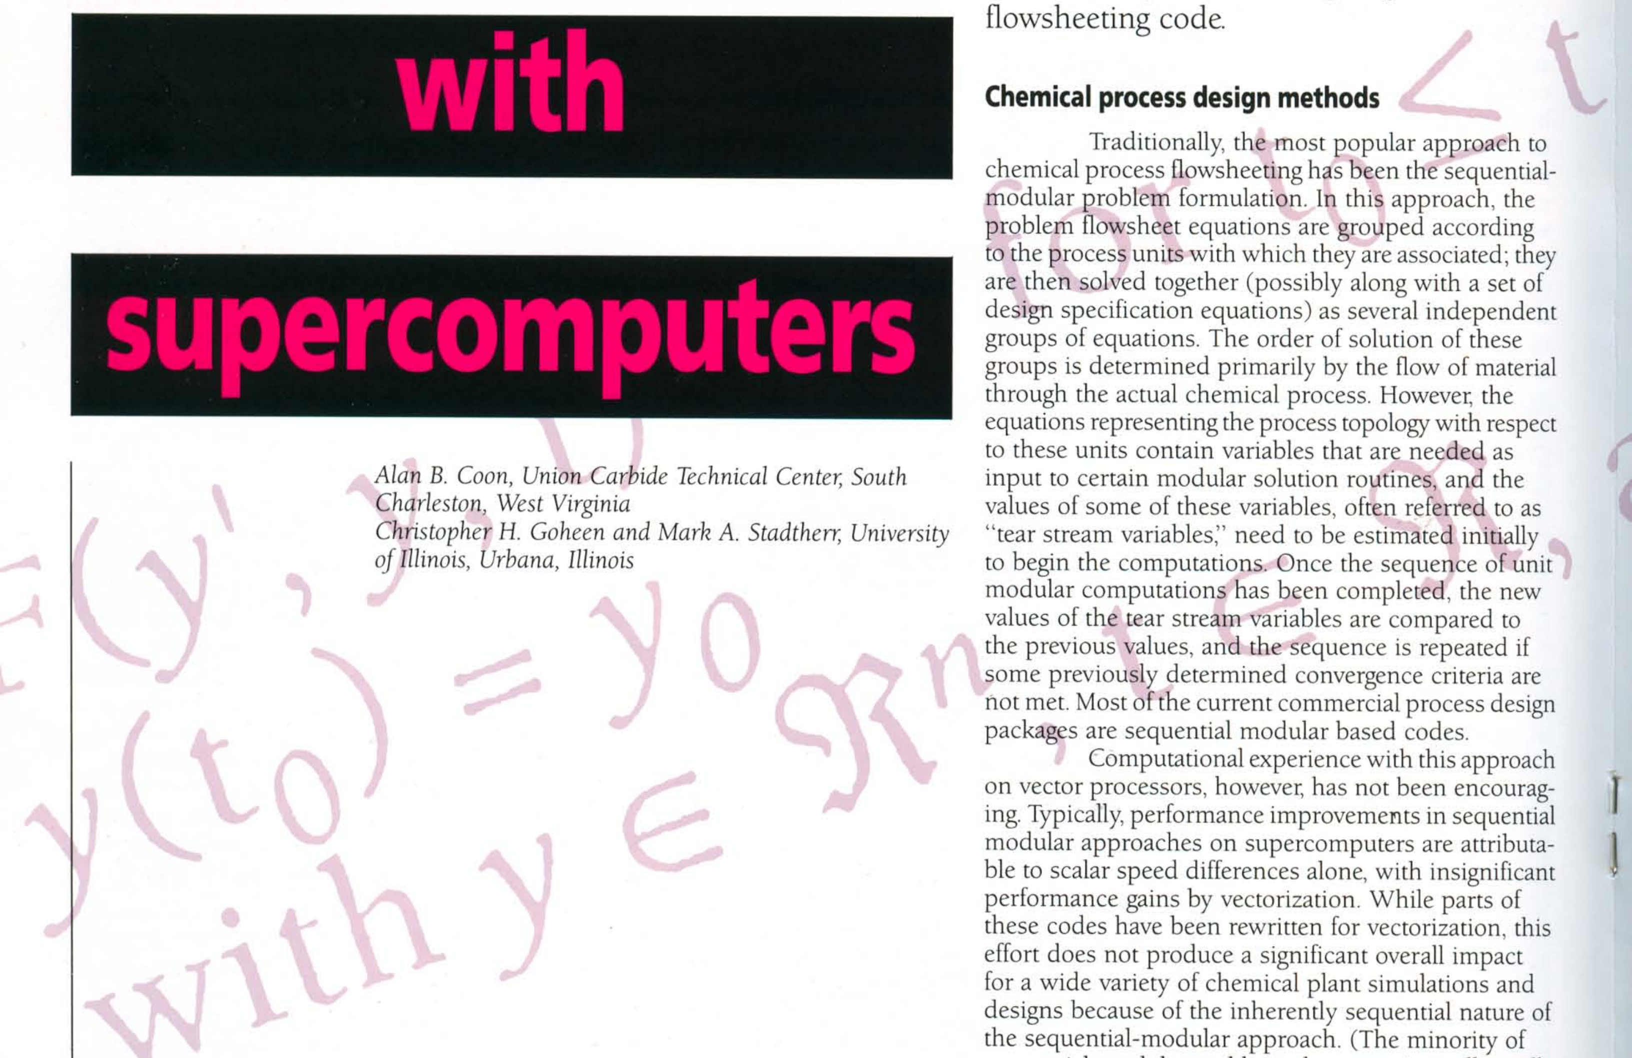

Although the equation-based (EB) methodology is the most attractive approach to chemical process design and simulation on supercomputers, it often results in a computational bottleneck. Successful implementation of the EB approach requires a linear solver that makes efficient use of parallel and vector architectures. This efficiency is critically dependent on the reordering phase of the equation-solving step of the simulation when direct sparse solvers are used. Therefore, the reordering algorithm is a crucial factor in exploiting the supercomputer's resources. A new reordering algorithm was teamed with a CRAY-2 computer system at the National Center for Supercomputing Applications to produce significant improvements in the linear equation solving step for an EB flowsheeting code.

Chemical process design methods

Traditionally, the most popular approach to chemical process flowsheeting has been the sequential-modular problem formulation. In this approach, the problem flowsheet equations are grouped according to the process units with which they are associated; they are then solved together (possibly along with a set of design specification equations) as several independent groups of equations. The order of solution of these groups is determined primarily by the flow of material through the actual chemical process. However, the equations representing the process topology with respect to these units contain variables that are needed as input to certain modular solution routines, and the values of some of these variables, often referred to as "tear stream variables," need to be estimated initially to begin the computations. Once the sequence of unit modular computations has been completed, the new values of the tear stream variables are compared to the previous values, and the sequence is repeated if some previously determined convergence criteria are not met. Most of the current commercial process design packages are sequential modular based codes.

Computational experience with this approach on vector processors, however, has not been encouraging. Typically, performance improvements in sequential modular approaches on supercomputers are attributable to scalar speed differences alone, with insignificant performance gains by vectorization. While parts of these codes have been rewritten for vectorization, this effort does not produce a significant overall impact for a wide variety of chemical plant simulations and designs because of the inherently sequential nature of the sequential-modular approach. (The minority of sequential-modular problems that vectorize well usually include one process module that is modeled by solving a very large set of simultaneous equations; solution of this module dominates the overall computation time, and in this case the sequential modular formulation is actually more equation oriented than module oriented.) The potential for using this approach on parallel processing machines is also significantly limited. Successful parallelization of code, like successful vectorization, requires that there be some independence in the order of computation.

In the equation-based approach to chemical process flowsheeting, the problem equations describing both process units and topology are solved simultaneously, rather than dividing the problem into smaller unit modular subproblems to be solved sequentially. This approach allows design specifications and constraints to be incorporated easily in the problem. It also significantly increases computational efficiency and enhances flexibility in the types of problems that can be solved. In the general case, equation-based process flowsheeting requires the solution of large, sparse differential-algebraic equation (DAE) systems. These systems result from the algebraic equations, the partial differential equations (PDE) (the presence of which usually requires a discretization of the spatial domain), and ordinary differential equations (ODE) that describe the various unit processes, the process topology, and design specifications. The general DAE problem can be written as

$$F(y', y, t) = 0 \text{ for } t_0 \le t \le t_f \qquad (1)$$

$$y(t_0) = y_0 \qquad (2)$$

with $y \in \Re^n$, $t \in \Re$, and $F : D \subset \Re^n \times \Re^n \times \Re \rightarrow \Re^n$

Here F, the nonlinear flowsheet equations, and y, the time dependent variable set including both differential and algebraic variables, are assumed to be sufficiently differentiable. Equations 1 and 2 represent an implicit ODE system if the Jacobian ($\partial F/y'$) is nonsingular for $t_0 \le t \le t_f$. Otherwise, ($\partial F/y'$) is singular for some $t_0 \le t \le t_f$, in which case the reduction of the DAE to an ODE requires some subset of Equation 1 be differentiated. In either of these two cases, and provided that a set of consistent initial conditions can be established, the system in Equations 1 and 2 is reduced to the form

$$f(y) = 0 \qquad (3)$$

with $y \in \Re^n$, and $f : S \subset \Re^n \rightarrow \Re^n$

and the nonlinear system of equations represented by Equation 3 remains to be solved at various time steps. This set of equations is linearized, and the resulting sparse Jacobian matrix (more often, an approximation of this matrix) is solved as part of the nonlinear solution method. The resulting linear equation set can be written

$$Ax = b \qquad (4)$$

where $A = (\partial f/\partial y)|_{y = y_k}$, or some appropriate approximation

$x = y_{k+1} - y_k$

$b = -f(y_k)$

The solution of this large, sparse system of equations is the single most computationally intensive step in the overall flowsheet solution process. Equation 4 usually is solved several times for each nonlinear equation solution, with the coefficients in matrix A updated at each iteration, and the process continuing until some convergence criterion is satisfied. The nonlinear equation-solving step, in turn, must be executed for each time step in the simulation. In some problems, the linear solution routine may require nearly 90 percent of the overall flowsheeting computation time.

Direct sparse solvers

Most general-purpose direct sparse linear solution routines use some type of indirect addressing, which can degrade vector performance drastically. In these instances, the use of hardware gather/scatter in vector machines can be quite effective. Another approach is to use variable bandwidth solvers, thereby eliminating the need for indirect addressing at the expense of executing vector operations that may include a great deal of zero-operand arithmetic. The unnecessary operations occur because these methods treat all coefficients between the bandwidth boundaries as nonzeros. While this approach enhances vectorization, popular measures of such enhancement, such as MFLOPS rating, can be misleading because the overall solution will be quite inefficient if the bandwidth is not kept reasonably small.

The frontal method has been a popular means of exploiting sparse matrix structure on vector processors, because the difficulties inherent in indirect addressing on vector processors can be circumvented in much the same way as with variable bandwidth solvers. However, the frontal method and its variants tend to work well only for unsymmetric matrices that can be reordered to a nearly structurally symmetric form. One approach to improving performance of the frontal method on EB matrices that cannot be reordered easily to a nearly structurally symmetric form is to reorder them to a block structure that is nearly structurally symmetric. The reordering algorithm considered here was originally developed for a coarse-grained parallel sparse solver that can exploit the natural block structure of EB flowsheeting matrices.[1] This generalized block tridiagonal reordering produces a symmetric block structure. This structure is achieved through a recursive two-way partitioning of the vertices in the underlying bipartite graph of the sparse matrix. For a good frontal method ordering, the frontal matrix size must be kept small. This can be achieved by minimizing the number of graph edges that are cut by the partition.

Graph partitioning algorithms have several applications, most notably VLSI placement and routing and computer program paging, and they are used in various divide-and-conquer strategies. As a result, many algorithms for finding graph partitions that minimize the size of the cutset (number of edges cut) have been proposed over the past two decades. However, the problem of partitioning the vertex set into two sets of the same order with the minimum number of nets cut is NP-complete, and consequently these algorithms attempt to select a good approximation of such a partition using various heuristics. EB flowsheeting matrices present a further complication, however, for these matrices are usually highly nonsymmetric and their sparsity is not well approximated by that of their symmetric parts. Therefore, a bipartite graph model must be used to include the asymmetry of the matrix. In spite of this additional complexity, a linear time-partitioning algorithm for unsymmetric EB flowsheeting matrices has been developed which drastically improves

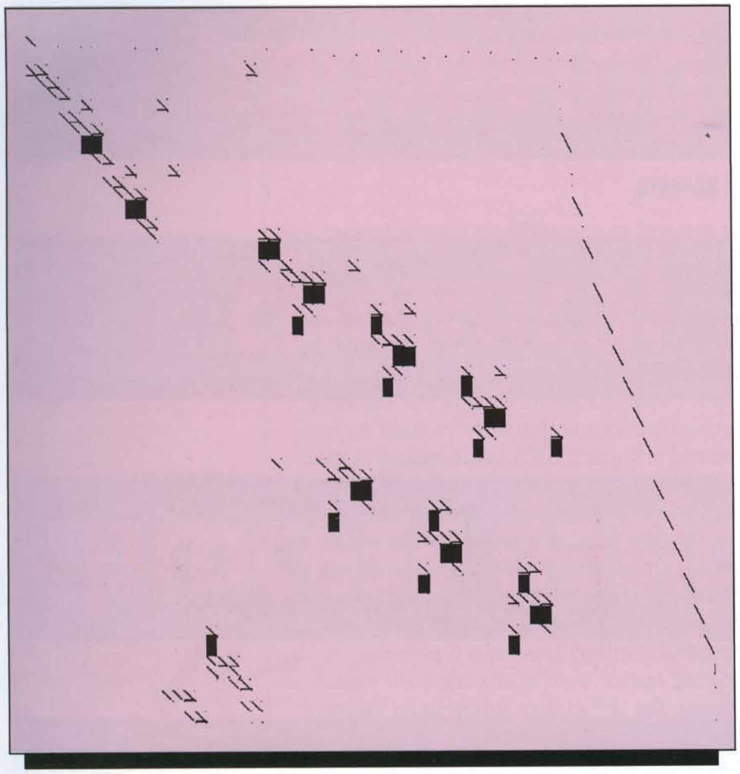

Figure 1. (above). Occurence matrix for Problem 1 before new reordering.

