PITTSBURGH
SUPERCOMPUTING
CENTER

Projects in Scientific Computing

1988-89

Cover: Images in the collage, rendered by Steve Palmer, represent the research projects presented in this book.

Projects in Scientific Computing

1988-89

Published June 1989 by the Pittsburgh Supercomputing Center
Mellon Institute Building
4400 Fifth Avenue
Pittsburgh, PA 15213

Foreword

The scientific impact of the Pittsburgh Supercomputing Center continues to grow. Research at the Center has led to over 500 papers in virtually all fields of science and engineering. As you can read in the following pages, this work has advanced fundamental science, spurred the development of new computational technologies and improved the competitiveness of American manufacturing. The studies highlighted here, only a sampling of the rich diversity of projects undertaken, confirm that computational science forms a productive complement to the more traditional experimental and theoretical approaches.

These developments have come about through combining the genius of the American scientific and engineering community with the expertise and dedication of the Pittsburgh Supercomputing Center's staff. They could not have occurred without the unflagging support of Carnegie Mellon University and the University of Pittsburgh, the continued cooperation and experience of the Westinghouse Electric Corporation, the commitment of the Commonwealth of Pennsylvania and the vision and continued funding of the National Science Foundation.

This year marked the first in a series of significant upgrades at the Pittsburgh Supercomputing Center designed to enable even greater scientific breakthroughs. The new CRAY Y-MP is three times more powerful than our workhorse for the past three years, the CRAY X-MP. Not only does the new machine increase the Center's computational capacity; more importantly, it enhances capability. Scientists and engineers can now attack problems that were beyond their previous scope.

The Pittsburgh Supercomputing Center's Y-MP is the first supercomputer of such power available to the entire national research community. Now American researchers can apply the most powerful supercomputing capability in the world to their problems. The Center's upgrade plans assure them that, as they work at the Pittsburgh Supercomputing Center, they can continue to count on the most advanced supercomputing technology in tackling projects of major scientific significance.

Michael J. Levine, Scientific Director

Ralph Roskies

Ralph Z. Roskies, Scientific Director

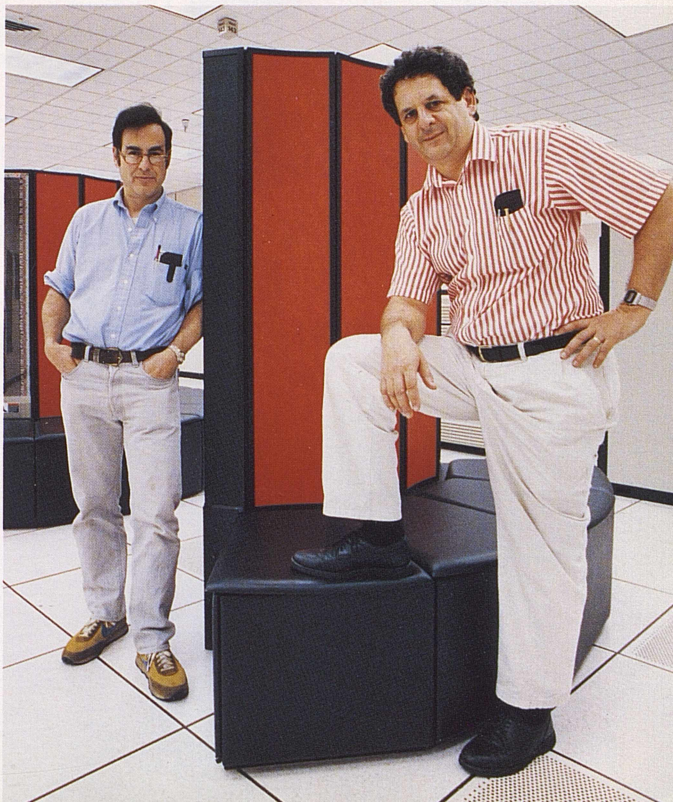

Mike Levine & Ralph Roskies at Westinghouse Energy Center, site of the Pittsburgh Supercomputing Center's computing facilities. By April 1989, the Center's newly installed CRAY Y-MP/832 had completely replaced this X-MP.

Contents

Secretaries Pat Guzzie and Jenifer Barbera often greet visitors to the Pittsburgh Supercomputing Center.

Pittsburgh Supercomputing Center at Three Years

Beverly Clayton has overseen day-to-day activities at the Pittsburgh Supercomputing Center since before it opened in June 1986—three years of building increased support for advanced computational research. "To date we've supported over 1700 scientists in 46 states. As demand for our service has grown, we've added good, experienced people. We now have 61 full-time staff, and with part-timers and students, we've essentially doubled since last year. This year we significantly upgraded our hardware capability to do more and better science. To do this, we relied on the energy, dedication and hard work of over 100 individuals."

KA

Beverly Clayton, Executive Director, with Jack Worlton of Los Alamos National Laboratory and Janet Brown, the Center's Education Coordinator, at the 1988 Summer Institute banquet. Worlton spoke on "Technology Forecasting for Supercomputers" at the Summer Institute Symposium.

KA

Press conference, August 11, 1988, announcing the agreement with Cray Research. Speakers included Pennsylvania Governor Robert Casey, presidents Richard Cyert and Wesley Posvar of Carnegie Mellon and the University of Pittsburgh, Paul Rotar, director of the National Science Foundation supercomputing program, John Rollwagen, president of Cray Research and Ted Stern, executive vice-president, Westinghouse Electric.

Assured State-of-the-Art Computing

On August 11, 1988, the Pittsburgh Supercomputing Center and Cray Research, Inc. announced an agreement with a major impact on American science. The Center would receive the first production models not committed to government sites of Cray's next two supercomputer models, the Y-MP and CRAY-3. "This agreement takes National Science Foundation supercomputing several steps along the path laid out three years ago," noted Richard Cyert, president of Carnegie Mellon. "It's a bold and necessary advance for the national science effort."

Cyert with the full support of University of Pittsburgh president Wesley Posvar actively joined the Center's scientific directors, Mike Levine and Ralph Roskies, and Jim Kasdorf of Westinghouse in working out terms with Cray. Their participation represents one of the Pittsburgh Supercomputing Center's strengths, institutional support from two outstanding universities and a major American corporation.

In December, Cray delivered the Y-MP/832 to Pittsburgh. Cray technicians assembled "the Y" on-site and passed control to the experienced staff at Westinghouse, for whom the Y represented their fourth CRAY supercomputer installation. Two-hundred guests joined in the gala ribbon-cutting and dedication on March 7. With three times the power and four times the memory of the X-MP, the Y soon prompted a response from the research community nationwide. New applications for large-scale projects have confirmed what Levine and Roskies expected: researchers waiting in the wings with important work that simply exceeded the range of pre-Y computers.

"Our users tell us," says Roskies, "that they'll benefit from this upgrade by attacking new problems and by deriving new insights with a scaled-up approach to current problems." The CRAY-3 will amplify these benefits. The long-term nature of the agreement—two generations of supercomputing—is itself a major benefit, notes Levine: "It inspires confidence in the scientific community that NSF supercomputing is an ongoing national priority. It's here to stay, and it includes a commitment to state-of-the-art capability. This will stimulate fresh thinking in almost all branches of science and engineering."

JV

The Cray installation team at work on "the Y." Laced piping in the processor carries the liquid that cools the Y's circuitry.

UNICOS: The Great Migration

The Y-MP agreement gave new impetus to a project already well underway. By early 1987, the writing was on the wall: UNICOS is the future. This powerful systems software offers many advantages for scientific processing. Nevertheless, for users it meant a major change in computing environment. Under the direction of Assistant Director Marvin Zalevsky, the Center's systems and user services groups planned carefully to make the transition as smooth as possible. By early 1988, UNICOS became a "guest operating system." For four hours a day on one processor of the X-MP, it served "friendly users," researchers already acquainted with the software.

DA

Marvin Zalevsky, who coordinated UNICOS conversion and the overall Y project. "In one way or another, this effort involved nearly every person in the Center."

By mid-1988, the Center had begun an intense training effort—extensive on-line help and example files, workshops at the Center and at remote user locations. The objective: to move more than 1500 users to UNICOS by the time the Y takes over. By September, UNICOS was running in "native mode," fully occupying X-MP processing for several hours a day. In increments announced to the user community well in advance, the Center increased the UNICOS share of the X-MP's time step-by-step. In February, the transition took a big step with implementation of Common File System archiving. As of midnight, April 19, the great migration was complete.

KA

Lori Smith, user consultant, worked full-time on UNICOS for over a year. "These were major changes in software and hardware. We surveyed users at various stages, and the great majority felt the process went smoothly. I think we did a pretty good job of taking the bumps out of an inherently bumpy ride. Many users now say how much they like the new environment. We have a widely-used operating system on one of the most powerful computers available."

Graphics: Seeing is Knowing

"A substantial chunk of the brain is dedicated to seeing, but the amount dedicated to interpreting piles of numbers is relatively small," says Joel Welling. "If you take those numbers and make a movie, you can often see in a few seconds what would otherwise take hours or days." Welling, a physicist and one of the Center's scientific specialists, coordinates the Center's computer graphics. "We support a number of programs that produce a CGM file (computer graphics metafile), and GPLOT will feed the metafile through our animation system. For the cost of videotape, we'll send you a movie." Facilities added this year include a Titan graphics workstation, "a sort of graphics supercomputer."

GPLOT is an especially versatile graphics program written by Phil Andrews, scientific specialist with a plasma physics background. "GPLOT will read any CGM file a user creates and translate it to any of the Center's viewing devices." Andrews has adapted GPLOT for VMS, UNIX and UNICOS operating systems. It is the most general and probably fastest CGM translator

DA

Phil Andrews, author of GPLOT.