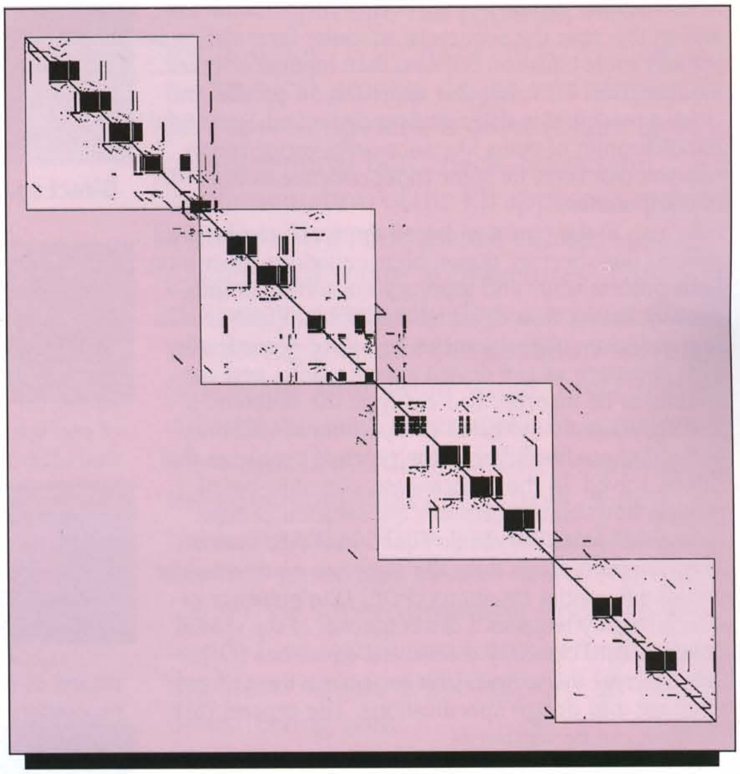

Figure 2. (above right). Occurence matrix for Problem 1 after new reordering.

a CRAY-2 system implementation of an EB flowsheeting code.[2]

Test problems

Two flowsheeting problems were used to test the effectiveness of two new variants of the reordering algorithm, along with the original reordering.[3] The two new orderings differ in the amount of imbalance allowed in partitioning the underlying graph of the matrix. SEQUEL-II, a prototype EB flowsheeting code developed at the University of Illinois and implemented on the CRAY-2 system at the National Center for Supercomputing Applications, was used to generate sets of equations for these problems.[4,5] The first flowsheeting test problem was a light hydrocarbon recovery process with 20 chemical components. The SEQUEL-II formulation required 28 units, 48 streams, 1477 equations, and 18,592 nonzero coefficients. The second flowsheeting problem was a simulation of a natural gas processing plant with 20 components. The SEQUEL-II formulation for this process consisted of 21 units, 39 streams, 1235 equations, and 16,868 nonzero coefficients. Both problems were quite sparse, with a nonzero density of approximately 1 percent. For the numerical experiments on these test problems, two versions of the frontal method that differ in the rank of the update to the frontal matrix were used.[3] Both the rank-one and the rank-four update routines were coded in Cray Assembly Language for increased efficiency.

Results

Comparisons on the CRAY-2 system using SEQUEL-II show substantial improvement with the new partitioning scheme over the original reordering

Table 1. Maximum frontal matrix sizes and total linear solution times with rank-one update.

			Original reordering (reverse P4)		New (balanced) reordering		New (unbalanced) reordering	
Pblm	Order	Nonzeros	Maximum frontal size	Total linear time (sec)	Maximum frontal size	Total linear time (sec)	Maximum frontal size	Total linear time (sec)
1	1477	18,592	60 x 942	6.49	111 x 235	4.38	111 x 236	4.48
2	1235	16,868	239 x 887	25.48	130 x 217	5.66	109 x 177	5.44

Table 2. Maximum frontal matrix sizes and total linear solution times with rank-four update.

			Original reordering (reverse P4)		New (balanced) reordering		New (unbalanced) reordering	
Pblm	Order	Nonzeros	Maximum frontal size	Total linear time (sec)	Maximum frontal size	Total linear time (sec)	Maximum frontal size	Total linear time (sec)
1	1477	18,592	60 x 944	4.68	114 x 235	3.31	111 x 236	3.61
2	1235	16,868	241 x 888	15.37	132 x 219	4.68	109 x 177	3.89

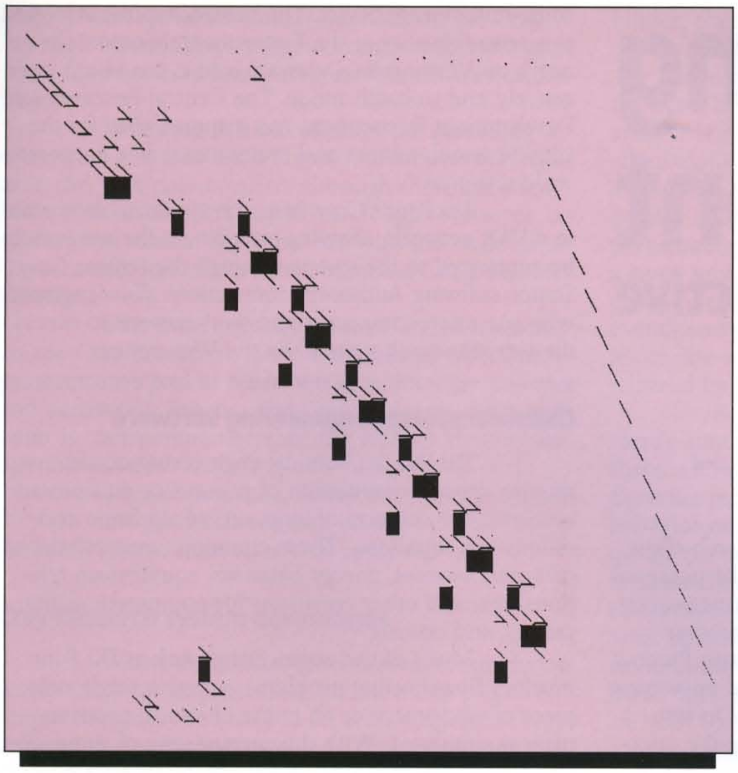

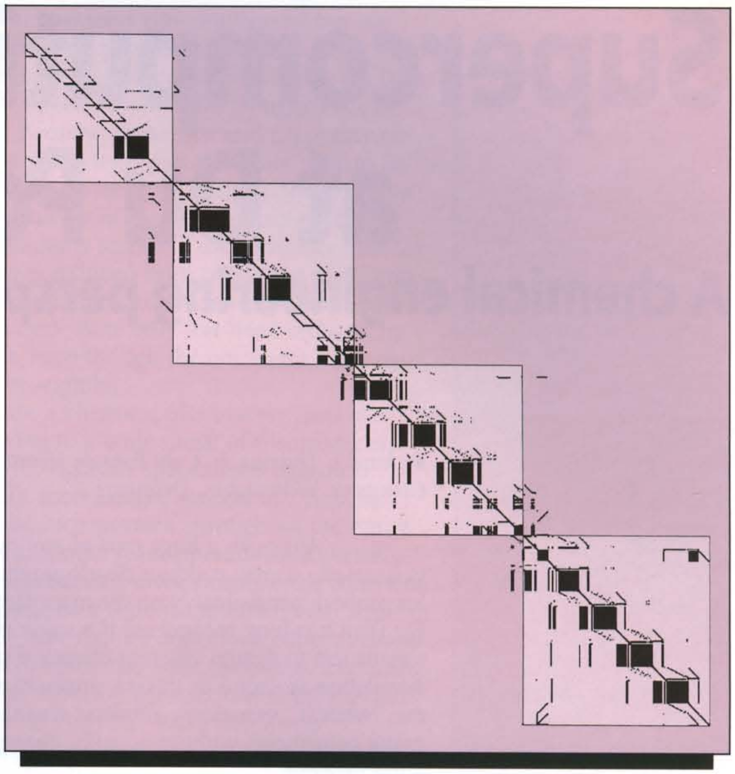

Figure 3. (above left). Occurence matrix for Problem 2 before new reordering.

Figure 4. (above). Occurence matrix for Problem 2 after new reordering.

used with the frontal method in SEQUEL-II. The frontal matrix sizes given by the new reordering were significantly smaller than those given by the original. Table 1 indicates that balanced partitioning is slightly better than the unbalanced partitioning for Problem 1, and the reverse for Problem 2. Table 2 shows the improvement in performance gained by switching to a rank-four update resulting from a more efficient use of the single path to memory on the CRAY-2 system. Figures 1-4 are the occurence matrices for both problems, before and after reordering.

Conclusions

Initial experience with the new generalized block tridiagonal orderings indicates that substantial improvements in the linear solution stages of equation-based process simulation are possible even for methods that already vectorize well on supercomputers. In the case of the frontal method, the frontal matrix sizes and amount of floating-point arithmetic involving zero-valued operands were strongly dependent on the method used to reorder the matrix. Based on this finding, the equation-based simulator was improved by developing new algorithms for the reordering, a stage that represents an insignificant fraction of the overall solution time.

Equation-based chemical process flowsheeting codes are the appropriate choice for vector and parallel machines, because the potential for concurrency in the traditional sequential modular method is inherently limited. However, efficient use of such advanced computer hardware also depends on a suitable choice of algorithms for the equation-based methodology. This study demonstrates that the appropriate combination of methodology and algorithms can fully exploit the power of a supercomputer such as the CRAY-2 system.

About the authors

Alan B. Coon is a senior chemical engineer at Union Carbide Corporation in the Applied Mathematics and Processs Simulation Group. He received his B.S.Ch.E. degree from the University of Texas at Austin, and his M.S. and Ph.D. degrees in chemical engineering from the University of Illinois at Urbana-Champaign.

Christopher H. Goheen received his B.Ch.E. degree at Georgia Tech and his M.S.Ch.E. degree at the University of Illinois. He is completing Ph.D. studies in chemical engineering at the University of Illinois. He will join the Technical Computing Section of Mobil Research and Development Corporation in Princeton, New Jersey, in the fall of 1991.

Mark A. Stadtherr is a professor at the University of Illinois at Urbana-Champaign in the Department of Chemical Engineering. He also works with the National Center for Supercomputing Applications. He received his B.Ch.E. degree from the University of Minnesota and his Ph.D. degree from the University of Wisconsin.

References

1. Coon, A. B. and M. A. Stadtherr, "Parallel Sparse Matrix Solvers for Equation-based Process Flowsheeting," *Institution of Chemistry Engineering Symposium Series,* Vol. 114, pp. 165-176, 1989.
2. Coon, A. B., "Orderings and Direct Methods for Coarse Granular Parallel Solutions in Equation-based Flowsheeting," Ph.D. Thesis, University of Illinois, Urbana, Illinois, 1989.
3. Zitney, S. E., "Frontal Algorithms for Equation-based Chemical Process Flowsheeting on Vector and Parallel Computers," Ph.D. Thesis, University of Illinois, Urbana, Illinois, 1989.
4. Stadtherr, M. A. and C. M. Hilton, "Development of a New Equation-based Process Flowsheeting System: Numerical Studies," *Selected Topics on Computer-Aided Process Design and Analysis,* R. S. H. Mah and G. V. Reklaitis, Eds., AIChE Symposium Series, Vol. 78, pp. 12-28, 1982.
5. Zitney, S. E. and M. A. Stadtherr, "Computational Experiments in Equation-based Chemical Process Flowsheeting," *Computers and Chemical Engineering,* Vol. 12, pp. 1171-1186, 1988.

Supercomputing at Du Pont

A chemical engineering perspective

Richard J. Thomas, E. I. du Pont de Nemours and Company, Wilmington, Delaware

Although a large part of any research effort is focused on new product development, the potential for profit depends highly on the manufacturing process. Du Pont has long recognized the value of process simulation to design efficient chemical plants. Process simulation is useful in design, analysis, and answering the "what if" questions, allowing engineers to test many conditions without actually changing the operating process.

In 1988, Du Pont purchased a CRAY X-MP computer system to be used primarily for computational chemistry. However, benchmarks also were run against a VAX 8800 on two process simulation packages: PROCESS, a commercially available flowsheet simulation package from Simulation Sciences; and CPES, a software program developed by Du Pont. Of the 34 real-world simulations, 28 were run on PROCESS and 6 were run on CPES. Results showed a three- to tenfold speed advantage on the Cray Research system with costs being one-half to one-third less than the VAX system. Du Pont realized it was a cost-effective tool in complex simulations and day-to-day process modeling. Now that chemical engineers at Du Pont have been using a Cray Research supercomputer for more than four years, Du Pont wanted to find out what, if any, advantages the Cray Research system contributed to the chemical engineers' work. This article discusses the engineers' responses and gives their suggestions for future directions in supercomputing capabilities.

Supercomputing environment at Du Pont

Scientific computation at Du Pont is primarily distributed. Several Du Pont sites have a VAX-class machine and many have a VAXcluster of machines. Most of these machines are networked, and electronic mail, data, and input files used in process simulation can be exchanged easily. Scientific and/or technical software is installed at several sites as well. Users who do not have a computer on site can access the software via communications to and from the Du Pont location. The result is a very flexible system facilitating the transfer of electronic information.

Du Pont's current Cray Research system is a CRAY Y-MP4/332 supercomputer consisting of three CPUs and 32 Mwords of memory with 41 gigabytes on disk drives. It is supplemented by a 128-Mword SSD solid-state storage device. The supercomputer is located at a central facility at the Experimental Station site north of Wilmington, Delaware, and is used both interactively and in batch mode. The Central Research and Development Department has responsibility for the Cray Research system and maintains it as a corporate-wide resource.