KA

Joel Welling at the August 11 press conference demonstrating computer graphics for Governor Robert Casey (seated) and several hundred thousand others—via the 11 o'clock news on Pittsburgh's Channel 11, WPXI. Wesley Posvar, John Rollwagen and Richard Cyert look on (standing left-to-right).

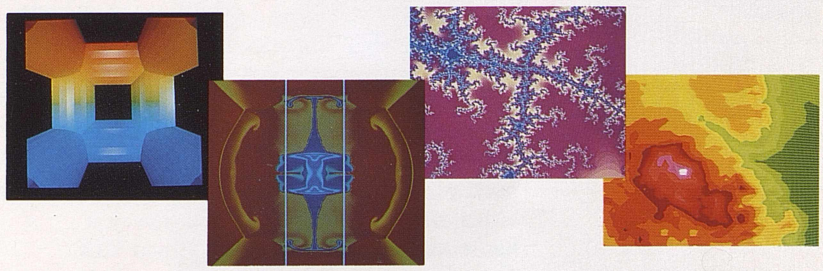

Images produced on the Center's graphics facilities: temperature inside a snow crystal (Mark Christon), galaxy formation (Allen Schiano, Arthur Wolfe & C. A. Chang), a Mandelbrot set (Matt Tolbert) and air pollution in the Los Angeles basin (Greg McRae).

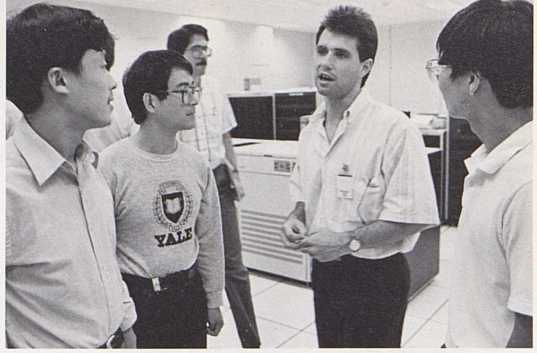
KA

Chuck Maiden leads a Summer Institute tour of Pittsburgh Supercomputing Center facilities at Westinghouse Energy Center. Maiden, a user consultant and training specialist, has instructed in many of the Center's educational programs. During the past year alone, these programs have included more than 50 seminars and training sessions with over 1500 participants.

Maiden is one among several full-time staffers who began their supercomputing careers as a student "hotline" consultant. "I was majoring in computer science at Pitt, and what I enjoyed about being here was how it pulled together all the various courses—like an ideal final class, except you don't get that in an academic environment."

The Center has employed more than 60 undergraduates. A number have moved on to research and industry positions that draw from their Pittsburgh Supercomputing Center background.

Distribution terms: free to anyone. At last rough count, over twenty sites—including several other supercomputing centers—had installed Andrews' code.

Networking: NSFNET & PREPnet

"Our users all over the country experienced noticeably better connectivity," says Gene Hastings, who manages the Center's communications and operations, referring to the NSFNET upgrade last July. The Pittsburgh Supercomputing Center is one of thirteen major hubs on the National Science Foundation's nationwide communications link, the so-called backbone, which connects over 700 local and geographical networks. "We dismantled our communications and reconfigured. Then we reconnected and were back in production by month-end. Throughput and stability have been excellent."

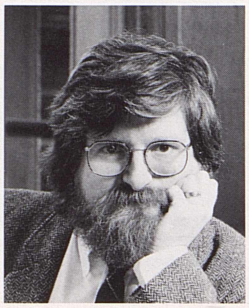
DA

Gene Hastings

Along with locally coordinating the NSFNET upgrade, Hastings played a lead role in designing the Pennsylvania Research and Economic Partnership Network (PREPnet), a statewide high-speed network that began operating last year. PREPnet connects business and research sites and seven universities to each other and to the Pittsburgh Supercomputing Center. Through the Center, PREPnet users gain access to NSFNET, ARPAnet (the Department of Defense's Advanced Research Projects Agency Network) and other regional networks.

From Science to Technology

Technology transfer—putting research knowledge to practical use—is an important part of what the Center does, says Ralph Roskies. It happens in several ways: through the Center's industrial affiliates, ALCOA being a prime example [p. 22], and through applied research [see pp. 8, 16-19, 24-31 & 44]. "Another vehicle, one which has been productive for us, is cooperative development projects with computer vendors." A software project with DEC led to the VAX SDE/Science package. With DEC and Cray Research, the Center tested the VAX Supercomputer Gateway, a hardware channel between CRAY and VAX computers.

The productive work with DEC expanded into a project called The Living Lab. "It's a real-life computing environment," says Chris Maher, who coordinates the Center's development work, "and it allows us to define and solve the problems you have with a mixed bag of machines and operating systems." DEC has stationed software engineer Dennis Carleton at the Center to work closely with Maher and others on the Living Lab project.

Other current development projects include work with Cray to improve UNICOS and the GAUSSIAN quantum chemistry software and projects with Ardent and IBM. "The results of these projects propagate back into the scientific community," says Levine, "through the increased variety and capability of vendor offerings. These products also boost vendors' ability to compete internationally and sustain the superiority of American high-performance computing."

User Training: The Summer Institute

"There I was, a statistician among engineers, physicists, astronomy people . . . " says Phil Wood. "It stimulates a lot of creative thinking." Wood, a post-doctoral research fellow at Penn State, attended the Center's second Summer Institute last July. "It's an incredibly fermenting experience, absolutely eye-opening to the capability of the machine."

Among Summer Institute participants, comments like Wood's unsolicited encomium are the rule rather than the exception. The two-week program, funded by NSF, "completely immerses participants in every aspect of scientific computing," says Janet Brown, the Center's Education Coordinator. Brown has planned and organized the Institute the past two years, and she's looking forward to another this July.


KA

1988 Summer Institute participants taking in the Pirates and Dodgers.

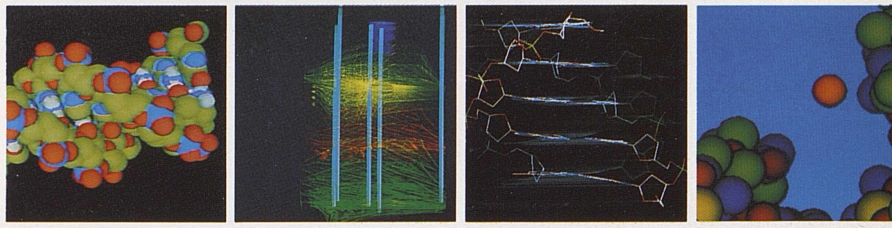

Molecular images produced by Hugh Nicholas & Dave Deerfield at the Pittsburgh Supercomputing Center. From left: DNA, macrophage, DNA, t-RNA and magnesium ion.

Biomedical Supercomputing

The Center's training effort this past year also included a series of workshops on biomedical supercomputing. The National Institutes of Health's Biomedical Research Technology Program funded these workshops through the only grant of this kind awarded to a supercomputing center. Two four-day workshops introduced supercomputing techniques and biomedical software. Other workshops focused on particular applications: molecular mechanics and dynamics, macromolecular structure refinement, and DNA and protein sequence analysis.

Hugh Nicholas and Dave Deerfield, who implement the Center's biomedical program, have planned another round of workshops this year that includes two new application areas: medical imaging and fluid flow in biological systems.

MF

Dave Deerfield & Hugh Nicholas, biomedical scientific specialists. Deerfield works in molecular mechanics and dynamics and supports the three major software packages in this field—Amber, CHARMM and Gromos. "We've also had a lot of success with Axel Brünger's X-PLOR program, which uses molecular dynamics to speed structural refinement."

Nicholas works in crystallography and macromolecular sequence analysis. "We're establishing a complete facility for nucleic acid and protein sequence analysis."

"Biological sciences traditionally have formulated theory without a lot of math and computing," says Nicholas. "But developments in molecular biology in particular have forced us toward computers just to deal with the quantity of information. You can't number crunch on a CRAY without learning how, and to be productive you have to learn quickly. That's what the workshops do."

Hardware, Software & Peopleware

"Through the Y-MP and CRAY-3 agreement," says Clayton, "our high-energy scientific directors have taken bold steps to bring hardware capability to Pittsburgh that will make this Center a magnet for research of the highest quality, which is why we're here. Yet many people may not appreciate that the foundation for everything we've accomplished is our staff—talented, motivated people. To do a major hardware-software conversion and bring it off smoothly requires the patience and experience to thoroughly plan, the technical competence to deal with unforeseen problems and the communication skills to effectively coordinate the many tasks and people. With the Y-MP under our belt, we'll start thinking about the CRAY-3 in a few months. It takes human resources every bit as leading-edge as our hardware to be able to do this, and we're proud of that."

Pittsburgh Supercomputing Center's CRAY Y-MP/832 and the X-MP/48 that it replaced. The Y-MP makes greater computing capability available for both new and experienced users. Currently, 40% of the Center's active grants of computing time are starter and training grants to users with little or no supercomputing experience. These grants account for only 3% of total usage, allowing the Center to train new supercomputing researchers without impairing progress on larger projects.

LR

Order Out of Chaos.
Among the jumble of molecules in a cell, each enzyme selects its own specific reactant or substrate and converts it to products which are the next intermediate in the metabolic chain. Enzymes bind with their substrates at pockets or creases in the protein chain called the active site. Enzyme catalysis is often inhibited by non-substrate molecules, inhibitors, which tie-up the active site, making the enzyme unavailable to the substrate.

A Model of Self-Control

Structure and Function of Aspartate Transcarbamylase

J. Eric Gouaux & William N. Lipscomb, Harvard University

(a)

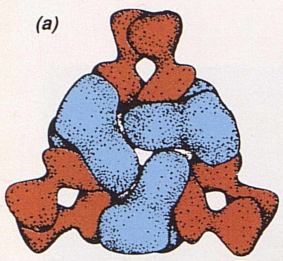

(b)

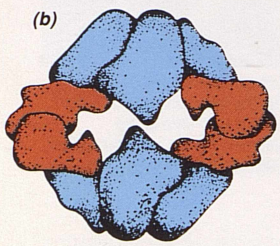

Structure of Aspartate Transcarbamylase
The six R-chains (red) are arranged in three two-chain groups (dimers) sandwiched in a mid-layer between two three-chain groups (trimers) of the C-chains (blue). In (a) the threefold symmetry (three 120° rotations) of the R-chain pairs is evident. In (b), with the structure rotated 90°, the pairing between C-chains in different trimers can be seen.

Enzymes and Egg White

Because we have enzymes, we live. Without them biological reactions would happen so slowly as to be almost nonexistent. Nearly all enzymes are proteins acting as highly specific, highly potent catalysts. Each one usually affects only one chemical reaction, and that only in a narrow range of temperature and pH. When these conditions are met, enzymes accelerate the reaction a million times or more. Consider egg-white boiled twenty hours in a vat of concentrated acid. That's the equivalent of only two hours in the human digestive tract—no boiling or concentrated acid, just pepsin and a few other enzymes to catalyze the reactions.

"We're interested in enzyme mechanisms," says William Lipscomb, "how the active site functions and, in particular, why enzymes speed up reactions by such large factors." Lipscomb travels two parallel paths in his research at the Pittsburgh Supercomputing Center. With John Stanton and Rodney Bartlett, he has extended his career-long investigation of boron hydrides, work which garnered a 1976 Nobel Prize. He also applies his interest in structure-function relations to large proteins, especially enzymes. "Our major work since 1967 has been the allosteric mechanism in aspartate transcarbamylase" (ATCase). With the CRAY X-MP as catalyst, Lipscomb and co-worker Eric Gouaux have produced new knowledge of this important enzyme system.

ATCase and Allosteric Regulation

When binding at one site changes another site on the same enzyme, the interaction and the change are *allosteric*. These "cooperative binding" interactions, as they're also called, either increase or reduce substrate binding at the second site. In effect, they're a control mechanism—a second level of inhibition and activation—by which the enzyme regulates its own activity.

Allosteric Interaction
Binding of a molecule at one site in the enzyme leads to conformational changes that propagate to a distant site.

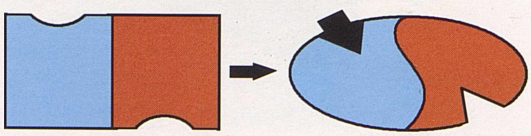

ATCase, for instance, is a large, intricately symmetric protein with six catalytic or C-chains and six regulatory R-chains. Each R-chain pairs with a C-chain, and each C-chain relates structurally to another C-chain. ATCase catalyzes a critical metabolic reaction, one that begins a committed path to production of pyrmidines, essential for DNA replication. If this path overproduces, its product—cytidine triphosphate (CTP)—in effect circles back to the R-chains, where it binds and triggers inhibition at the corresponding C-chain active site.

ATCase in Pyrmidine Biosynthesis
The metabolic pathway catalyzed by ATCase. The reaction is feedback inhibited by CTP.

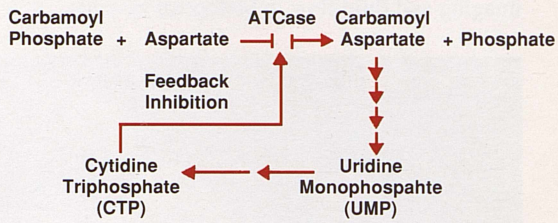

It's a sophisticated feedback control loop. CTP competes with its purine counterpart ATP for access to the R-chain site. When CTP in the cell is low, ATP binds to the R-chain. This activates the C-chain and boosts formation of CTP. Remarkably, these allosteric interactions occur over a distance of about 60 angstroms, halfway across the protein. "The question," says Lipscomb, "is how does the information get propagated from one site to another?"

All-or-Nothing or Step-by-Step

CTP and ATP bind loosely to the regulatory site, notes Gouaux, and cause only subtle, difficult-to-detect structural changes. ATCase's other cooperative interaction, which occurs among C-chain active sites, shifts the enzyme dramatically from one structure to another. As a result, it has been more responsive to the entreaties of the Harvard researchers. The shifted structure—the R (relaxed) state—is a much more active catalyst than the initial structure—the T (tense) state, and the shift in effect turns the raeaction on. "They're really just complicated switches," says Gouaux, "and

we're trying to figure out how molecularly they're thrown." With Evan Kantrowitz of Boston College, one of his former students, Lipscomb reported a speculative model of this mechanism last August in *Science*.

Substrate binding at one C-chain tightens its active site and rearranges the interface with its C-chain counterpart. Both chains shift to the R-state, and because of the enzyme's symmetry, the other chains follow. The transition is "concerted"—binding at one site converts the whole enzyme. "It's all or nothing," says Lipscomb of his model, "not much in between." The article also noted that research (by Gouaux) hinted at an intermediate between the T and R state. "It's remarkable," says Lipscomb. "It wasn't expected from any previous work." Their recent work has shown more evidence of intermediates, findings that tend to qualify the fully concerted model. Lipscomb nevertheless believes the shift is still best described primarily in terms of two states.

Practical Results of Enzyme Research

Lipscomb's work with enzymes has exemplified fruitful interaction between basic and applied research. Phosphonacetyl-aspartate (PALA), an ATCase inhibitor that Lipscomb and Gouaux have used to bind the active site of ATCase, has been released as a cancer treatment. "There are about 3000 genetic diseases, each of which is a failure in an enzyme or some other protein," says Lipscomb. "If you take a critical enzyme which is regulated and you can control it, then you gain some control over the metabolic pathway." In earlier work, scientists at Upjohn developed Captopril, the best-selling drug for high-blood pressure, from Lipscomb's structural work on carboxypeptidase, a digestive enzyme. "If you really know the structure of enzymes and how they work, you can design inhibitors that are much more effective."

References:

Evan R. Kantrowitz & William N. Lipscomb, "Escheria coli Aspartate Transcarbamylase: The Relation Between Struture and Function, *Science* **241**, 669 (1988).

J. Eric Gouaux & William N. Lipscomb, "2.5 A Structure of a Single Amino Acid Mutant of Aspartate Carbamoyltransferase: Implications for the Cooperative Mechanism," *Biochemistry* (1989), in press.

National Institutes of Health grant GM06920 supported this research.

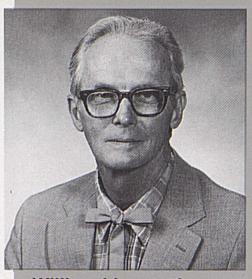

William Lipscomb

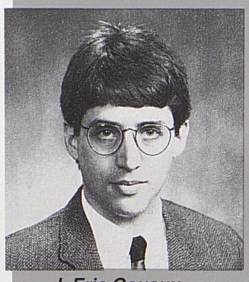

J. Eric Gouaux

The Role of Supercomputing in Enzyme Research

Lipscomb and Gouaux have used Axel Brunger's XPLOR program to great advantage as a complement to "site directed mutagenesis." The primary structure of a protein, coded by the organism's DNA, is its linear sequence of amino-acids. A single substitution in that sequence can sometimes radically alter the enzyme. Refining the X-ray crystallography data from laboratory mutants with XPLOR and comparing the structures to the native enzyme allows invaluable insights to structure-function correlations. "It's a really useful technique," says Gouaux. "There aren't surprises all the time, but they're certainly frequent enough to make one realize we don't understand these systems thoroughly yet."

XPLOR is computationally demanding, about 100 hours a run notes Lipscomb. "But we'd never be able to do these calculations without it. We can work on much more complex problems, problems we never would have tried, when we know the supercomputer is available. That's the real place that it contributes."

T-State and R-State of ATCase

The two states are shown along one of the two-fold axes of symmetry, C-chains in yellow, R-chains red. Along the threefold (vertical) axis, the R-State is 12 angstroms longer than the T-state.

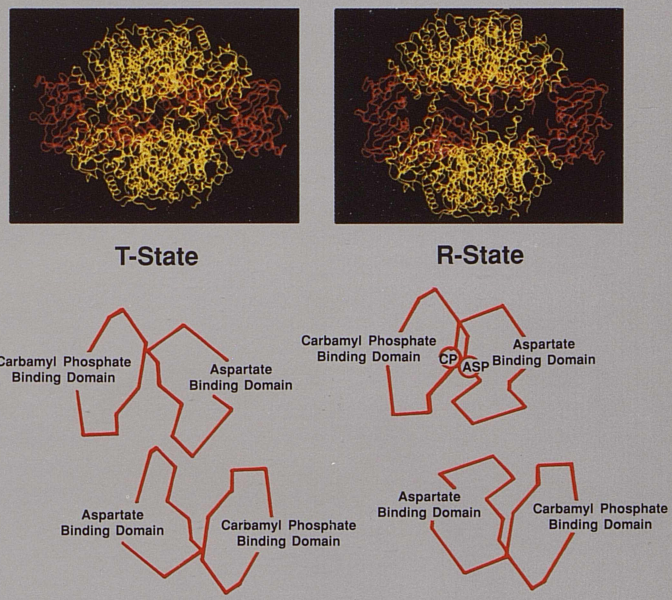

William Lipscomb's model for the cooperative interaction between two C-chain subunits as they shift from the T- to the R-state when the substrates, aspartate (Asp) and carbamyl phosphate (Cp), bind at one T-state active site. Lipscomb concludes that substrate binding triggers a primary level (amino-acid residues) interaction at the interface between these C-chain related pairs that is key to the shift.