Du Pont's Cray Research system is connected to a VAX network, allowing any site on the network to be connected to the system through the remote Cray Station software. Additional connections allow engineers who use microcomputers and workstations to access the Cray Research system via the VAX system.

Chemical process engineering software

Du Pont's chemical engineering needs require rigorous simulation of processing equipment involving the solution of large sets of algebraic and differential equations. These equations may consist of material balances, energy balances, equilibrium relationships, and other equations for equipment sizing, ratings, and costing.

Much of the engineering work at Du Pont involves flowsheeting problems, where a single unit piece of equipment or an entire chemical process plant is simulated. With the current state of simulation technology and the many needs of the diversified business groups, these modeling requirements cannot be met by any one simulator. Du Pont's engineers and consultants can choose from the following software packages to solve their flowsheeting problems:

- ☐ PROCESS and PRO/II, products of Simulation Sciences, Inc., Fullerton, California
- ☐ ASPEN PLUS, a product of Aspen Technology, Inc., Cambridge, Massachusetts
- ☐ CPES, software written by Du Pont
- ☐ SPEEDUP, a product of Prosys Technology Ltd. in Cambridge, England

Feedback from engineers cited later in this article applies to all the above software packages except SPEEDUP, which has since become available for use on Cray Research systems. (See Applications Update in this issue.)

Making the Cray Research system transparent

To encourage users to take advantage of Du Pont's supercomputer resources, the Central Research and Development group formed a small Cray Research project team with the VAX-based engineers to make simulations on the Cray Research system as easy as possible to run. The project team decided that the interface to the Cray Research system should be as transparent as possible for engineers familiar with the VAX or PC environment.

In the VAX environment, running a simulation requires answering a series of questions. For example, to run a typical flowsheeting program (in this case, PROCESS) the user types the command PROCESS *filename,* where a file named *filename*.DAT exists in the directory. The user is then prompted with a series of questions that includes the type of simulation to run and a cost code to charge. The Cray Research project team created a similar interface on the Cray Research system. The analagous command for the same input file on a Cray Research system is CRAYPROCESS *filename*.

Following the execution of the preceding command, the user must respond to similar questions as in the VAX environment, although there are a few more of them on the Cray Research side. However, on subsequent runs of any input file with the same name, the user only has to issue the command CRAY SUBMIT *filename*.

Therefore, subsequent runs are even easier on the Cray Research system than on the VAX. The project team's goal of push-button technology was met and surpassed. Because the Cray Research supercomputer is transparent, users do not have to learn a new operating system and computer editor to run their jobs correctly the first time. Today there are essentially no more complaints about ease-of-use on the Cray Research system than on a VAX computer or PC.

Cray Research system advantages

When chemical engineers analyze simulation results of complex processing plants, many pieces of output data must be checked for reasonableness. Typically, an engineer must be able to determine a wide range of possible input variables. Quick transmission of large amounts of output data as well as moderately sized input files is therefore necessary. The combination of the Cray Research system and the fast intercomputer communications of the Du Pont VAX system was cited as a factor that has helped make the use of supercomputers for chemical engineering a success companywide. In general, Du Pont engineers believe the major advantage of the Cray Research system is speed, resulting in a variety of user benefits.

One engineer said he could finish more work in a given amount of time. Others said impressive schedule improvements resulted from their ability to run several jobs on the Cray Research system during the time it would take to run one job on the VAX or a tenth of a job on a PC. Some engineers thought their resulting designs were better because they could do more case studies in a given amount of time. For example, if a job takes three hours to run on the VAX, it can be run only two to three times per day. With the Cray Research system, however, 18 to 21 runs can be achieved in a normal eight-hour day.

Many engineers prefer the Cray Research system for even short runs of one to five minutes. Frequently, they do not even have time to edit another file while waiting for the results. Because of this quick response, they can maintain continuity of thought. This helps reduce the number of projects they must keep active simultaneously and allows them to complete projects sooner.

One engineer summarized the Cray Research system advantage as, "The essence of our business is timely computer simulations; therefore, the Cray Research system improves my effectiveness."

Future needs

With the speed advantage of the Cray Research system and the improved software calculation capabilities, chemical engineers can simulate all but the most unusual chemical processes in a reasonable time. The level of technology is now at a point where Du Pont is bound by the speed at which their engineers can analyze results. Because of the size and capability of simulations, many separate pieces of data have to be available for review. To improve designs, they need to be able to look quickly at the effect of input variable perturbations on many results. If the computer time for a given problem is reduced from two hours to 15 minutes and the analysis of results takes eight hours, the overall analysis cycle time only has been reduced by about one-sixth, even though the computer time was reduced by seven-eighths.

Because a chemical plant rarely runs at steady-state, owing to a wide range of disturbances, there is a need for rigorous simulations based on dynamic models. Increasingly, aspects of dynamic behavior must be incorporated throughout the entire design process to model plant operations accurately. Because these simulations must calculate the transient behavior of a plant over time using differential algebraic equations, they require enormous computational resources to solve industrial-scale problems. Executing these full process dynamic simulations is only practical with the use of powerful supercomputers.

Process optimization also is needed. Individual reactors, distillation columns, condensers and other pieces of equipment often have been objects of optimization programs, but the entire process is rarely, if ever, optimized. It is becoming clear that profit can be gained by ensuring that processes, even entire plants that include several processes, are run at some optimal point. Process or plant optimization requires that computer systems be able to run large simulation models at current plant conditions.

Often an important process variable is unavailable using current measurement and analytical techniques. A simulation model may be used to calculate a particular concentration, temperature, pressure, or flow condition, and this often can be done in the local process monitoring and control computer with simple simulation models, heuristic models, or statistical models derived from the full models. As networking capabilities improve, we anticipate that full models may be run on the Cray Research system using current process conditions to supply not only these process variables, but also to advise on current optimal performance set points. On-line use of full process simulation models is just beginning.

As chemical simulations become increasingly rigorous, more and more problems will require the power of supercomputers, such as the CRAY Y-MP system, to deliver accurate flowsheeting simulation, graphical analysis, dynamic modeling, and on-line optimization. Du Pont's engineers have found their Cray Research system to be a practical tool, delivering accurate answers to their most difficult questions faster than any other hardware in their network.

About the author

Richard Thomas is an engineering consultant in the Process Engineering section of Du Pont and is responsible for some large chemical plant simulations. He received his B.S.Ch.E. at the University of Maryland.

Computational brain dynamics

Visualization of activity flow in the cerebral cortex

Robert E. Wyatt and Jesse W. Driver, University of Texas, Austin, Texas

The cerebral cortex of the mammalian brain is the most complex structure on Earth. It possesses many classes of neurons that are interconnected in very intricate ways. The connections are both highly convergent — many neurons synapse onto any given cell — and highly divergent, in that each neuron contributes synaptic output to many other neurons. In addition, the signaling properties of the neurons themselves are highly nonlinear. A knowledge of the anatomical connectivity of the system and the signaling properties of its constituent neurons is insufficient to reveal the collective dynamic behavior and function of the cortical network. The dynamic behavior of systems of interacting neurons is also difficult to study experimentally.

In recent years, computational modeling has emerged as a powerful tool for understanding complex systems, including neural systems. A number of investigators have developed computational models of subcortical brain structures that incorporate a substantial degree of anatomical and physiological realism. These structures include the cerebellum, the hippocampus, the olfactory cortex, and the olfactory bulb. These models have been useful in elucidating the mechanisms of oscillatory and epileptic activity in the hippocampus. They also have suggested mechanisms for motor program storage in the cerebellum and distributed storage and information access in the olfactory cortex.

The computational study described here focused on the first large-scale dynamic analysis of signal flow in a slab of cortical tissue. The visual cortex, in particular, has been the focus of experimental studies for 100 years and is the subject of our model. The model used in this study was written using Fortran code and implemented on the CRAY Y-MP8/864 supercomputer at the University of Texas Center for High Performance Computing.

An accurate computational model of a patch of visual cortex would be useful for several reasons; it would help us understand the dynamics of cortical activity flow, the generation of cortical oscillations by recurrent excitation, and the emergence of the receptive field properties of cortical cells from their connectivity and physiological properties.[1-3] Among the challenges to overcome is the generation of a network whose connectivity incorporates available knowledge of the anatomy and physiology of mammalian visual cortex. Here we describe a method of generating an appropriate pattern of connections for a model of the vertical organization of visual cortex. The network generated by these methods is then subjected to stimulation of the input fibers, and the activity of the component neurons is followed by integrating a large system of coupled first-order nonlinear differential equations. The primary output is the activity of each neuron for hundreds of discrete time steps.

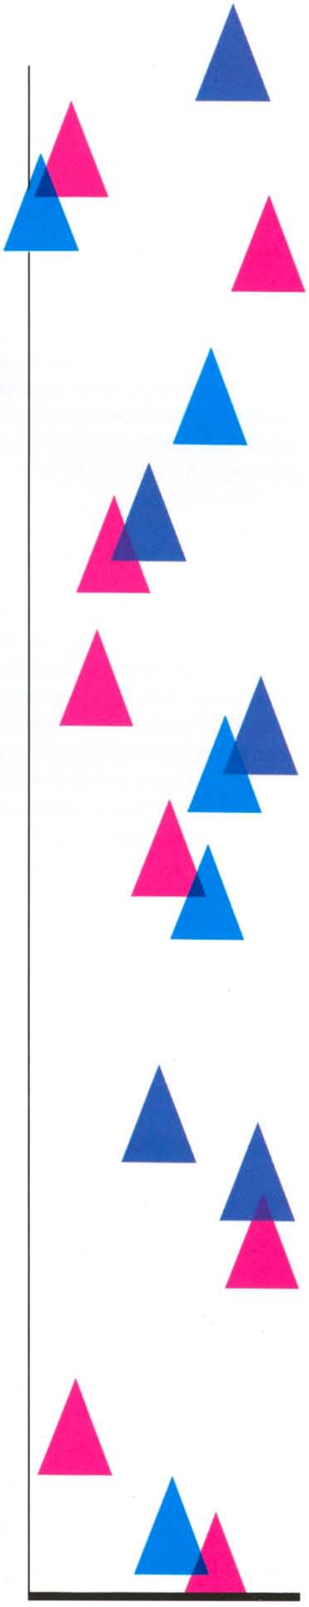

A model of cortical connectivity

The model is of a slab of cortical tissue bounded laterally by a square 0.5 mm on each edge and extending 1.5 mm from the surface of the brain down to the white matter. The cerebral cortex is a laminated structure with thickness and cell density that varies within the layers of different cortical areas. The laminae are numbered such that layer 1 is at the top or outer surface while layer 6 is at the bottom of the cortex, adjacent to the white matter. The thickness assigned each lamina and sublamina approximates experimental data for cat visual cortex. About 20,000 cortical neurons lie beneath this patch of cat visual cortex. This volume contains about 116 million synapses, or an average of 6000 connections per cell.

The cortex comprises many types of neurons, but for simplicity they can be divided into pyramidal cells (so named because of the shape of their cell bodies) and interneurons, which are nonpyramidal cells. The pyramidal cells provide for vertical signal flow within the cortex as well as producing output signals for other cortical or subcortical regimes.

To model each morphological cell type, each neuron in the model was given a schematic three-dimensional axonal and dendritic arbor consisting of line segments. Seven types of pyramidal cells, with differing morphologies, were placed in cortical layers 2-6 (except for the lower half of layer 4). Eleven morphological classes of nonpyramidal cells are recognized in the visual cortex. Based on descriptions and drawings, stick arbors also were devised for each of these cell classes. The construction of each cell type required about 100 lines of Fortran code.

In each of the six model cortical lamina, neurons were distributed in three dimensions according to pseudorandom numbers chosen for the x, y, and z coordinates. The abundance of each morphological cell class in the model was estimated using data from the experimental literature.

Two major classes of input fibers (denoted X and Y) convey signals from the retina through the thalamus to layers 4 and 6 of the visual cortex. These afferents have been included in the model, but they are modeled differently than the intrinsic neurons. To achieve a broad dispersal of axon terminals throughout such an arbor, we chose to abandon the convention of modeling the axonal arbors as sticks. Instead, each afferent arbor is modeled as a series of cylindrical volumes extending through each lamina or sublamina in which terminal arborization has been reported.

To model the interactions between neurons in the model cortical network, the synaptic connectivity of the network must be specified. The location of each connection with respect to presynaptic and postsynaptic neurons must be specified so that propagation times and signal strengths can be assigned. In the model, synaptic connections are assigned based on proximity of the axonal arbor of the presynaptic neuron and the dendrite or soma of the postsynaptic cell. To make these assignments, the model cortical slab is divided into 3000 cubical boxes, each labeled by a unique integer index. For each model neuron, the indices of the boxes through which the axonal arbor line segments pass are stored in an array AXON. Similarly, the indices of the boxes through which the dendritic arbor line segments pass are stored in an array DEND. The box index of each cell body is stored in an array SOMA. A particular neuron might have ramifications passing through hundreds of boxes. A connection is assigned whenever one cell's AXON array contains an index which is also a member of another cell's DEND or SOMA array. A connection may be thought of as representing one or several synapses.

Numerical values of the number of connections among the cell classes show, first, that all possible connections between cell classes do not occur in the model: about 41 percent of all possible pairs of cell classes exhibit no interconnections. The pattern of connections generated by the model possesses specificity as well as structure. Second, when connections do occur between cell classes, the number of connections for any one cell class pair tends to be small: 771 pairs of connected cell classes each exhibit less than 0.25 percent of all connections. The only connections that are more abundant than this are those between classes of pyramidal cells and those between pyramidal cells and input fibers from the thalamus.

Modeling cortical cell physiology

A great deal of experimentally derived physiological data were used to simulate the properties of cortical neurons in this model. All membrane voltages were scaled to the range 0.0 to 1.0. These scaled voltages are referred to as *activities*. The lower bound, $V = 0.0$, corresponds to the lowest membrane voltage reached by a cell following development of the action potential. The upper bound, $V = 1.0$, corresponds to the average value at the peak of an action potential spike. In the dynamic equations, all activities were initialized at rest, which corresponds to an activity of $V_0 = 0.1$. The firing threshold corresponds to the activity range $V = 0.29 \pm 0.07$. The threshold activity for each neuron was chosen at random from within this interval.