Allosteric Regulation of CTP: Bound Configuration

The allosteric site (extreme left) of an R-chain (yellow) is bound to a CTP molecule. About 60 angstroms away, the C-chain active site (blue & light green) is bound to phosphonoacetyl aspartate (PALA), the bisubstrate analog of carbamyl phosphate and aspartate, the two substrates which ATCase catalyzes.

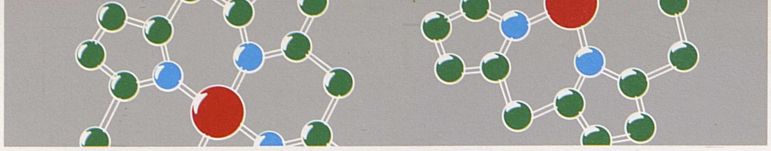

With This Ring

Theoretical Investigations of Heme Proteins

Gilda H. Loew, Jack Collins & Frank Axe, SRI International and The Rockefeller University

Long-Haul Proteins

Hemoglobin is the best-known member of a ubiquitous family of globular proteins called heme proteins. It is also the most fittingly named. "Heme" comes from *haima*, Greek for blood, and hemoglobin, which resides in the red blood cells, carries a precious cargo on its journey along the arteries. Its single iron atom, centered in a porphyrin ring, binds to oxygen in the lungs and discharges it in the capillaries before circling back via the veins to recharge. This iron-oxygen bond gives blood its deep-red pigment, probably the strongest visceral signal in the entire vocabulary of color.

In the late 50s, Max Perutz and John Kendrew at Cambridge won a Nobel prize for determining the structure of hemoglobin and its cousin, myoglobin. These were the first protein structures solved, and research since has from this base built a formidable, though far from complete, body of knowledge. "Heme proteins are quite unusual," says Gilda Loew. "Their active site has no amino acids. Instead it's a metal-organic complex unlike any other part of the protein." Loew, who directs molecular theory at SRI International, has for over twenty years investigated structure, function and spectra relations in heme proteins. Using Pittsburgh's CRAY, she and colleagues Collins and Axe have identified factors that help to explain a unique puzzle in the heme protein family tree, the relation between myoglobin and cytochrome C peroxidase (CCP).

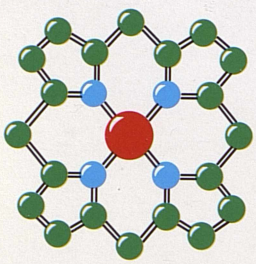

The Heme Unit
The heme, active site of heme proteins, is a disk-shaped, intricately symmetric structure formed of an organic part, the porphyrin ring, *and an iron atom in its center. In the heme plane, the iron bonds to four nitrogen atoms, each of which forms a* pyrrole ring *with four carbons. The four pyrroles link to form the porphyrin ring. The iron can also bond on each side of the heme plane.*

Probing the Structure-Function Enigma

CCP presents a paradigm case among heme proteins. It functions as an enzyme even though its heme unit is virtually identical to myoglobin. "One of the enigmas about heme proteins," explains Loew, "is that they all have essentially the same active site, the heme unit, and their biological function centers on this unit, most often on the iron itself, and yet they have three distinct biological functions." Hemoglobin and myoglobin exemplify the oxygen transport group. The cytochromes are a second group, important as electron transfer agents in the respiratory chain. The third group, prominently CCP and cytochrome P450, are enzymes, catalyzing various metabolic reactions.

With active sites so similar, why do they act so differently? Several factors may be involved: 1) the number and type of *axial ligands* — molecules that bind with iron above or below the heme plane, 2) the oxidation state of the iron — usually either ferrous or ferric (+2 or +3 valence), 3) the number of electrons with unpaired spins on the heme unit and 4) amino acid residues near the heme unit. "Differences in any one or all of these can result in very diverse chemical behavior. But how these factors combine to bring about a particular function isn't clearly understood." Loew's ongoing work aims at understanding these interactions, knowledge that can form the basis for designing more effective, less toxic drugs.

Results: Structure Determines Function Determines Structure . . .

Loew and her coworkers frequently use quantum chemistry techniques to calculate the electromagnetic properties of heme proteins. Like biological function, these properties center on the active site. Quantum studies of the heme unit therefore can establish a one-to-one correspondence between structure and observed spectroscopic properties. Calculated molecular properties then can help explain function. Loew has found that heme proteins are subject to sensitive interactions between electromagnetic properties and function. "It's an exquisite fine-tuning of these subtle properties — like electronic, spin and oxidation states. Often they change during function and determine function."

Loew's attention to these interactions highlights her work with CCP and myoglobin. The two heme units are nearly indistinguishable, geometrically and chemically. This includes their axial ligands — water on the upper site and imidazole, a side-chain from a histidine amino acid, on the lower. "We made a thorough study of the heme unit itself and concluded that the small differences we found don't sufficiently explain these differences in function." Among other findings, Loew's study confirms and provides insight to experimental indications that a negative charge on the imidazole, not present in myoglobin because of different nearby amino acids, may be a factor in the functional difference.

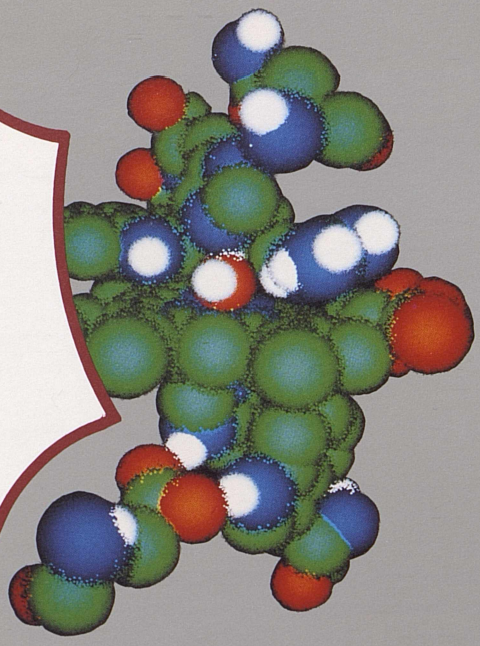

This space-filling model indicates actual packing of the molecules for the same structure show in the line model.

Three Representations of the Active Site

A water molecule (H red, O white) is shown bound to iron (orange) at its distal site. The spherical surface represents the water's van der Waals energy surface. This H_2O is bound in a network of hydrogen bonds to three other water molecules (red & white), and four polar amino acid residues (green with H in red). This polar network, absent in myoglobin, is crucial to CCP's function.
On the proximal side, the five-membered axial ligand is imidazole, side-chain to a histidine residue which interacts via hydrogen bond (dotted yellow lines) to an aspartate anion hydrogen, bonded in turn to another amino acid residue, tryptophan, to its right. Loew's research has supported experimental work indicating that this aspartate interaction with the histidine ligand is also crucial to the enzyme function.

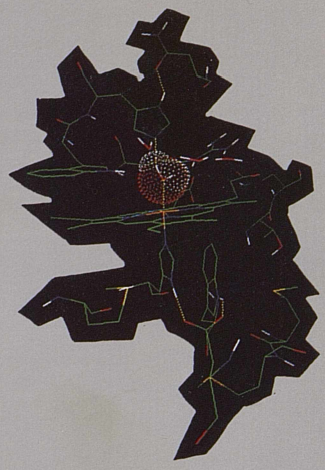

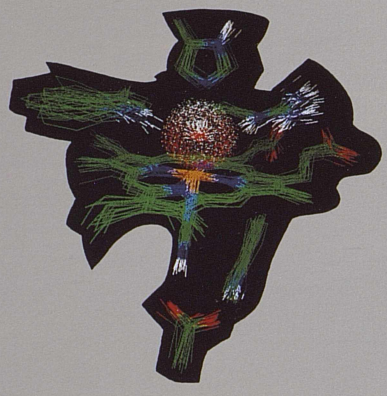

A 20 picosecond (one-trillionth of a second) molecular dynamics simulation using AMBER with a snapshot each picosecond. "Note that despite the flexible movement of the water ligand and the amino acids in the binding pocket," say Loew, "they maintain their H-bonding interactions." On the proximal side, the protein also maintains the crucial H-bonding from histidine to aspartate (oxygen in red).

Cytochrome C Peroxidase: The Resting State and the Transient Peroxide Complex

Cytochrome C peroxidase, like other peroxidases, performs its enzyme function only in an unstable, intermediate form, the ferryl state, to which it transforms in stages from its low-energy, resting state. First, hydrogen peroxide (H_2O_2) replaces the water molecule bound to the distal site of the heme unit. The peroxide then rapidly loses water until only a single oxygen atom is attached. This is believed to be the first characterization of this intermediate state.

Gilda Loew and Jack Collins

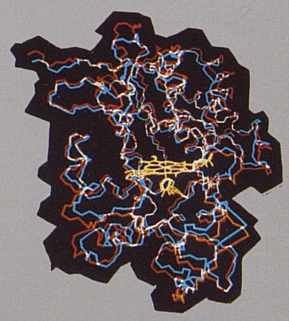

Comparison of the X-Ray Crystal and Energy Optimized Structure
The energy-optimized structure (blue) of the protein resting state overlays the X-ray crystallographic structure (red) with the heme unit in yellow.