Deviations from the rest activity arise once signals start to arrive externally from the thalamus and intracortically from other spiking neurons. Excitation or inhibition, however, does not reach a postsynaptic neuron instantaneously. The model incorporates a delay in signal arrival comprising the time taken to travel along the presynaptic axon and a synaptic delay.

The equations governing the time dependence of the activity for each neuron depend on whether the activity is below or above the firing threshold of the neuron. These two dynamic regimes are functionally significant in the network. In the suprathreshold regime, the neuron broadcasts signals over its axon to other cells, but does not receive signals.

In the subthreshold regime, however, the neuron receives signals from other neurons, but cannot broadcast signals. Within a short time, signals arrive from a relatively small number of neurons.

In the subthreshold regime, all of the incoming signals to the neuron are summed in the cell body. The change in activity is computed by numerically integrating a system of coupled first-order nonlinear differential-delay equations. The dynamic equations, together with the initial conditions and specification of the thalamic input signals, define an initial value problem. In this study, the thalamic input fibers are pulsed between 0.3 msec and 0.4 msec after the start of the simulation.

In the suprathreshold regime, when the activity for the neuron exceeds the threshold value, the differential equation for this neuron is not integrated during the firing sequence. Instead, the activity is computed from an analytic spike function. The spike function values depend on the type of neuron and the firing state that it is in. Three types of action potential patterns have been described.[1-3]

After firing, each cell was given a refractory (or resting) time of 1 msec during which it could not fire. The cell remained hyperpolarized with its activity below the rest activity during this period. After the end of the absolute refractory period, the cell's activity was again governed by the subthreshold regime equations and, in the absence of other inputs, recovered from hyperpolarization in about 9 msec.

Dynamic simulation

Our results are the product of a series of simulation runs in which we sought to characterize the response of the cortical slab to a brief stimulus to the thalamic input fibers. The stimulus was a 0.1 msec pulse delivered simultaneously to each thalamic afferent at the base of the cortical slab. This stimulus paradigm was intended to model the delivery of an electrical stimulus to the fibers leading from the thalamus to the cortex, as occurs in experiments to determine neuronal response latencies. The model slab contained 3 percent of the true number of neurons in such a slab of cat visual cortex, or 584 cells.

The program DYNAMIC computes the activity of each neuron in the model as a function of time, and the activities of selected neurons can be displayed.[1] The behavior of the model cortical slab in response to thalamic afferent stimulation also can be shown in terms of average activity levels rather than numbers of firing cells. However, the best way we have found to visualize the dynamic data is by making use of the radiative tetrahedral lattices (RTL) volume visualization technique.[4]

Visualization via radiative tetrahedral lattices

The RTL method evolved from our experience with other visualization techniques. The first visualization of the cortical dynamics was done on a Silicon Graphics workstation, using GL graphics calls. The cortical slab was displayed as a rectangular parallelepiped outlined with a wire frame boundary. Horizontal lines divided the volume into its distinct lamina. Neuron cell bodies were represented as small disks located within this volume. Neuron activities were displayed by turning the disks from blue (rest activity) to red

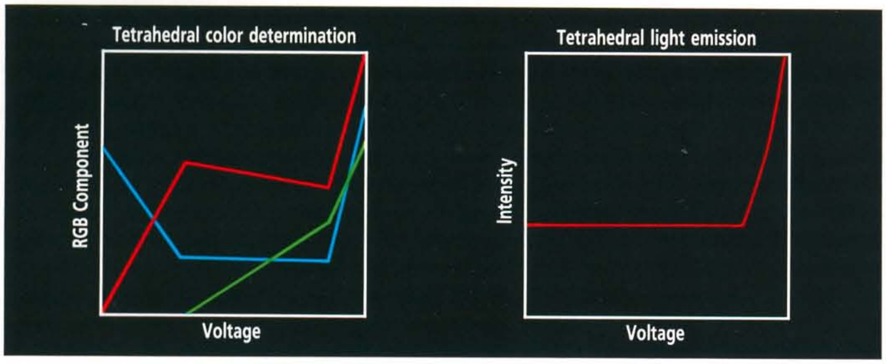

Figure 1. Red, green, and blue components of neuron color as a function of the neuron voltage (left) and emissivity (brightness or intensity) of tetrahedra as a function of neuron voltage (right).

(maximum firing activity). This method was quite fast, allowing a high degree of user interaction.

A second method along these lines used ray tracing techniques. The volume representing the cortex was enclosed by semi-transparent walls, while neuron cell bodies were represented as opaque spheres. The coloring scheme for the spheres was blue to red as in the first visualization. Ray tracing made the spheres appear more realistic, which enhanced the viewer's ability to interpret the scene.[5] However, at this point it appeared that the limit had been reached in terms of the ability to visualize the data using the three cues: position, shape, and color. It was unclear whether allowing each sphere to emit light would improve the visualization. Also, using spheres to represent neurons presented a problem, because the radiosity implementation we planned on using required a sphere to be tesselated into constituent polygons before the energy balance was computed. The tetrahedron was chosen instead because of its geometric simplicity. (The fact that some neurons have roughly pyramidal cell bodies is coincidental.)

Using RTL, each neuron is modeled as a tetrahedron whose locus corresponds to the centroid of the neuron cell body. As a neuron's activity increases, the neuron changes color from dark blue to red and then to white. In addition to its change in color, each tetrahedron emits a radiative burst as it reaches peak activity. In energy-balance rendering methods, light sources can be rendered within the scene. Hence, we model the neuron cell body as a light source, located within a volume, changing its color and "glow level" as a function of neuron activity. It is the ability of the tetrahedra to emit light that differentiates this method from those previously used.

Radiosity methods based on radiative heat transfer theory calculate the energy flow between small areas of a scene called *patches*. Given the emissivity, reflectivity, and relative geometry of a pair of patches, these methods determine the amount of energy exchanged. Energy flow through a scene can be computed and the results used to calculate the amount of energy that reaches the viewer's eye from each patch.

The most popular radiosity algorithms suffer from quadratic complexity in the number of patches.[6] Computation times are rather significant as a result. We have chosen I-nets, a new algorithm introduced at SIGGRAPH '89, whose complexity is only linear in the number of patches.[7] The use of this method significantly reduces computation time.

The scene geometry does not change, because the tetrahedra do not move from one frame to the next. Only the color and emissivities of the tetrahedra change with time. This characteristic is exploited by the I-nets algorithm to reduce further the computation time required for each frame of an animation. In the preprocessing phase of I-nets, lattices of links are built to connect patches between which energy can be transferred. Each polygon in the scene is divided into patches in the I-nets module, and the energy transfer between the patches is computed. For each time increment, the color and emissivity associated with each tetrahedron in the volume are computed as a function of the tetrahedron's activity. The color and emissivity are sampled in red, green, and blue light.

Figure 1 shows the functions we have found that model the emissivity most effectively. For example, the emissivity of the red band is computed from the composition of the function R(a) (red component computed as a function of neuron activity) with E(a) (emissivity as a function of activity). Thus the tetrahedron emits light that is the color of its surface. To make the firing neurons more noticeable, emissivity is held constant until the cell activity reaches 0.8, after which the emissivity increases exponentially with voltage until maximum emissivity has been achieved.

Once energy has been emitted from each neuron, inter-reflections must be computed among all the neurons and the walls of the volume. This is done until all the light has been absorbed or reflected out of the scene. At this point the amount of light energy that reaches the viewer is computed and stored in the vertices of the patches. The vertex and connectivity information is then rendered, and a new frame is stored on an optical disk.

The 584-neuron simulation computes the I-nets phase at four frames per second using a Stardent P3/G3 workstation with a 96-Mbyte memory. Frames are computed and stored on the optical memory disk recorder, which plays them back to produce an animation. Work is underway to port the I-nets module to the CRAY Y-MP system (which may make the network between the supercomputer and the workstation a bottleneck for the algorithm). Running the I-nets module on the CRAY Y-MP system will allow for interactive control of the cortical simulation by sending the output of the I-nets module from the supercomputer to the workstation. The colors computed for the vertices of the tetrahedra and the walls of the volume will be compressed and sent down to the workstation for each step in the interactive visualization.

Dynamics of the cortical slab

Using the RTL visualization technique, we have produced a videotape showing the vertical spread of activity through the cortical slab. Four frames from this video are shown in Figures 2 to 5. The individual neurons appear to glow, a result of using the RTL visualization technique. Sooner than 0.5 msec after the start of the simulation (not shown), the whole slab appears blue, because most neurons have activities near the rest value. When afferent fibers from the thalamus bring input to the middle and bottom layers (4 and 6), the activity of many of the neurons increases and their colors change to red, as shown in Figure 2. The yellow activity plot to the immediate left of the slab shows the average activity increasing in the layers that receive

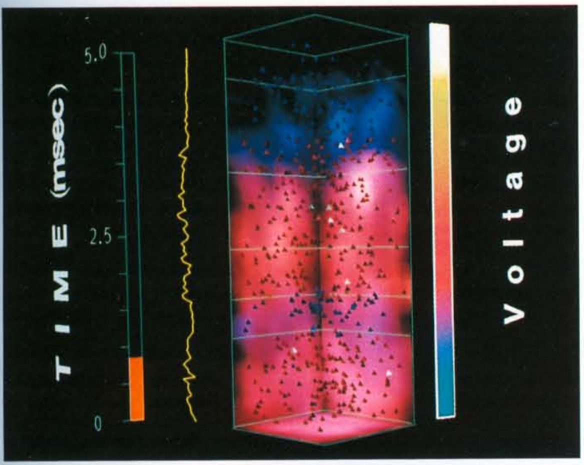

Figure 2. Neuron activity in the cortical slab 0.8 msec after stimulation. The neuron cell bodies are shown as tetrahedra which change color from blue to white as the cell voltage increases from the rest value to the peak firing value. Input from the thalamus enters the middle and bottom layers (4 and 6) of the cortical slab. The vertical color bar on the right shows how the cell voltage is related to the color of the neuron: blue corresponds to a low (or resting) voltage; yellow to white corresponds to high voltage. The time in msec is shown on the vertical time bar to the far left of the cortical slab. The average neuron activity in each of 30 horizontal slices through a slab is shown to the immediate left of the slab. At this relatively early time, the activity is primarily in the middle and bottom layers.

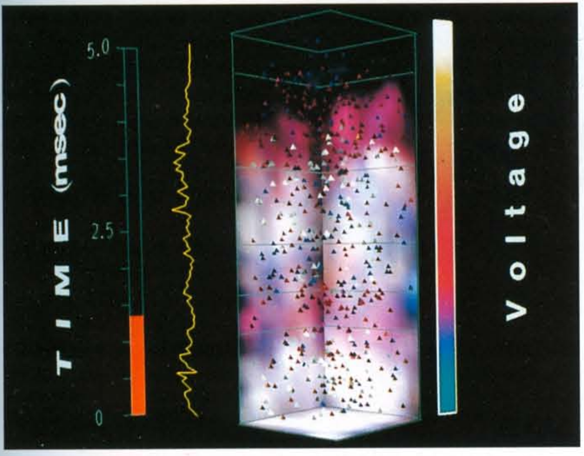

Figure 3. Neuron activity in the cortical slab 1.4 msec after stimulation. Many neurons in the middle (4) and bottom layers (6) are firing.

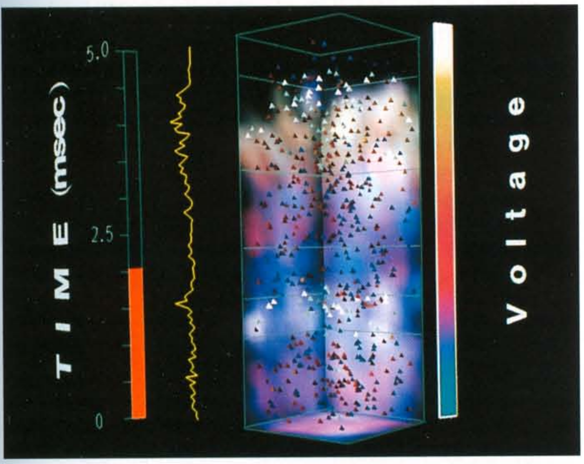

Figure 4. Neuron activity in the cortical slab 2.1 msec after stimulation. The activity spreads up toward the top layers (2 and 3) and into layer 5, which does not receive direct input from the thalamus.

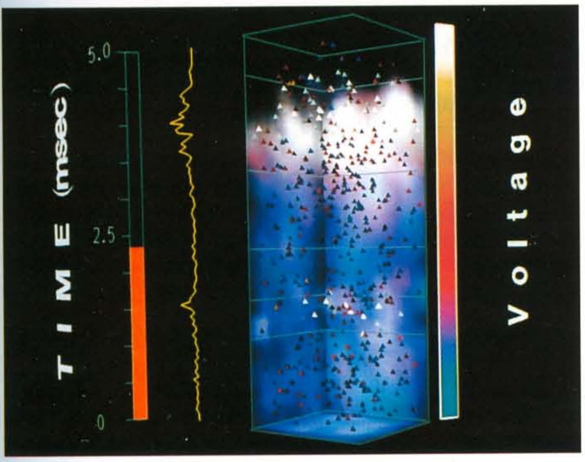

Figure 5. Neuron activity in the cortical slab 2.4 msec after stimulation. The activity decreases in the middle (4) and top layers (2 and 3).

thalamic input. At later times (Figures 3 to 5) the activity gradually builds up in the superficial layers (1 to 3) and in layer 5, but the activity simultaneously decreases in layers 4 and 6. At later times (not shown) the activity decreases throughout the slab. The series of pictures conveys the flow of activity through this slab of cortex. A sound track is being created for a new version of the videotape in which neuron type, location, and activity govern auditory variables, such as volume, pitch, and voice. The sound track is being created with the BioSymphonics sonification algorithm.[8] The combination of the RTL visualization technique and sonification will provide a much better understanding of the complex dynamic phenomena that occur in the brain.

Acknowledgments

The authors thank Paul Patton and Elizabeth Thomas for many discussions during the course of this work and Matthew Witten for encouragement and stimulating conversations. For access to the CRAY Y-MP8/864 computer system, the authors thank the University of Texas System Center for High Performance Computing and Cray Research, Inc., for two years of grant support.

About the authors

Robert E. Wyatt received his Ph.D. degree in theoretical chemistry from The Johns Hopkins University. Following postdoctoral studies in England and at Harvard University, he joined the faculty of the University of Texas at Austin. He is currently the W. T. Doherty Professor of Chemistry and Director of the Institute for Theoretical Chemistry.

Jesse W. Driver obtained a bachelor's degree in computer sciences from the University of Texas at Austin. He has worked for Lockheed Austin and Merit Technologies of Dallas on computer displays for flight simulators and the display of cartographic data. He is currently group leader for the Scientific Visualization Group at the Center for High Performance Computing. He is nearing completion of a master's degree in computer sciences and is president of the Doré Users Group.

References

1. Patton, P., E. Thomas, and R. Wyatt, "Computational Dynamics of Signal Propagation in the Visual Cortex," in *Proceedings of the Twenty-Sixth Semi-annual Cray User Group Meeting,* K. Winget (Editor), Cray User Group, Inc., Bethpage, New York, pp. 22-38, 1990.
2. Patton, P., E. Thomas, and R. Wyatt, "A Computational Model of the Vertical Anatomical Organization of Primary Visual Cortex," *Biological Cybernetics,* (to be published).
3. Patton, P., E. Thomas, and R. Wyatt, "Computational Model for Space-time Signal Propagation in the Visual Cortex," *Biological Cybernetics,* (to be published).
4. Driver, J. and C. Buckalew, "Radiative Tetrahedral Lattices," *Proceedings SPIE* (Symposium On Extracting Meaning From Complex Data: Processing, Display, Interaction), Vol. 1459, February 1991.
5. Cohen, M. and D. Greenberg, "A Radiosity Solution for Complex Environments," *Computer Graphics* (SIGGRAPH '85 Proceedings), Vol. 19, No. 3, pp. 31-40, 1985.
6. Buckalew, C. and D. Fussell, "Illumination Networks: Fast Realistic Rendering with General Reflectance Functions," *Computer Graphics* (SIGGRAPH '89 Proceedings), Vol. 23, No. 3, pp. 89-98, 1989.
7. Glassner, A., "Introduction to Ray Tracing," *ACM SIGGRAPH Course Notes 7,* 1988.
8. Witten, M. and R. Wyatt, "Neural Orchestration: From Cortical Simulation to Cortical Symphony," USENIX Summer Technical Conference Proceedings, Nashville, June 1991.

Announcing new releases of the Cray Research compiling systems

Scientists and engineers long have recognized Cray Research compiling systems as the most powerful programming environment in the industry. Building on its reputation as the world leader in this arena, Cray Research offers the CF77 5.0 compiling system, Cray Standard C Compiler 3.0, and Cray Ada 2.0, the highest performing, most dependable, and easiest-to-use compiling systems available.

CF77 5.0 compiling system

The CF77 compiling system outperforms all other commercially available vectorizing compilers. It is the only Fortran compiler in the industry with reliable default optimization. On the Perfect Benchmarks, a suite of application programs, CF77 realized 13 times overall speedup with default optimization, achieving a factor of 40 speedup on one of these programs. With extra compiler optimizations, a speedup of 130 times was obtained.

Release 5.0 of CF77 complies with ANSI 3.9-1978 and the MIL.SPEC 1753 extension as well as extensions that support implicit data conversion between Cray Research systems and other vendors' systems. Using a method transparent to the user, CF77 automatically translates data to or from other popular data formats, such as those in DEC, IBM, and CDC environments, as well as IEEE floating-point data. Cray Research offers the only compilers with this capability.

CF77 was designed to be easy to use. With one simple command, all or parts of the compiling system can be invoked. Scalar, vector, and multitasked optimization is automatic, and the programming environment includes a wide selection of scientific library routines that are highly optimized. Further performance gains easily are accessible with Cray Research performance analysis tools. New visualization tools based on the X Window System, such as Atexpert and Perfview, provide simple, easily understood displays of performance data in a point-and-click environment. These tools provide general program performance data as well as loop-level analysis and offer recommendations for further optimization.

Release 5.0 improves on the world's most powerful compiling system with the following enhancements:

- ☐ The UNICOS message system increases users' productivity by delivering a single error message format for the entire system. Simple, concise messages give users a clear understanding of errors. In addition, on-line help is available via the *explain* command.
- ☐ Autotasking debugging support with CDBX 6.1 allows users to debug tasked code by setting breakpoints. This enables optimized debugging from the original source, rather than the intermediate source created by the Autotasking Fortran Translator, commonly known as FMP.
- ☐ Compiler information files provide a mechanism for users to create additional tools (such as a listing generator) that are tailored to their specific needs based on compiler output. These files provide more flexibility and are generated on request. Full support for this feature requires version 7.0 of the UNICOS operating system.

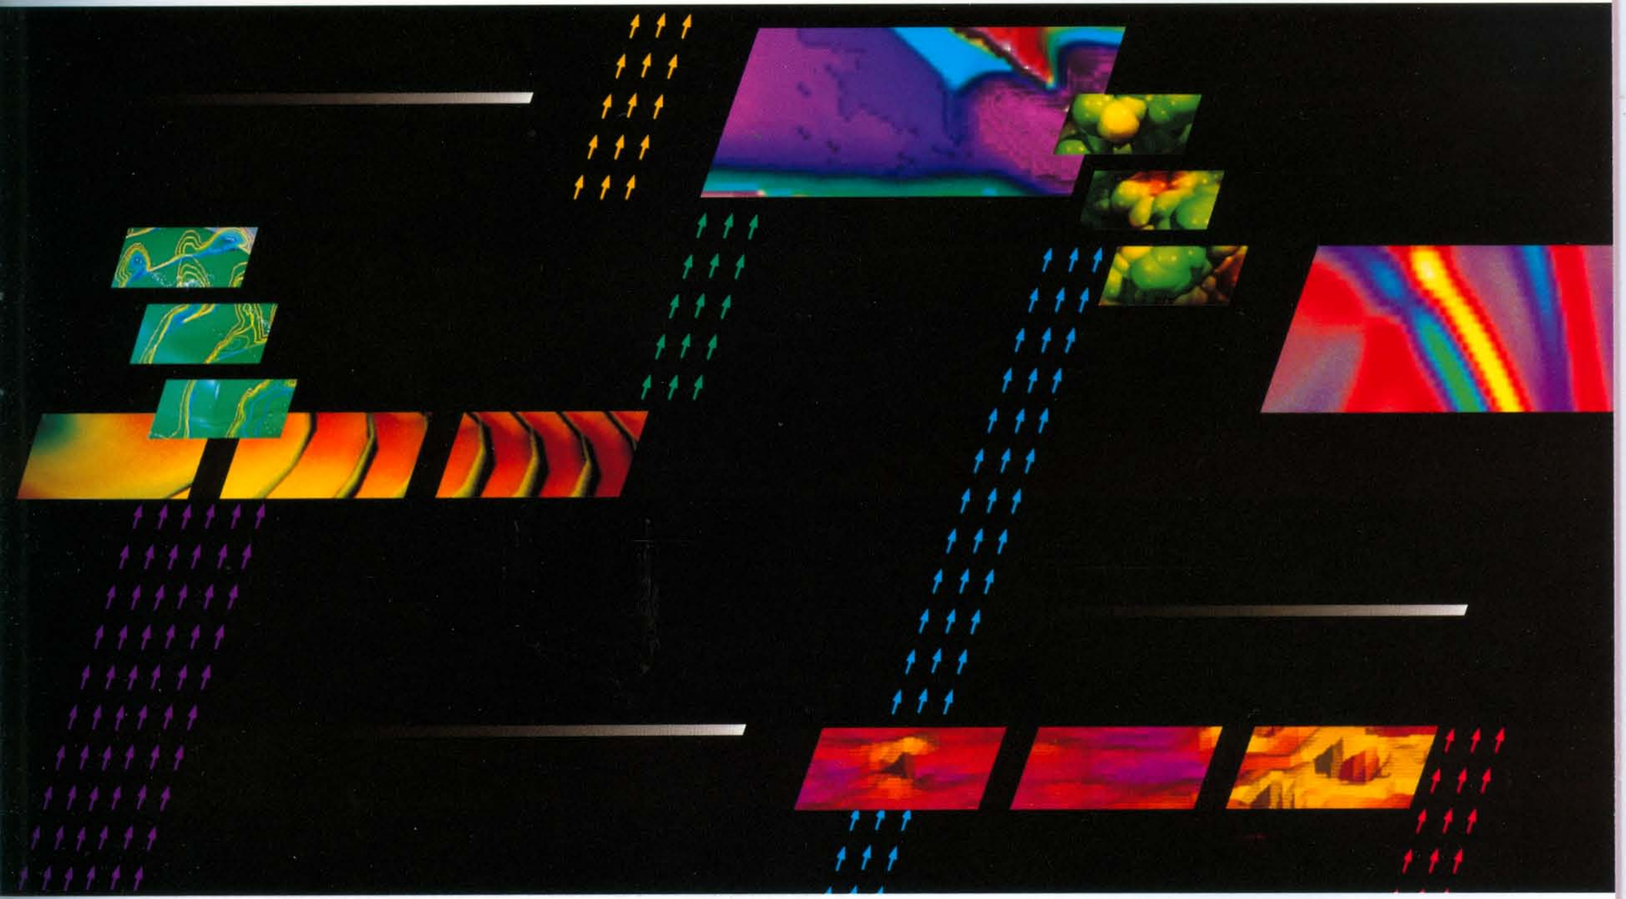

☐ Argument checking provides users with runtime argument checking through the use of a compiler option. The compiler aids debugging by performing a runtime comparison between the number and types of arguments passed to a procedure against the expected number and types of arguments.
☐ Vector reduction loop code sequence modification improves performance for code with vector lengths less than 64 elements.
☐ Nested loop unwinding allows small, deeply nested loops to vectorize.
☐ Conditional vectorization has been significantly improved.
☐ Vectorization of loops containing indirect references to arrays (patent pending) is now allowed.
☐ Extended scope of parallel regions improves performance. The Fortran Autotasking Dependence Analyzer (FPP) now can examine the code between loops, and between a loop and its outer loop, to determine if it is beneficial to execute this code within a parallel region. Thus, larger parallel regions are created when possible, reducing overhead and yielding greater efficiency.
☐ More kernels are automatically recognized and replaced with improved calls to highly optimized Cray Research scientific library routines.
☐ Enhanced Autotasking scheduling improves performance in a multiprogramming environment.

To take advantage of the power of CF77 5.0, version 6.0 or later of the UNICOS operating system must be running.

Cray Standard C Compiler 3.0

The Cray Standard C Compiler exploits the inherent parallelism of Cray Research system architecture by employing the most advanced parallel detection techniques in the industry. Through scalar optimization, vectorization, Autotasking, and automatic inlining, users can realize the potential of their high-performance hardware automatically without code modification. Cray Standard C Compiler 3.0 is the only C compiler that delivers reliable optimization by default, resulting in significant speedups without user intervention.

Release 3.0 complies with ANSI X.3 150-1989, ensuring that programs do not have to be rewritten to take advantage of new language features. Cray Standard C supports the standard C library and all other UNICOS libraries. Through the libraries, it is also able to perform reads and writes on IEEE data formats.

For computationally intensive applications, Cray Research is setting the industry standard with numerical extensions to C, such as complex data types, variable-length arrays, and restricted pointers. These extensions allow programs produced with Cray Standard C to take full advantage of the performance features of Cray Research supercomputers.

As with CF77 5.0, further performance gains easily are accessible with Cray Research performance analysis tools if optimization beyond the level provided automatically is desired. New visualization tools based on the X Window System, such as Atexpert and Perfview, provide simple, easily understood displays of performance data in a point-and-click environment.

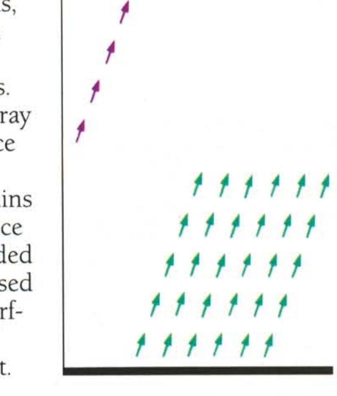

These tools provide general performance data as well as loop-level analysis, and identify areas that would benefit most from further optimization. Debugging also is easier with CDBX in an X Window System interface.

Release 3.0 improves on the world's most powerful C compiler with the following enhancements:

☐ A new block scheduler improves performance.
☐ Autotasking automatically recognizes parallelism and directs processing across all available CPUs to yield performance gains without user intervention on CRAY Y-MP systems.
☐ Character vectorization delivers performance gains with vectorizing character arrays.
☐ New bottom loading speeds up execution of scalar loops.
☐ Inline tasking code has replaced internal tasking function calls, eliminating the function call overhead.
☐ Conditional vectorization allows small, deeply nested loops to vectorize.
☐ More powerful techniques for code restructuring enhance scalar and vector optimization on a CRAY Y-MP system. For example, with one simple command-line option, one loop may be split into two loops to allow one of the resulting loops to be vectorized.
☐ Restricted pointers are a Cray Research extension that send information to the compiler to allow further optimization.

☐ Enhanced debugging tools include support for symbolic debugging of microtasked codes using CDBX.
☐ The UNICOS message system increases users' productivity by delivering a single error message format for the entire system. Simple, concise messages give users a clear understanding of errors. In addition, on-line help is available via the *explain* command.
☐ Function inlining has been enhanced to give users more control over which routines will be expanded inline.
☐ New pragma syntax has been added for use in macro definitions. This allows user directives whenever the macro is used.
☐ A new command option, *-h port,* is helpful in porting from earlier versions of Cray Standard C Compiler as well as other hardware platforms.
☐ New C tasking directives exploit parallelism and improve program performance.