The Supercomputer as Laboratory

"Computers are our laboratory. The numerical experiments we do complement actual experiments. Experiments can measure the properties of the active site, but this data alone doesn't allow us to explain the origins of biological function." Supercomputing contributes in two ways. It allows structure studies of the entire protein and in-depth investigation of the active site. "Even for heme proteins, where there's such a well-defined active site, we're finding that we can't ignore the protein. To do this, we need large-scale computational capability. Only then can you look at huge systems like the heme proteins with 2500 atoms. We also need to continue looking at the active site alone. Supercomputing lets us do this in much more depth."

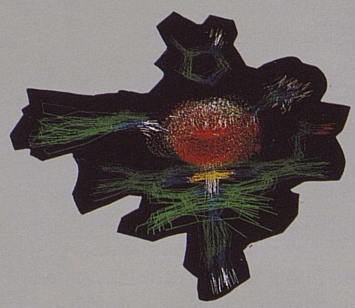

The Transient CCP-Peroxide Complex: Dynamic Model of the Active Site
This peroxide state is so transient, notes Loew, that it's never been isolated experimentally. This is the first attempt to characterize it. Loew replaced the H_2O distal ligand with H_2O_2 and re-optimized the structure. The simulation is 20 picoseconds, a snapshot each picosecond. "The greatly increased flexibility of the peroxide relative to the water in the resting state is apparent. Amino acids surrounding the peroxide and the heme unit also appear more flexible."

References:

F. U. Axe, A. Waleh, L. Chantranupong & G. H. Loew, "A Comparative Analysis of the Electronic Structure and Molecular Properties of the Resting States of Cytochrome C Peroxidase, Metmyoglobin, and Catalase," *Int. J. Quant. Chem.* **35**, 181 (1989).

F. U. Axe, L. Chantranupong, J. R. Collins & G. H. Loew, "Ground State Properties of Heme Complexes in Model Compounds and Intact Proteins," *Computational Chemistry, The Challenge of d- and f- Electrons* (Am. Chem. Soc. Symp. Series, 1988), in press.

NSF grant PCM8410244 and NIH grant GM27943 supported this work.

BIOLOGICAL SCIENCE

A Germinal Model

Simulation of *Bacillus Subtilis* Growth and Differentiation

Jin Wook Jeong & Mohammad M. Ataai, University of Pittsburgh

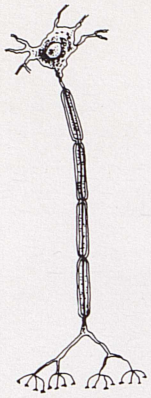

The uniformity of the earth's life, more astonishing than its diversity, is accountable by the high probability that we derived, originally, from a single cell. It is from the progeny of this parent cell that we take our looks; we still share genes around, and the resemblance of the enzymes of grasses to those of whales is a family resemblance.

—Lewis Thomas, The Lives of a Cell

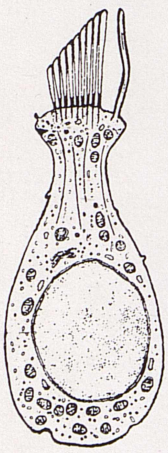

Bacteria and Human Cell Division

Twenty hours at 212° F. and still alive? Nothing unusual about that, not for ornery *Bacillus subtilis*, a one-celled bacterium with survival skills that make it a practical problem in milk processing as well as an intriguing, important research subject. When the temperature goes too high or the amount of food too low or too much of a toxic chemical enters *B. subtilis*' environment, it does what amounts to a dramatic tactical retreat—it sporulates. The life process draws inward and condenses within an extremely hardy form. The "mother cell" self-destructs as it releases this new form, the spore, which bides its time waiting for a chance to grow again.

Since the 1960s, *B. subtilis* has been the subject of extensive experimental work. The underlying premise has been that sporulation is a basic case study in cell differentiation, a process that's involved in the development of all higher life forms. All cells come from pre-existing cells, but some cells create new cells different from themselves. "If we understand differentiation in its simplest form," says Mohammad Ataai, "in *Bacillus* for example, we should be able to look at higher cellular systems—like our body. If you want to replace cells that don't function properly, to know how the cells differentiate is extremely important."

A Comprehensive Cell-Growth Model

A biochemically-oriented chemical engineer, Ataai has been modeling bacterial growth since his 1986 thesis with Michael Shuler at Cornell, and he has now joined supercomputing to his knowledge of *B. subtilis*. Experimental work with this bacteria has sought to isolate a triggering mechanism for sporulation. Some results point at nucleotides, yet interpreting the data plunges researchers into a jumble of inter-related processes. "Except in obvious cases, it's almost impossible to sort out and assess the scenarios that can occur from changes in this system. A mathematical model can give us a way to look at components in isolation." Ataai believes his model has the potential to give a coherent picture of experimental data and, even more importantly, to allow investigations of *B. subtilis* which can't be done experimentally.

"As far as we know, it's the most comprehensive model of bacterial growth that's been attempted. And it's the first attempt to model the growth process of *Bacillus subtilis*." The model has 39 components, each representing a distinct reactant, product or intermediate in the reactions that regulate cell growth. It includes major metabolic pathways such as glycolysis, purine metabolism, the Krebs cycle (TCA cycle), and gluconeogenesis. The differential equations that represent net rate of production for each component include a biosynthesis term for increase, a degradation term for decrease, and terms for reactions with other cellular materials.

Experimental work has produced a pool of data from which Ataai has drawn in developing the 200 rate constants for his equations—a formidable and inexact process. A deficiency of other cell growth models, in Ataai's view, has been that the demanding computations to evaluate these parameters haven't been done. For this purpose, he has written code that does a sensitivity analysis simultaneous with solving the rate equations. This tells him which constants have the greatest effect on the outcome and to what extent changing their value will vary the results.

Simulating the Passage to Sporulation

Once a *B. subtilis* cell develops its growth enzymes and has sufficient nutrients, it reproduces by cell division—the exponential growth phase—until nutrients deplete or toxic materials accumulate. At this "stationary phase," usually preparatory to sporulation, reproduction stops. During this transition, the biochemistry of the cell becomes radically unstable. Ataai sees simulating this transient period as the most demanding test of his model. Results from initial work have been encouraging.

Ataai modeled cell growth in a culture with glucose as food source. Though initially ample, the glucose depleted with time, and as it dropped to essentially zero, exponential growth stopped. Computation time increased dramatically as the time-steps compressed in response to rapid fluctuations of the model components. DNA, RNA and protein concentrations generally agreed with experiment as did other key components. Some results differed from experiment, and Ataai

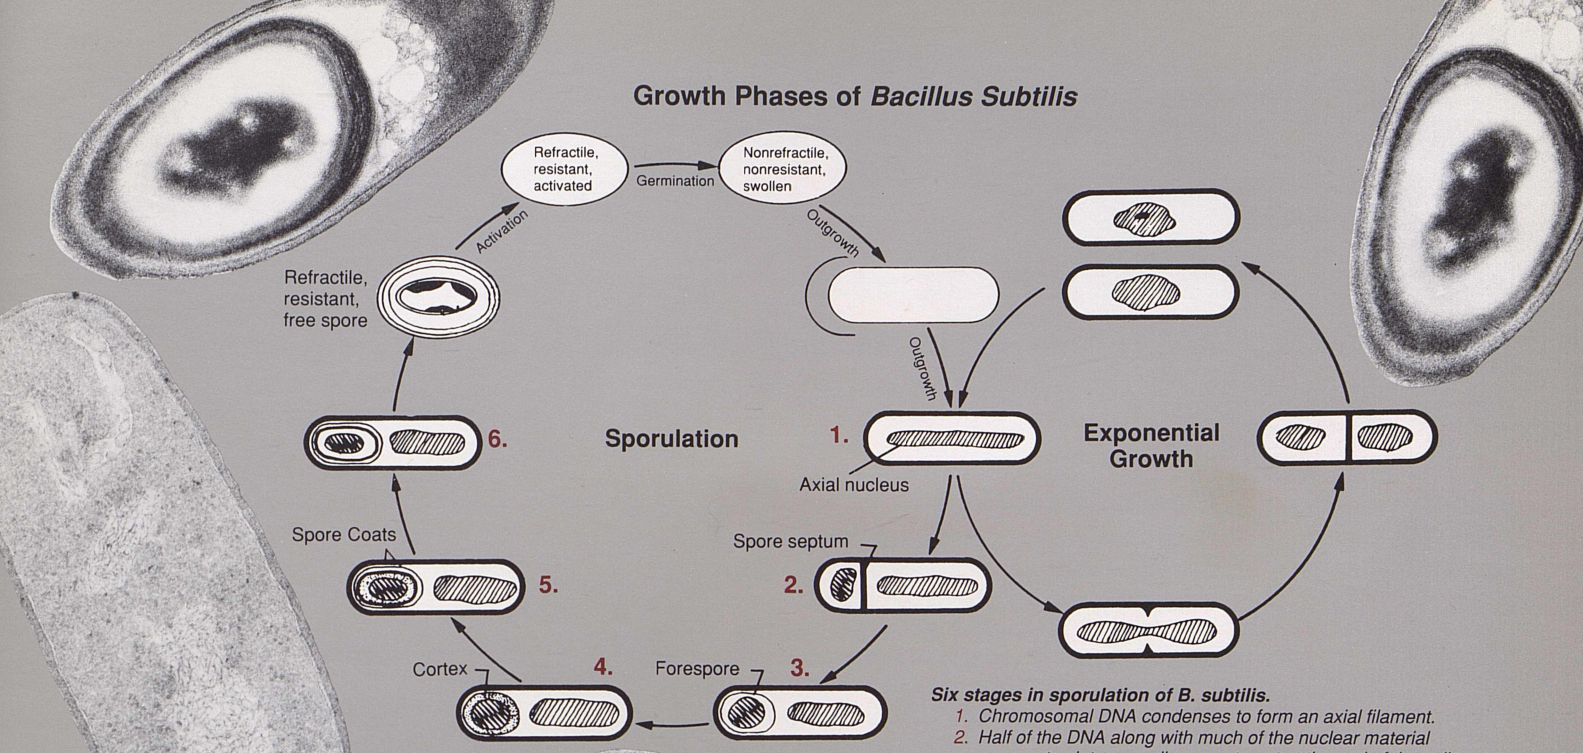

Six stages in sporulation of B. subtilis.

1. *Chromosomal DNA condenses to form an axial filament.*
2. *Half of the DNA along with much of the nuclear material segregates into a small compartment at the end of the cell, surrounded by the spore septum, which generates from the cytoplasmic material.*
3. *The compartment develops into a primordial spore.*
4. *New envelops begin to form around the spore, including the cortex and one or two spore coats.*
5. *The spore goes through a series of chemical alterations which increase its ability to resist heat, dessication and toxic chemicals.*
6. *Through autolysis, a self-degenerative process, the "mother cell" releases the spore.*

will continue to analyze these. "But we can simulate the transition from exponential growth to the stationary phase," says Ataai, "and this is very difficult. There's no numerical instability. And the fact that we can do this means we should be able to see which changes cause differentiation."

Next: Modeling Laboratory Mutations

The next test is mutant cells. If the model gives good predictions in these cases, which will involve modifying the *wild type* (as found in nature) reaction rates, it would "strongly support the model characteristics." Existing experimental data isn't broad enough to unambiguously evaluate these tests, so Ataai has planned his own experiments, a process he expects to last at least another year. "It's through the interaction of modeling and experiment that a unified picture of experimental work can emerge."

Supercomputing plays an essential role. "The sensitivity analysis couldn't be done without the CRAY." Because of the sheer number of rate constants, this calculation makes major demands on computing. For the simulations themselves, Ataai's well-vectorized code takes slightly less than an hour per run. Pittsburgh's newly installed CRAY Y-MP should assure that Ataai has the computing capability he needs to realize the potential of this ambitious and innovative project.

Mohammad Ataai

Figure 1. Cell Growth, Glucose & Macromolecule Concentration versus Time.
The first two hours, 0 through 2, correspond to the last two cycles of exponential growth. Cycling in RNA concentration follows the cyclic process of the exponential phase. At about 2 hours, extracellular glucose concentration becomes very low, and cell growth rate decreases, indicated by the leveling of cell mass. From 2 to 3.5 hours, the cell undergoes a transition from exponential to stationary growth. During this transient period, the model predicts intracellular concentrations for all the components modeled.

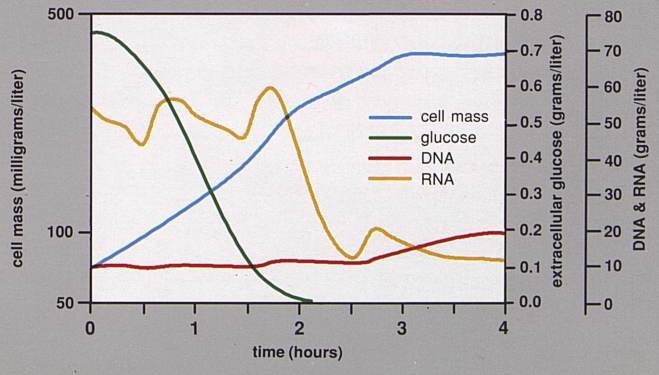

Reference:

J. W. Jeong, J. Snay & M. M. Ataai, "A Mathematical Model for Examination of Growth and Differentiation in Bacterium *Bacillus Subtilis*," *Biotechnology and Bioengineering* (1989), in press.

NIH Division of Research Resources grant U41RR04154 and a University of Pittsburgh Young Investigator start-up award support this research.

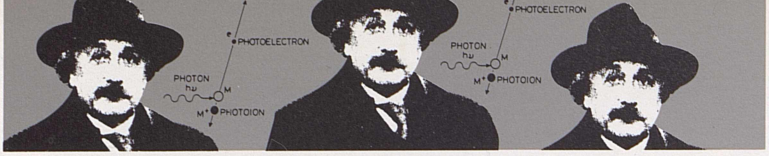

Charge of the Light Particles

Dissociation Dynamics of Energy Selected Ion-Dipole Complexes

J. Charles Morrow, University of North Carolina, Chapel Hill

GAUSSIAN: User Friendly Quantum Mechanics?

It's convenient that Doug Fox, a Scientific Specialist at the Pittsburgh Supercomputing Center, does his research with John Pople's group at Carnegie Mellon. The chemistry department and Pople's office on the 8th floor of Mellon Institute are only an elevator ride from Fox's work station connection to the CRAY. This relation has also been convenient for Pittsburgh Supercomputing researchers using GAUSSIAN, though convenient might not be Fox's term-of-choice. He feeds on a steady diet of GAUSSIAN trouble-shooting, strategizing and hand-holding. Nevertheless, he insists that one of the program's strengths is user accessibility.

"In terms of a pure programming approach, the choice was made long ago, a decade or so, to make this program accessible to people who aren't really theorists, people who haven't spent their lifetime figuring out how to use it. It's keyword driven. The user says "OPT" when he wants to do an optimization, and he says "MP" if he wants to do a Møller-Plesset calculation. It's very user friendly in that sense. And it does a fair bit of checking to prevent you from going off half-cocked doing meaningless calculations. GAUSSIAN is very good at looking at the lowest energy state of a molecule."

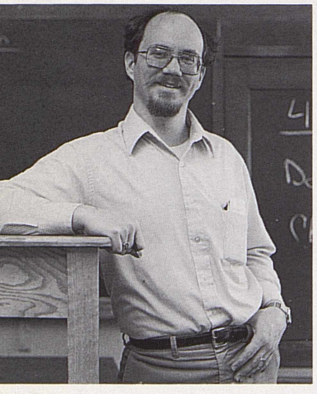

Doug Fox

Photons, Ions and PEPICO

Put some gas in a vacuum-sealed chamber and zap it with a beam of ultraviolet light. Since light is made of particles—photons, as Einstein named them in 1905, the beam striking a gas molecule is roughly like the breaking shot in a game of pool. Since light is also a wave, the beam's energy depends on its frequency, and if the energy is high enough—more than the *ionization potential* of the molecules—a photon colliding with a gas molecule will dislodge an electron. The molecule will become a positively charged ion.

Photoelectron photoionization coincidence spectrometry (PEPICO) is a relatively new experimental method based on this phenomenon. At the University of North Carolina, Charles Morrow has brought his quantum chemistry background into collaboration with one of only several groups in the world using PEPICO. With his colleague Tomas Baer, who leads the PEPICO group, Morrow has been able to map detailed reaction pathways for several ion systems. Their work on ethyl chloride and 1-propanol has revealed structural forms—ion-dipole complexes—not previously identified in the reaction path of these ions, a discovery that sheds new light on the anomalously slow *dissociation rate* of so-called *metastable* ions.

Dissociation Rates and Metastable Ions

Basically, the PEPICO apparatus measures dissociation rates—how long it takes ions in the vacuum chamber to decompose into fragments. Often an ion rearranges into isomers, different structures with the same atoms. With more energy, the ions and their isomers dissociate into non-isomeric forms, fragments which may also form isomers and further decompose. An ion's dissociation rate is usually an index of stability—slow dissociation indicating high *activation energy*, the energy needed to trigger dissociation. Metastable ions, however, dissociate slowly even though their activation energy is low—a phenomenon that has not been well explained.

Dissociation rates from PEPICO lab work are Morrow's raw material. It is data that poses as many mysteries as it solves. "It's as though I went to the race track and told you I saw cars that went 1200 mph," says Morrow. "You'd say *That's remarkable, what kind of car was it?* And I'd say *I don't know, it went too fast for me to see it.*" The ions in the PEPICO chamber have extremely low concentrations and very short lifetimes, typically a millionth of a second or less. To identify these ions with their measured reaction rates, Morrow relies on Rice-Marcus theory, an accepted theory for unimolecular rate of decay, and quantum mechanical calculations.

PEPICO, GAUSSIAN & CRAY

Based on what they know of the chemistry, the researchers make an informed guess about what ion produced an observed rate. Applying rate theory, they can calculate a rate and test their guess against the PEPICO data. But the calculation requires the activation energy and vibrational frequencies of the hypothesized ion; for this information, Morrow turns to *ab initio* computations with GAUSSIAN—"the great workhorse of quantum chemistry," says Morrow with a nod to John Pople and colleagues at Carnegie Mellon, who developed this package of programs.

From a molecule's basic geometry—the distance and angles between its bonded atoms—GAUSSIAN determines its potential energy, essentially the energy to maintain its structure. It uses the Hartree-Fock method, a familiar approach that calculates each electron's interaction with the rest of the molecule as an averaged field. GAUSSIAN users specify a *basis set* — a set of single-electron functions that in sum approximate the quantum mechanical wave function for the molecule. Morrow also uses GAUSSIAN's capability to fine-tune electron interactions with Møller-Plesset perturbation theory.