To take advantage of the power of Cray Standard C 3.0, version 6.0 or later of the UNICOS operating system must be running.

Cray Ada 2.0

The Cray Ada 2.0 compiler provides unsurpassed performance when combined with the fast processing speeds and large memories of Cray Research systems. This performance is achieved through vectorization, scalar optimization, and instruction scheduling, while retaining a compatible user interface to earlier releases of Cray Ada.

The Cray Ada 2.0 compiler offers significant performance improvements over previous releases. Thanks to automatic vectorization capabilities, Release 2.0 is 50 times faster than Release 1.0 on Livermore Loop benchmarks. Cray Ada vectorization performance now approaches that of powerful Fortran compilers such as Cray Research's CF77 compiling system.

The United States Department of Defense requires all new and upgraded software systems to use validated Ada compilers for their work. The Ada Validation Facility publishes the Ada Compiler Validation Capability (ACVC) test suite to verify the functionality and completeness of Ada compilers as specified in the Ada Language Reference Manual (LRM). Cray Ada 2.0 has been validated using the ACVC 1.11 test suite and passes all validation tests with optimizations enabled.

Cray Ada 2.0 includes the following new features:

☐ Automatic vectorization capabilities can provide vector performance 50 or more times the performance of scalar code.
☐ Loop block, cross block, and intra-block instruction scheduling capabilities can increase performance by as much as 15 percent.
☐ Scalar performance has been improved by as much as 100 percent with such optimizations as register tracking, full data flow analysis, enhanced inlining, loop invariant code motion, arithmetic strength reduction, strength reduction on loop induction variables, and common subexpression recognition and elimination.
☐ Pragma export capabilities allow routines written in Fortran, C, Pascal and the assembler to call Ada subprograms if an Ada routine is the "main" subprogram.
☐ Cray Ada language features allow users to access Fortran common blocks and C global variables directly for Ada.
☐ A library interface package allows users to reference highly optimized Cray library routines.
☐ Calls to multitasked Fortran and C language routines are supported.
☐ Exception handling provides support for all interrupts issued by UNICOS and has the ability to trap and toggle UNICOS signals within Ada programs.

Cray Research's Ada environment incorporates a set of tools and libraries that provide for flexible, project oriented software development. Release 2.0 conforms to the Ada language as defined in the American National Standard Reference Manual for the Ada Programming Language(ANSI/MIL-STD 1815 A).

CORPORATE REGISTER

Four first-time customers order entry level systems

Four new customers in the United Kingdom, the United States, Spain, and Japan have ordered Cray Research's entry level computer systems.

The **Defence Research Agency, Maritime Division, Admiralty Research Establishment (ARE),** Dunfermline, a United Kingdom Ministry of Defence site located in Scotland, has ordered an entry level system to be used by scientists in the organization's vulnerability and structures group.

Altair Engineering, Inc., an automotive engineering consulting company, installed an entry level system at its facility in Troy, Michigan. The new system replaces a Convex computer system. Approximately 50 Altair engineers are using the Cray Research system for finite element analysis associated with automotive and automotive parts design.

Empresa Nacional del Uranio, S.A. (ENUSA), a national company for nuclear engineering located in Madrid, has ordered a Cray Research entry level system. ENUSA conducts mining, manufacturing, and management of nuclear fuel and will use the new system to run applications related to nuclear engineering and nuclear fuel management.

Nagaoka University of Technology in Japan has ordered a Cray Research entry level system. The university was founded to help develop a new approach to science and technology in Japan and to study issues surrounding Japanese industry and the need to protect the environment. Nagaoka's system is being used for research in energy and environment, materials science, and information science.

The **National Cancer Institute (NCI)** has installed a CRAY Y-MP8 computer system at the NCI Frederick Cancer Research and Development Center in Frederick, Maryland. The new system will be used for computational chemistry to help design new drugs and to search for cures for cancer, acquired immune deficiency syndrome (AIDS), and other diseases. In 1986, NCI became Cray Research's first chemical/pharmaceutical/biomedical customer.

Tokyo-based **Asahi Chemical Industry Co. Ltd.,** Japan's largest chemical company, has installed a CRAY Y-MP2E system at its research laboratories in Shizuoka-ken. Asahi is the first Japanese customer also to license UniChem, Cray Research's easy-to-use, interactive chemistry software environment. Asahi's Cray Research system is connected to scientific workstations throughout Japan. The CRAY Y-MP2E system is the nucleus of Asahi's computing network. "UniChem is a first in the high-performance computing industry," said John Rollwagen, chairman and chief executive officer of Cray Research. "We look forward to working closely with Asahi researchers on further enhancing the UniChem product."

The **University of Rochester's Laboratory for Laser Energetics (LLE),** Rochester, New York, has ordered its first Cray Research supercomputer, a CRAY Y-MP2E system. LLE researchers will use the system to investigate nuclear-fusion energy sources using in-house computational fluid dynamics software codes. The LLE is one of only four U.S. laser facilities actively engaged in inertial fusion studies; it is the primary focus in the United States for direct illumination of fusion targets for energy research. The university is a unique training ground for students and researchers in areas ranging from x-ray microscopy and plasma physics to x-ray laser technology. The fusion research program is supported jointly by the U.S. Government, the state of New York, industry, utilities, and the university.

The **Konrad-Zuse Center for Information Technology Berlin (ZIB)** has ordered a CRAY Y-MP2E system to be networked to the center's CRAY X-MP system purchased in 1987. ZIB provides supercomputer services to German research institutes and universities. The new system will be used for theoretical chemistry and structural analysis applications, as well as other general academic research.

The **Regional Computer Center of the University of Stuttgart (RUS),** Stuttgart, Germany, has ordered a CRAY Y-MP2E system to be used as a high-speed file server. The Cray Research system runs the UNICOS Storage System software, providing a high-performance network storage management system that allows high-speed access to large numbers of files on a variety of storage media. The file server will be networked to RUS's CRAY-2 system and to high-performance graphics workstations used for scientific visualization. The graphical representation of the results produced by the CRAY-2 system will be moved across a local area network to the individual workstations at a speed of up to one gigabit per second. This large bandwidth makes it possible to directly transfer high-resolution, moving images in real time, allowing users to work interactively with their applications.

The **National Center for Atmospheric Research (NCAR),** Boulder, Colorado, has ordered a CRAY Y-MP supercomputer to be one of the few supercomputers in the world dedicated solely to climate research. NCAR will establish a numerical climate research laboratory with funds provided by the Model Evaluation Consortium for Climate Assessment (MECCA), a newly developed international partnership of industry, government, and academic groups from the United States, Japan, Italy, and France. MECCA was formed to evaluate the reliability of current global climate simulation models. NCAR researchers will use their Cray Research system to run climate simulations and improve climate model physics.

Yamaha Motor Co., Ltd. has ordered a CRAY Y-MP2E system to be installed at the company's headquarters in Shizuoka-ken, Japan. Yamaha develops and manufactures a wide range of products including motorcycles, boats, yachts, snowmobiles, and automobile engines. The CRAY Y-MP2E system will be used in engine design, structural analysis, and computational fluid dynamics to develop more fuel-efficient, high-performance products.

Korea's **Agency for Defense Development (ADD)** has ordered a CRAY Y-MP2E system to be installed late this year at ADD's facility in Daejeon City, Repubic of Korea. This will be ADD's first Cray Research system. The system will be used for structural analysis, computational fluid dynamics, and high-speed computer visualization. As part of the order, ADD will obtain a license for Cray Research's Multipurpose Graphic System (MPGS) software visualization tool.

Bayerische Motoren Werke AG (BMW) has installed an eight-processor CRAY Y-MP system, the largest supercomputer offered by Cray Research and the largest of its systems currently installed at a European automotive company. The CRAY Y-MP system was installed at BMW's Research and Development Center in Munich, Germany. The system replaced a CRAY X-MP system and will be applied to structural analysis and computational fluid dynamics problems, including engine design.

ACS and Cray Research announce chemistry fellowships

The Division of Physical Chemistry of the American Chemical Society (ACS) and Cray Research have announced a national competition to award three graduate fellowships in computational chemistry for the 1991-1992 school year and again in the two years following. Cray Research will provide funding for three fellowships each year, each of which will include a cash stipend of $5000. The Subdivision of Theoretical Chemistry of the ACS will manage the selection and awarding process.

The goal of the fellowship program is to encourage doctoral degree studies in computational chemistry by focusing on innovative developments and applications of methods in theoretical chemistry or related disciplines, with emphasis on large-scale numerical calculations or simulations. The Cray Research Fellowships will be awarded on a competitive basis. Applicants must have completed Ph.D. course work and exams, and be ready to pursue or already be pursuing dissertation research.

Use of advanced scientific computing equipment and visualization techniques is central to the subject of this fellowship. Supercomputing must be a critical aspect of the proposed research. Within reasonable limits, Cray Research will provide the student awardees with additional resources in the form of computer time and technical assistance. Announcement of each year's fellowship winners will be made in connection with ACS's national awards symposium.

Cray Research announces collaborative project with Silicon Graphics

In a cooperative venture, Silicon Graphics, Inc. (SGI) and Cray Research have developed an application-building environment that SGI is calling IRIS Explorer. The SGI product will be distributed in December of 1991. Cray Research's version of the product, tentatively called Cray Explorer, will be available in the first quarter of 1992 to run distributively between SGI and Cray Research platforms.

IRIS Explorer is a visually based, object-oriented application-building environment that blurs the line between writing and using applications. Computational scientists and engineers with unique visualization requirements can integrate their existing data sets and algorithms for fast custom viewing without writing any code. This is more than just a visualization environment, IRIS Explorer provides a visual programming interface to develop all types of applications including those that do not use computer graphics, making it a very flexible testbed environment.

Cray Research believes IRIS Explorer represents leading-edge technology, and therefore wants to be the first to offer it on a platform that runs distributively on SGI hardware. Like AVS, a product of Stardent Computer Inc., and apE, developed at the Ohio Supercomputer Center, IRIS Explorer lets users construct programs by linking visual icons. However, the technology is more advanced in that it was designed from the outset for distributed processing. This next-generation technology provides a new way to build applications through the modular writing and sharing of code. IRIS Explorer technology also offers the following advantages:

- ☐ Application size and complexity scales up very well, unlike other systems that have inherent data flow bottlenecks.
- ☐ All parameter controls are visible at once and are visually "connected" to the modules they control; other systems limit parameter access to one at a time or produce additional disassociated windows.
- ☐ IRIS Explorer was designed to support distributed processing in heterogeneous environments and to support machines without virtual memory and different machine arithmetic.
- ☐ Visual tools are provided for integrating new modules and foreign data types into the system. With these tools, "wrapper" code for user-supplied modules is automatically generated. Other systems require this code to be written manually for each new module.

Initially, IRIS Explorer will run on Cray Research and SGI hardware platforms only. However, SGI plans to make a porting kit available. This was an important consideration in Cray Research's participation in this cooperative effort.

IRIS Explorer demonstrations will be held at the fall 1991 meeting of the Cray User Group and at the Cray Research booth at IEEE Supercomputing '91 (November 18-22, Albuquerque, New Mexico).

Cray Research selected as finalist for NASA quality award

Cray Research announced in July that the National Aeronautics and Space Administration (NASA) has chosen the company's manufacturing division as one of eight finalists for the 1991 George M. Low trophy, NASA's excellence award.

Last year, the NASA Excellence Award was renamed the George M. Low trophy in honor of the space pioneer. The award, established in 1984, recognizes NASA prime contractors, subcontractors, and suppliers who have demonstrated outstanding achievements in quality and productivity improvements within "total quality management" (TQM) programs. The goals of the award are to improve internal quality and productivity practices as well as TQM processes within NASA and its contractors' organizations and to transfer performance improvement methods of award recipients to other organizations, associated with NASA.

NASA has been a Cray Research customer since 1981, when the NASA Ames Research Center received a CRAY-1 system. Twenty Cray Research supercomputers currently are installed at seven NASA research laboratories. Cray Research is the first commercial supplier to become a finalist for the George M. Low trophy.

NASA will announce the final award recipients November 7, 1991, at the Eighth Annual NASA/Contractor Conference.

APPLICATIONS UPDATE

Cray Research and Prosys Technology develop chemical process simulator for supercomputer use

Cray Research, Inc., has signed an agreement with Prosys Technology Ltd. (ProsysTech), of Cambridge, England, to jointly develop a high-performance version of ProsysTech's SPEEDUP software for use on Cray Research supercomputers. Product availability is scheduled for first quarter 1992.

SPEEDUP is a software system used to calculate and solve problems associated with chemical process engineering. Based on original research conducted at London's Imperial College, SPEEDUP has proved to be a dynamic simulator that can improve quality and yield, reduce operating costs, and enhance the safety and environmental conditions of chemical processing plants. SPEEDUP is used by over 100 major process industry companies around the world.

The availability of the SPEEDUP system on Cray Research supercomputers will benefit numerous industries, including chemicals, petrochemicals, gas processing, oil refining, inorganics, pharmaceuticals, nuclear energy, electric power, minerals, pulp and paper, food processing, and waste and water treatment. With SPEEDUP, a process engineer can perform a variety of simulations associated with chemical processing:

☐ Dynamic simulation, design, and optimization
☐ Steady-state simulation, design, and optimization
☐ Parameter estimation
☐ Data reconciliation

Cray Research and ProsysTech will work jointly to optimize the performance of SPEEDUP on Cray Research systems. The Cray Research version of SPEEDUP will feature substantially enhanced abilities to solve larger, more complex industrial problems for design and chemical process operations. Du Pont and Company is the first Cray Research customer to purchase a license for the Cray Research version of SPEEDUP. Support from Du Pont for the product encouraged Cray Research and ProsysTech to convert the software to run on Cray Research systems. "We see dynamic simulation as increasingly important in process design and plant operation," said David Filkin, a senior consultant of applied mathematics at Du Pont. "The Cray Research version of SPEEDUP will be an integral part of Du Pont's process simulation efforts."