Morrow systematically varies geometry until he arrives at an energy-minimized configuration, the most stable form of a molecule. From this ground state, he can compute activation energy and vibrational frequencies—"fiendishly complicated computationally." It's here, he says, where the CRAY goes hand and glove with PEPICO. "No Pittsburgh Supercomputing Center, no activation energy, no anything. One of the great joys of the whole NSF enterprise is it gives you a chance to find these properties. Without supercomputing, you don't have a chance." Compared to earlier work, Morrow's computations show more accurate

"Look at it as if you're going over a mountain range, from one valley to another. Say you're in Denver and you want to go to Grand Junction. You have to go up before you go down, and you wind up in Grand Junction at a different altitude than Denver. The difference is dissociation energy, but you have to go over this barrier, an activation barrier, to get there. And from where you are to the top of the path is the activation energy."

— ***Charles Morrow***

relative energies and charge and spin data, a result he sees as directly related to supercomputing capability.

Results: Metastability and Ion-Dipole Complexes

Morrow's work has revealed structures, energies, and vibrational frequencies of ten ionic species of the metastable ion ethyl chloride ($C_2H_5Cl^{+\cdot}$). "The dissociation of the ethyl chloride ion by the loss of HCl looks like a simple reaction, but experimental results are hard to reconcile with a simple mechanism." Morrow found energy-minimized structures for each of five relatively stable isomers. He also found structures for the five related transitional states—energy peaks or "saddle-points" which must be crossed along the reaction pathway into the next low-energy valley. For two of the isomers and four of the transition states, Morrow's findings are the first reported.

He also found that one of the five low-energy isomers is an ion-dipole complex, ($C_2H_4^{+\cdot}$ - ClH). In these structures, an ion bonds weakly to a dipolar molecule, usually at an extended bond length. Earlier studies mischaracterized this structure. Morrow's results with ethyl chloride led him to explore a similar rate anomaly for the propanol ion ($C_3H_7OH^{+\cdot}$). Again he found a previously unreported ion-dipole complex, ($C_3H_6^{+\cdot}$ - OH_2). In both ion systems, Morrow concludes that metastability relates directly to the ion-dipole complex he discovered. In ongoing work, he is extending his analysis of this data.

References:

Jian-Dong Shao, Tomas Baer, J. C. Morrow, and Maria L. Fraser-Monteiro, "The dissociation dynamics of energy selected ion-dipole complexes. I. The cyclopropane ion-water complex [c-$C_3H_6^+$-OH_2]," *J. Chem. Phys.* **87**, 5242 (1987).

J. C. Morrow and Tomas Baer, "Ab Initio Calculations on $C_2H_5Cl^+$ and the Loss of HCl from $C_2H_5Cl^+$ isomers," *J. Chem. Phys.* **92**, 6567 (1988).

This research is supported by NSF grant CHE-8405743.

Reaction Pathway for the Photo-Dissociation of Propanol, C_3H_7OH

At a calculated activation energy of .47 eV, the ion begins to dissociate into cyclopropane, $C_3H_6^+$ and water. Morrow's calculations revealed the ion-dipole complex in an energy well after the activation barrier.

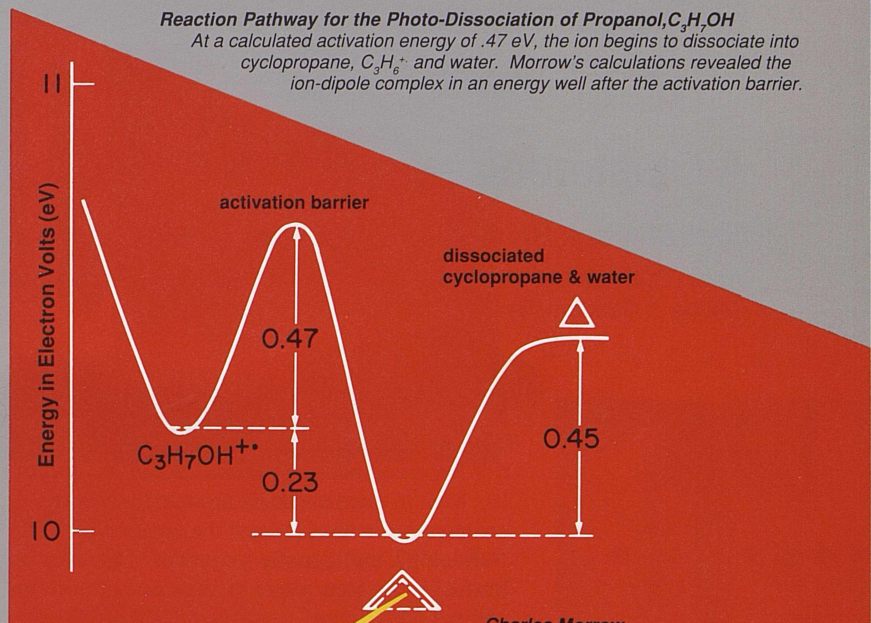

Charles Morrow

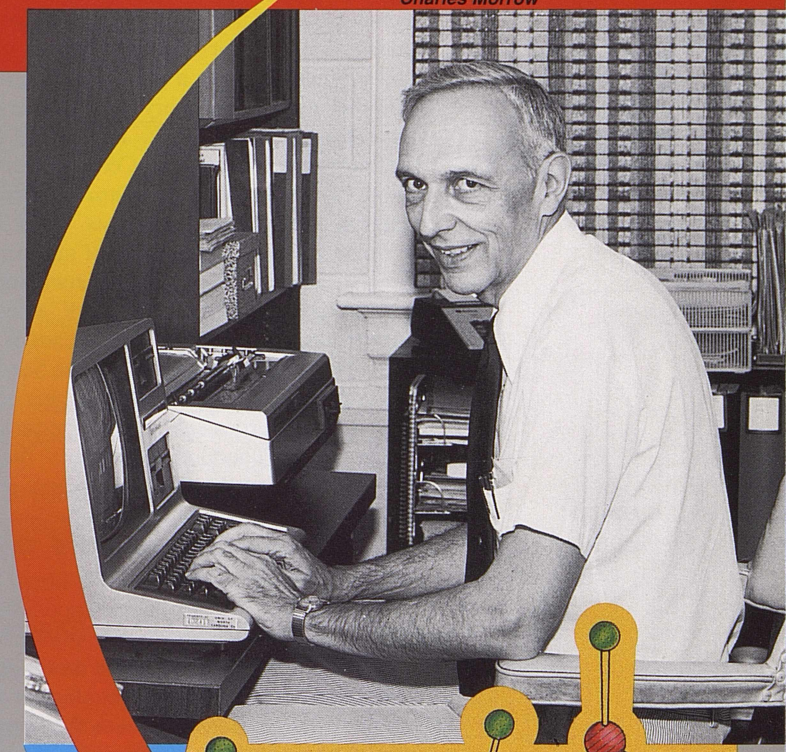

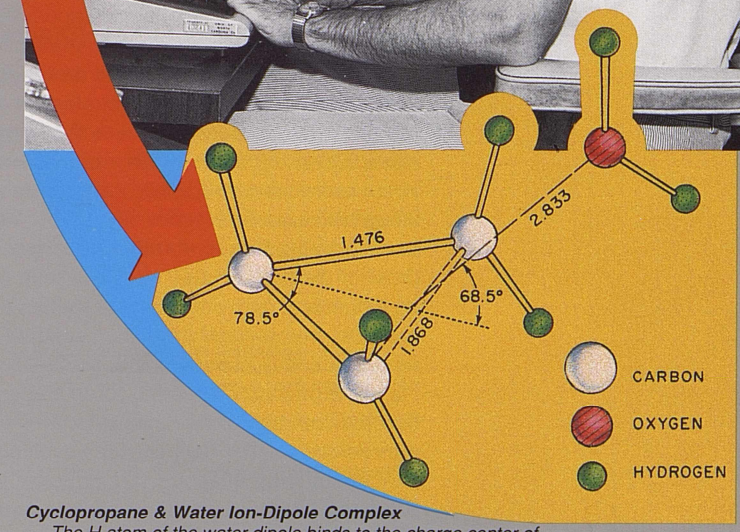

Cyclopropane & Water Ion-Dipole Complex

The H atom of the water dipole binds to the charge center of the cyclopropane ion, $C_3H_6^+$. Bond-lengths are in angstroms, a hundred-millionth of a centimeter.

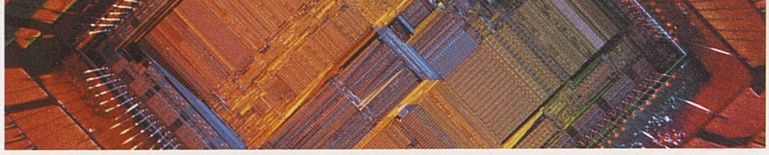

A Clustered Approach to Semiconductors

Hartree-Fock Studies of Anisotropic Semiconductor Clusters

Walter C. Ermler, Maria M. Marino & Makoto Sawamura, Stevens Institute of Technology
Claude J. Sandroff, Bell Communications Research, Inc.