Chemical plants rarely run at steady-state, due to changes in feedstocks, operational objectives and settings, and because many plants are batch or cyclic in operation. SPEEDUP provides a continuous view of a process in action by calculating the transient behavior of a chemical plant over time. Parallel processing on Cray Research supercomputers will offer significant advantages to make this technology a reality for large-scale industrial problems.

As part of the agreement with Prosys Technology, Cray Research is sponsoring a limited number of university licenses for the Cray Research version of SPEEDUP. With this approach, Cray Research and ProsysTech will assist engineering students in gaining experience with equation-based, dynamic simulation methods for chemical processing problems.

For more information on using SPEEDUP with Cray Research systems, contact Peter Winter, Prosys Technology, Sheraton House, Castle Park, Cambridge CB3 OAX, England; telephone: (0223) 312220; or Steve Zitney, Cray Research, Inc., 655-E Lone Oak Drive, Eagan, MN 55121; telephone: (612) 683-3680.

NMR "back-calculation" method available in chemistry packages

Nuclear magnetic resonance (NMR) experiments have become a powerful means for determining the structure of proteins in solution. Computation plays a critical role in the final stages of these experiments, which produce cross-peak intensities that are measured to obtain an estimate of the distance between the two nuclei responsible for a given peak. The distance between the nuclei is estimated using the independent spin approximation, in which the presence of nuclei other than the two involved in the peak is neglected. After the distances are determined, the structure of the protein is optimized computationally so that as many of the distances as possible are satisfied.

Researchers David Case and Ping Yip of the Research Institute of Scripps Clinic have demonstrated a new method that allows researchers to refine NMR data directly against cross peaks without using the indedendent spin approximation. They developed the method by deriving an analytic formula for the derivatives of the intensities with respect to the coordinates of the protein:

$$\nabla_{\mu}[\exp(-\mathbf{R}\tau)_{ij} = \sum_{\mathrm{rstu}} L^{\mathrm{T}}_{ir} \mathrm{L}_{rs} (\nabla_{\mu}\mathbf{R})_{st} L^{\mathrm{T}}_{tu} L_{uj} \left[\frac{\exp(-\lambda_r \tau_{\mathrm{m}}) - \exp(-\lambda_u \tau_{\mathrm{m}})}{(\lambda_r - \lambda_u)}\right].$$

Although the method requires two or three orders of magnitude greater computational power than the independent spin approximation, it provides a much truer model of the experiment because it takes spin diffusion into account. This new method has become popular among NMR groups and is available in several application programs that run on Cray Research computer systems:

☐ UCSF AMBER 4.0 — from David Case at the Research Institute of Scripps Clinic and Peter Kollman of the University of California, San Francisco.
☐ X-PLOR — from Axel Brünger at Yale University
☐ DIANA — from the lab of Kurt Wütrich at the Eidgenössische Technische Hochschule (ETH) in Zurich, Switzerland.

For more information on the NMR back calculation method and the programs in which it is used, contact John Mertz, Cray Research, Inc., 655-E Lone Oak Drive, Eagan, MN 55121; telephone: (612) 683-3674.

Cray Research becomes industrial affiliate of the Engineering Design Research Center (EDRC)

Cray Research has joined the Engineering Design Research Center (EDRC) of Carnegie-Mellon University as an industrial affiliate. The EDRC is a multi-disciplinary National Science Foundation (NSF) engineering research center that develops and integrates computer-based methods to support fast, efficient, and cost-effective design in United States industry. The EDRC spans all engineering disciplines as well as architecture and computer science. The center operates on the single premise that design practice can be improved dramatically at an early stage by applying formal methods that enable designers to:

☐ Anticipate life cycle concerns, such as manufacturing, testing, and reliability
☐ Convey thought processes that motivate design decisions
☐ Automatically generate and select design alternatives
☐ Develop design systems that are easy to learn and use

Cray Research joins the EDRC industrial membership of 39 companies, which represents the aerospace, electronics, computer, chemical, petroleum, automotive, medical equipment, construction, and materials industries. Many of the industrial affiliates are current and prospective Cray Research customers.

The EDRC's research program comprises three laboratories: Design for Manufacturing, Synthesis, and Design Systems. In the Design for Manufacturing Laboratory, faculty members from electrical and computer engineering, mechanical engineeering, and architecture are developing methods that enable designers to anticipate the downstream consequences of decisions they make while formulating a design. The Synthesis Laboratory, which combines expertise from chemical, civil, and mechanical engineering, as well as industrial administration, computer science, and architecture, is developing and integrating qualitative and quantitative methods for design optimization in layout design, separation system synthesis, preliminary design of engineering systems, automated learning, and combinatorial optimization. The Design Systems Laboratory draws upon architecture, chemical, civil and electrical engineering to create generic design tools and domain-independent systems to support the large-scale integration of dissimilar design tools. A major goal of this laboratory is to create schemes for integrating designers, CAD tools, and modeling facilities in organizations that support concurrent design.

Cray Research joins collaboration to improve chemical process engineering

In an effort to create a prototype system integrating the synthesis and simulation of chemical process plants using mixed-integer nonlinear programming (MINLP) techniques, Cray Research has entered into a collaborative project in chemical process engineering with three organizations representing different areas of expertise: Air Products and Chemicals, industrial applications; Aspen Technology, a flowsheet software vendor (ASPEN PLUS); and the Engineering Design Research Center (EDRC) of Carnegie-Mellon University, noted for algorithm research and development.

Process synthesis can be viewed as the thinking process by which a design engineer arrives at a flowsheet configuration to transform chemical reactants into desired products. Once one or several flowsheet configurations are determined, process simulation packages can be used to set design parameters for plant design. Process simulation can be used to help guide the process-synthesis stage. A single, efficient design method enables users to identify the optimal configuration and operating parameters.

Merging traditional process simulation with process synthesis requires the ability to handle both continuous and discrete decisions. Continuous design decisions, such as setting equipment operating conditions, are handled effectively with commercially available process simulation systems and nonlinear programming (NLP) software. Process synthesis invariably requires many discrete decisions such as the type and number of unit operations, feed tray locations, and number of distillation stages. These discrete decisions present an integer component that makes the overall synthesis/simulation optimization problem a MINLP problem. The MINLP approach involves solving an alternating sequence of NLP problems and mixed integer linear programming (MILP) problems.

Researchers at the EDRC have demonstrated the concept of using MINLP techniques to solve problems involving both discrete and continuous design decisions. Design applications include selection of optimal process flowsheet structures, synthesis of heat exchanger networks, and retrofits of steam and power plants, to name a few. Operations examples can apply to plants requiring startup and shutdown of units because of demand or feedstock fluctuations, multiproduct batch plant operation, and operation of steam and power generation plants with chemical process facilities.

Presently, no MINLP capability exists to tackle large-scale industrial problems even though the concepts have been clearly demonstrated by academic researchers. This is a case where technology exceeds practical applications, largely because the MINLP approach requires a giant leap from traditional design practices. In addition, advanced computer architectures employing parallelism will need to be used effectively to make this technology a reality for large-scale industrial process engineering.

The next step toward translating these techniques into industrial practice is to integrate MINLP capabilities in a framework where commercial simulators can be used to solve industrial problems. Such a goal requires computer software and hardware, algorithm development, and industrial applications as embodied in the MINLP project team of Air Products and Chemicals, Aspen Technology, Cray Research, and the EDRC.

For more information regarding EDRC activities, contact Richard LaRoche, Cray Research, Inc., 655-E Lone Oak Drive, Eagan, Minnesota, 55121; telephone: (612) 683-3696.

Cray Research and Dutch research organization to develop petroleum exploration software

Cray Research announced in July that it has signed an agreement in principal with the Dutch research organization TNO (Toegepast Natuurwetenschappelijk Onderzoek), Delft, The Netherlands, to jointly develop a new generation of commercial seismic application software for the petroleum industry. The resulting software will make seismic modeling on high-performance computers more accurate and useful in the discovery of oil and gas.

As part of the agreement, Cray Research will provide a CRAY XMS computer system, to DELSI, as the project is named. The system will be used as a platform for the development of commerical products based on prototype software developed by DELPHI, a seismic data processing research project headed by A. J. Berkhout at the Delft University of Technology. Cray Research has been associated with the DELPHI project for the past five years and is one of the 22 sponsoring corporations.

To help broaden the use of high-performance computing at all levels of the petroleum industry, three portions of prototype software code have been identified for commercialization. They are model-based algorithmic approaches, which focus on seismic energy to discern detailed geological characteristics of the potential hydrocarbon reservoirs. Using the CRAY XMS system, the DELSI team will make these codes easier to use through enhancements, applying IEEE standards, writing software documentation, and executing case studies.

The CRAY XMS system will be installed at TNO and, through a high-speed network, linked to the CRAY Y-MP system located at SARA (Stichting Academisch Rekencentrum Amsterdam), the Netherlands national supercomputer center. The DELSI project will make the commercialized software available on Cray Research's entire product line of high-performance computers.

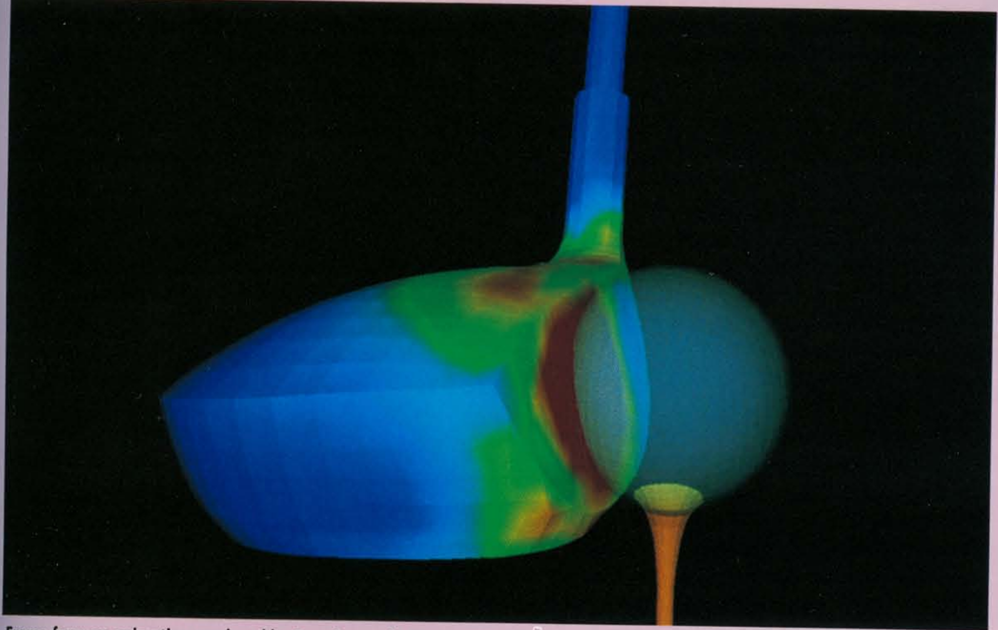

Frame from an animation produced by Cray Research, MacGregor Golf Company, and the MacNeal-Schwendler Corporation. The image shows the stresses in a titanium metal driver at a particular instant as it strikes a golf ball at 100 miles per hour. Simulations such as this can be used to help design stronger, more lightweight clubs.

MacGregor brings computing to the fore

MacGregor Golf Company, MacNeal-Schwendler Corporation, and Cray Research teamed up at the 1991 U.S. Open golf championship to demonstrate how supercomputers can be used to design and deliver better-performing golf clubs. Analysts from MacNeal-Schwendler used a CRAY Y-MP computer system and the MacNeal-Schwendler Corporation's MSC/DYNA structural analysis software package to simulate what happens when a MacGregor T920 titanium metal driver strikes a golf ball at 100 miles per hour — the approximate speed of a normal swing. Using Cray Research's MPGS graphics software, the Cray Research analysts produced video footage of the simulation showing an "ant's-eye-view" of the club head as it strikes the ball. The video animation was shown at the U.S. Open, held in June at Hazeltine National Golf Club in Chaska, Minnesota.

MacGregor commissioned the analysis to investigate the use of titanium in the manufacture of a club with a larger head, which would carry a larger striking surface and, ultimately, help golfers hit straighter shots. Because titanium is lighter than stainless steel, it allows a club to have a larger head without an increase in total weight.

It takes only three weeks to create, test, and analyze a new club design using supercomputer modeling, compared to three months using conventional prototype testing. Golf club designers often place a tiny, quarter-inch-square strain gauge onto the head of a prototype club prior to conducting physical stress tests. The device reads and records the stress only locally at its position on the head. The analysis done on the CRAY Y-MP system, however, revealed the stresses in the club head at 1000 surface positions. The supercomputer simulation produced results that were 25,000 times more accurate than physical testing would have produced.

Additionally, MPGS was used to visualize the output, a capability MacGregor engineers had never before experienced. They were impressed with the animation of their analysis runs and immediately began to rethink some of their design assumptions.

"For the T920 design we were dealing with a material and head size we hadn't worked with before," said Clay Long, vice president of research and engineering for MacGregor. "The CRAY Y-MP simulation not only helped us speed our design and test cycle, but also allowed us to see more closely how this material and larger head would react upon typical impact and show us if such a radical design would be possible." Long added that the project will result in a future product made of titanium.

Molecular dynamics study provides new insight into calcium-binding proteins

One of the more exciting applications of molecular dynamics is in molecular biology. Using high-performance Cray Research computer systems, researchers are able to study complex biological processes that previously could have been investigated only with costly and time-consuming experiments. The work of Harel Weinstein, a professor of physiology and biophysics at the Mount Sinai School of Medicine, is a case in point.

Weinstein used molecular dynamics to study calcium-binding proteins, such as calmodulin, troponin-C, and calbindin. These proteins play an important role in enzyme activation, muscle contraction, stress response, and many other cell functions triggered by calcium ion concentrations. All the members of this class of proteins have the same structural motif: a helix-loop-helix sequence called the "EF hand." By performing molecular dynamics calculations on the CRAY Y-MP computer system at the Pittsburgh Supercomputing Center, Weinstein was able to show that this motif is a direct consequence of intramolecular interactions and solvent forces (the interaction of the surrounding water molecules with the protein). Incorporating solvent forces in a computer simulation is time consuming and can be done only on a supercomputer, such as the CRAY Y-MP system.

Another question Weinstein answered is why calcium, as opposed to other similarly charged ions, such as magnesium, binds to these proteins. The solvent again plays a key role. Weinstein's calculations show that both ions bind well to these proteins but that magnesium binds to water much more strongly than calcium does. Because the structure of these proteins prevents water molecules from approaching the binding site, it is easier for calcium than for magnesium to leave the surrounding water and bind to the protein.

Weinstein's most exciting discovery involved computations on calmodulin. X-ray crystallographic studies, which indicate the structure of proteins in a special, high-density environment, showed that the calmodulin structure forms a "dumbbell" — two binding domains each with Ca^{2+}binding sites connected by an erect helical backbone. Other work indicated that when calcium-calmodulin complex binds to its target proteins, its "dumbbell"

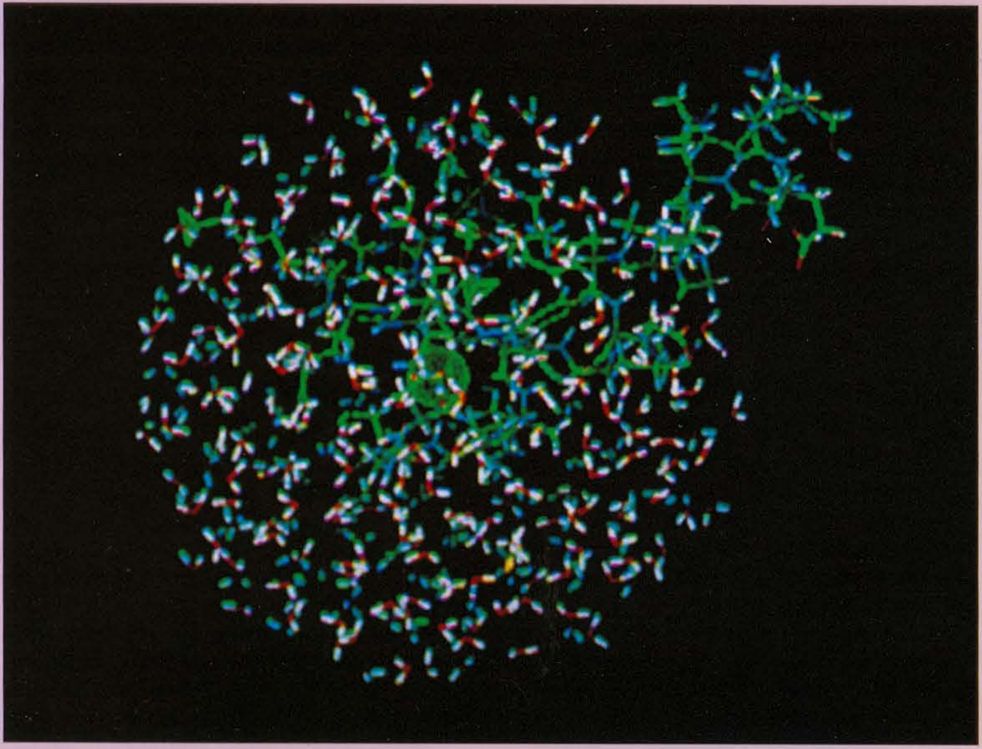

Snapshot from a molecular dynamics study of calcium-binding proteins conducted on a CRAY Y-MP system by Harel Weinstein of the Mount Sinai School of Medicine. The protein (green and blue) forms a helix-loop-helix structure called the EF hand. The solvent (white) is represented by a "drop" of water surrounding the protein.

structure becomes compacted. However, Weinstein's work shows that in water solution this compacted form is not merely a result of protein binding, but a consequence of changes in the helical backbone, helped along by Ca^{2+} binding. Weinstein notes that, "This shows the great value of computer simulations. The crystal structure is what people have been using to explain the activity of this very important molecule, and we now know that this can be misleading." (This article condensed from *Projects in Scientific Computing,* Pittsburgh Supercomputing Center, 1990.)

UNICOS TCP/IP extensions support high transfer rates

Van Jacobson of Lawrence Berkeley Laboratory, working with the Pittsburgh Supercomputing Center (PSC), achieved networking breakthroughs at the recent Educom National Net '91 conference in Washington, D.C. Before the conference, Jacobson replaced his Sun OS networking code with his own experimental TCP/IP code written according to RFC1072/RFC1185 extensions to TCP. The code is designed to expand the bandwidth of TCP/IP. He did not have a chance to test the code on his workstation (a Sun Sparcstation 2) before the conference. Jacobson's experiment involved connecting his Sun workstation and PSC's CRAY Y-MP system via TCP/IP via the NSFNet backbone to see if his code and a similar one developed independently at Cray Research were compatible and if they performed as designed.

Jacobson connected the workstation to a 10-Mbit/sec Ethernet that was connected to the 45-Mbit/sec NSFNet backbone. The PSC CRAY Y-MP system, 1000 miles away, was connected to a 100-Mbit/sec FDDI ring, which also was connected to the NSFNet backbone. When the connection was made, Jacobson's code and the enhanced UNICOS TCP/IP code on the PSC system, which were developed independently, worked together on the first try; the Sparcstation 2 and the CRAY Y-MP system began exchanging data flawlessly using a 512-Kbyte window. The data being transferred from PSC to Washington flooded Jacobson's local Ethernet, achieving a sustained transfer rate of 10 Mbits/sec — the limitation of the Ethernet to which the workstation was connected.

"Two developers who had not seen each other's work wrote code that talked on the first try," said David Borman of Cray Research's networking and communications group, who developed the enhanced TCP/IP that was resident on the PSC Cray Research system. "This speaks well for the specifications in the standard. To be able to flood the Ethernet from 1000 miles away means that the enhanced code allows the support of sustained transfer rates over long delay paths that cannot be achieved with standard TCP/IP. With standard TCP/IP, the sending side could only send 64 Kbytes of data before having to wait for an acknowledgment. For long-delay, high-speed paths, this can be a severe limitation. With enhanced TCP/IP, this limitation is removed; now the limitation is the rate of the wire and the rate of local processing."

The enhanced TCP/IP facility is included in version 6.1 of Cray Research's UNICOS operating system.

At Apple Computer, ANSYS tests product durability

During the design of the Macintosh portable computer, Apple engineers used the ANSYS structural analysis package from Swanson Analysis Systems, Inc., on a Cray Research supercomputer to provide timely results. The portable Macintosh computer comprises three main plastic parts: the bottom case, the top case, and the subframe, which houses the floppy and hard drives, batteries, and logic board. This subframe snaps into the bottom frame at various locations that act as areas of internal support.

Because of proposed logic board changes, engineers had to remove and change the positions of some internal supports and test the product's resulting durability. "We wanted to determine how these changes would affect the subframe during drop tests," said Ken Lown, structural engineering manager at Apple Computer. He explained that the tests determine how portable products respond when dropped with and without protective shipping packages.

Using ANSYS on a workstation and a Cray Research supercomputer, Lown found that the proposed changes would result in a 30 percent decrease in torsional stiffness and a 36 percent decrease in bending stiffness for the two drop tests. Then, based on this information, he determined the response stress levels for the subframe. The modified subframe showed excessively high stress and shock upon impact, so Apple engineers used the model to find better locations to snap the subframe into the bottom case. "We ran four trial cases and came up with two options that would bring stress and shock-response levels back to acceptable levels," said Lown.

"I was able to run this model, look at modifications in two days, and complete the entire project in a week and a half, which was really helpful," he said. "The design teams on all projects are under tremendous time pressure, and they need to make informed decisions on their designs whether they have analysis results or not. I was able to provide this information in a time frame that made it useful to them."

GALLERY

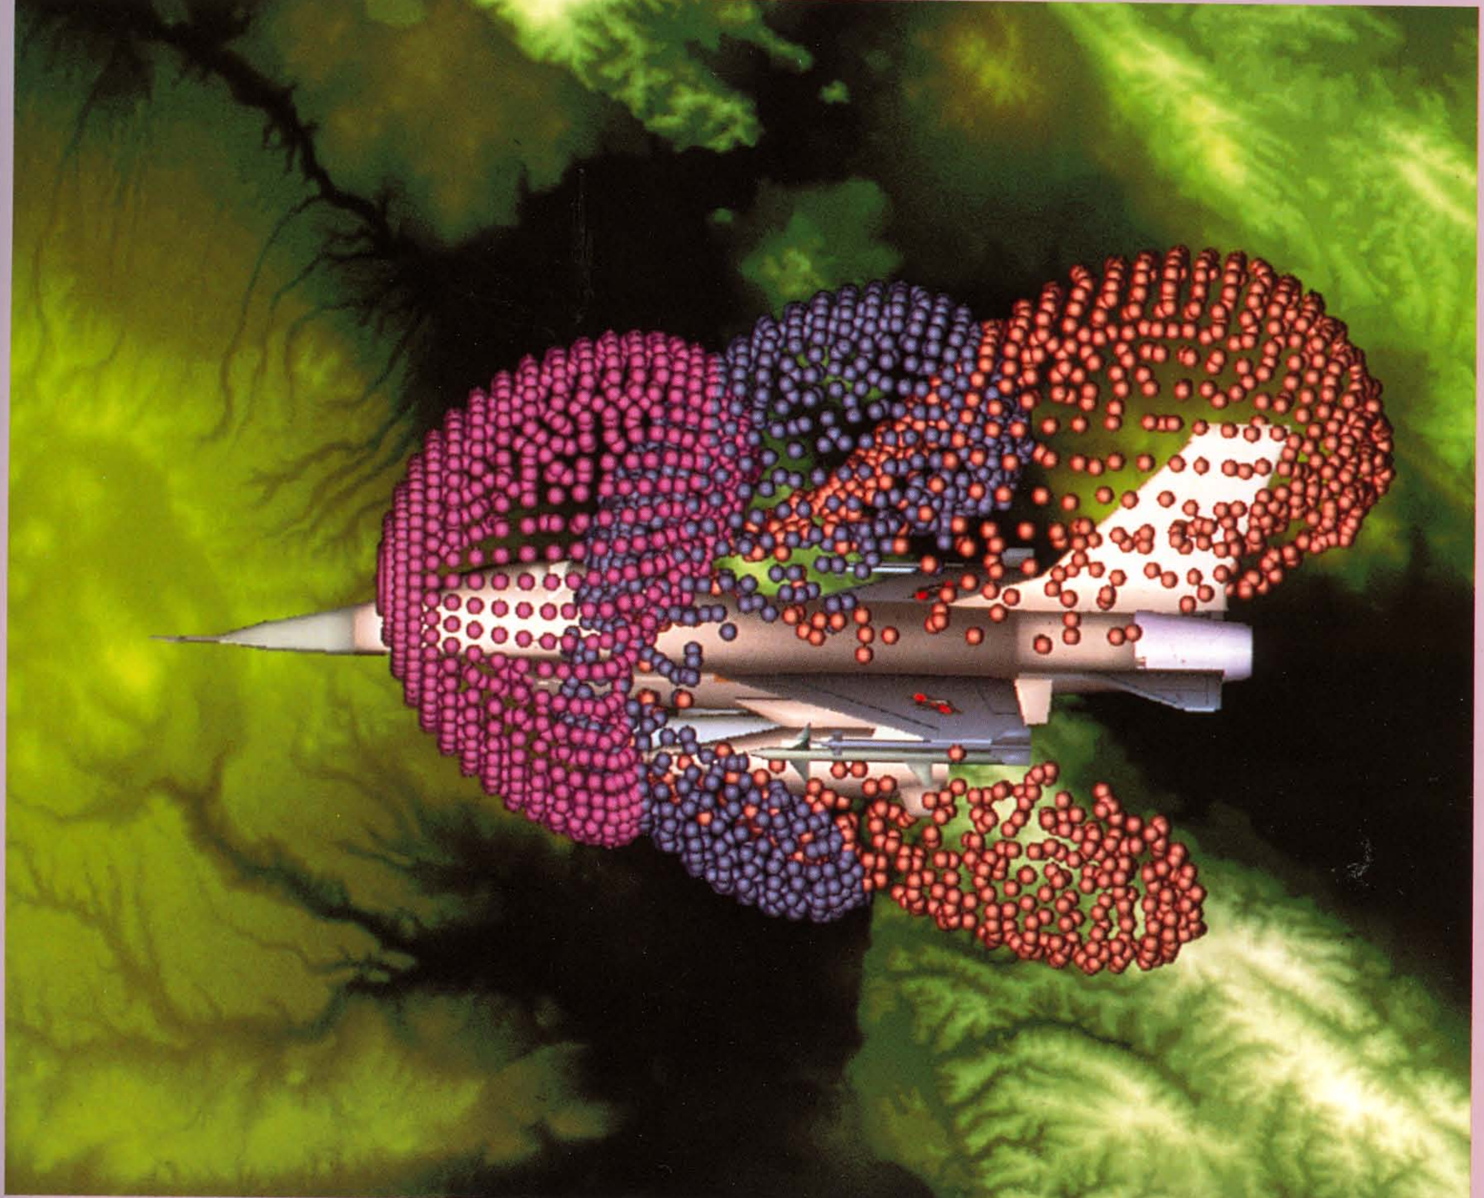

Computed tracer particles show air-flow patterns around an F-16 fighter plane. The terrain was generated from the 60-square-mile Albuquerque-to-Los Alamos section of U.S. Geological Survey data. The particle positions were computed and the image generated using the NEWTUN fluid dynamics program, CRAY Y-MP and CRAY-2 computer systems, a Silicon Graphics workstation, and an Ultra Network Technologies frame buffer. Image courtesy Laurence A. Feldman, Fluid Dynamics International.

CRAY
RESEARCH, INC.
Marketing Communications
1440 Northland Drive
Mendota Heights, MN 55120
612/683-7215

$F(y', y, t$ … for $t_0 \le t \le t_f$

$y(t_0) = y_0$ … nd $F: D \subset \Re^n \times$

with $y \in \Re^n, t \in$