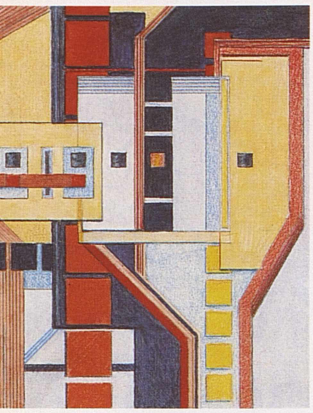

Converging on New Technology

From the onset of the National Science Foundation supercomputing program, one of its objectives has been to open channels between university-centered basic research and the creation of technology, traditionally a concern of industry and business. An area where the interests of pure and applied science tend to converge is the study of semiconducting materials. As the electronics industry searches for new ways to pack more circuitry in less space, it increasingly enters realms of knowledge where theory is a necessary complement to experimental work. In a project at the Pittsburgh Supercomputing Center, supercomputing has served as the medium for joining these practical and theoretical interests.

Walter Ermler leads a group of quantum chemists at Stevens Institute of Technology who have brought their theoretical work in heavy metal clusters into collaboration with a team of experimentalists at Bellcore led by Claude Sandroff. Sandroff has investigated relationships between electronic properties and cluster size for small clusters of semiconducting materials. To date, the cooperative effort has centered on lead-iodide (PbI_2). Ermler's group has performed extensive quantum calculations on several PbI_2 clusters modeled from Sandroff's work.

Clusters and Quantum Size-Effects

"Sub-systems of atoms based on the structure of bulk metal—clusters—have considerable theoretical interest," says Ermler, "because they represent an intermediate regime between solid-state and the atom, and neither solid-state nor molecular formulations adequately describe them." At what point, as clusters get bigger, do their properties become those of the bulk crystal? Besides such fundamental theoretical questions, cluster studies bear on practical, commercial interest in creating new semiconducting materials with finite-size effects. "In semiconductor research, for example, it has become common to fabricate structures with dimensions smaller than carrier wavelength as a way of studying quantum size-effects," says Sandroff. "Exploiting effects such as tunneling and resonance could greatly reduce the size and relative cost of microcircuitry as well as increase their speed and efficiency."

In his experimental work, Sandroff uses transmission electron microscopy, scanning tunneling microscopy and optical absorption spectroscopy to analyze small clusters prepared as a colloidal suspension. Results for PbI_2 indicate that the colloidal clusters retain bulk symmetry but with reduced interatomic distances—a result that suggests the possibility of increased semiconductor speed. The optical studies show "blue-shifts"—greater energy differences—in the band-gaps relative to the bulk. This means that the electrons need more energy to jump from one state to another, increasing the significance of quantum size-effects.

The Computational Approach

Ermler's group used a restricted Hartree-Fock approach to compute optimized geometries and electronic spectra for several PbI_2 cluster geometries. Even within the Hartree-Fock self-consistent field approximation of electron interactions, heavy metal cluster calculations present a "formidable task," says Ermler. "One cluster required 336 basis functions to describe 128 valence electrons." Ermler relied on relativistic core potentials and point group symmetry to keep the computing within reasonable limits. Core potentials, explains Maria Marino of Ermler's group, derive from the frozen-core approximation: "You assume that the innermost or core electrons play a negligible role in bonding compared to valence electrons."

Point group symmetry relies on the symmetry of a cluster's internuclear geometry. "The time required to evaluate integrals from Schrödinger's equation has a fourth-power dependence on the number of basis functions," explains Makoto Sawamura, another of Ermler's co-workers. "Basis-function integrals related by symmetry contribute equally to orbitals with symmetric integrands. This simplifies the mathematics by allowing us to compute only the unique integrals." Because the PbI_2 clusters are highly symmetrical, symmetry effects reduce the calculations appreciably and make it feasible to treat small and large clusters at the same level of theory.

Results and Future Supercomputing

Ermler's numerical study examined two different clusters derived from the PbI_2 bulk

Makoto Sawamura, Claude Sandroff, Walter Ermler and Maria Marino.
"The two research groups continually exchange information on what we're doing," says Ermler, "and the exchange influences what we do next. This type of collaboration would have been quite difficult prior to supercomputers. We couldn't do the calculations quickly enough for the results to be timely."

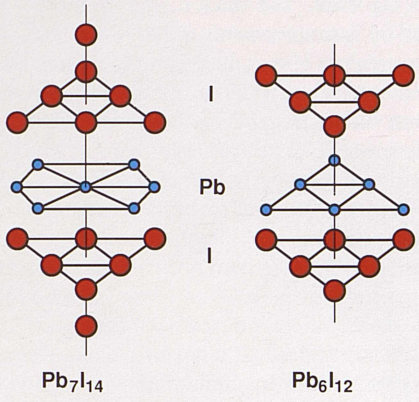

Two Lead-Iodide Clusters

crystal—Pb_7I_{14} and Pb_6I_{12}—over a range of geometries beginning with the bulk crystal and including "shrunken" geometries modeled on Sandroff's experiments down to 30% of the bulk size. For Pb_7I_{14}, quantum calculations indicate that bulk crystal symmetry is preserved with reduced interatomic distances. The most stable geometry occurs at 20% contraction in interlayer separation relative to the bulk with a corresponding lateral expansion of 10%. The most stable form of Pb_6I_{12} corresponds to the bulk in lateral interatomic distances with a 15% interlayer expansion. For both clusters, Ermler found that a 60% vertical elongation results in as much as a 10% lateral contraction in interatomic distances relative to bulk PbI_2.

Ermler's group also did calculations on a three-layered version of Pb_7I_{14}, Pb_7I_{12}—which lacks the "capping" iodine atoms of the Pb_7I_{14} structure. Results for this cluster showed three electron configurations with an energy range from 2.7 to 3.9 electron-volts (eV), considerably "toward the blue" relative to the PbI_2 bulk band-gap of 2.5 eV. Ermler notes the significance of this result, which supports Sandroff's finding: "This behavior appears to exist only for three-layered clusters, which correspond to the thickness of a unit-cell of PbI_2."

Future work between the two groups will focus on bismuth-iodide (BiI_3) as well as other configurations of PbI_2. Ermler also plans to investigate the nature of the bonding between the colloidal clusters and the graphite surface to which Sandroff transfers them, which should shed further light on Sandroff's work. "The close interaction between the group of theorists and experimentalists in designing meaningful calculations has been a crucial component of this effort," says Ermler. Because of the size of these systems, supercomputing plays a major role. Calculations on PbI_2 clusters took about six hours of CRAY X-MP processing time per cluster. "The same analysis on a VAX 11/780 would take over 500 hours. Without the supercomputer, this study isn't feasible."

References:

Walter C. Ermler, Richard B. Ross, Makoto Sawamura & Claude J. Sandroff, "Ab Initio Calculations and Experimental Studies of Anisotropic Semiconductor Clusters," Proceedings of Supercomputing Symposium '88, University of Alberta, Edmonton, Canada, June 1988.

R. N. Nottenburg, C. J. Sandroff, D. A. Humphrey, T. H. Hollenbeck & R. Bhat, *Applied Phys. Lett.* **52**, 218 (1988).

This research is supported by NSF grant CHE-8712315, a grant from the Air Force Office for Scientific Research and Bell Communications Research, Inc.

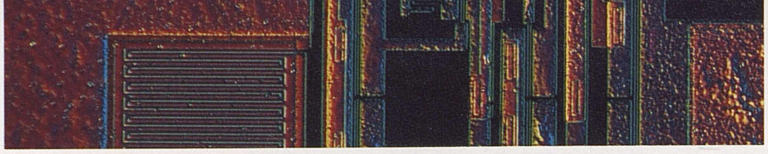

Dissipation in Small Chips

Simulating Charge Transport in a Gallium Arsenide Transistor

M. C. Yalabik, G. Neofotistos, K. Diff, H. Guo & J. D. Gunton
Center for Advanced Computational Science, Temple University

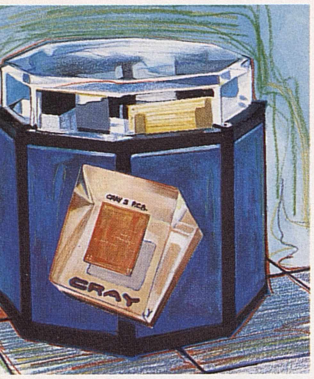

The CRAY-3: a prototype model and its gallium arsenide circuit board.

The Smaller, Faster Revolution

"Speak softly and carry a small, fast chip." It's not far-fetched to think of sophisticated electronics technology as a 1980s parallel to gunboat diplomacy's big stick. The stock market tremors induced by announcements of new Japanese supercomputers hint at the degree to which not only national security, but also economic strength and even such intangibles as collective self-confidence have sensitive ties to American superiority in high-performance computing. The incredible technological revolution fueling it has been large-scale circuit integration. In 1959, a semi-conducting chip held one transistor. Today, over a million fit on one chip—with a billion foreseeable, and circuits are 10,000 times faster.

The steady reduction in device size has opened new questions in semiconducting research. Standard methods of modeling current flow break down when the devices are very small. "Modern methods of modeling carrier transport are based on semi-classical theory—the Boltzmann transport equation," says James Gunton. "These methods have worked well, but they become invalid when devices get down to submicron size" (less than .00004 inches). Gunton, a physics professor and Dean of Arts and Sciences at Lehigh University, formerly directed the Center for Advanced Computational Science at Temple: "These very small devices aren't much farther across than the wavelength of the electrons going through them, which means you have to go back to first principles quantum mechanics."

Though a finite limit to size reduction hasn't yet been identified, its inevitability has spurred research on materials that may be faster than the silicon compounds now used almost exclusively. One material in particular, gallium arsenide (GaAs), with electron mobility six times faster than silicon, holds promise of a fresh push to computing speed. Gunton and M. Cemal Yalabik lead a research group using Pittsburgh's CRAY to explore both these areas: "Our long-term objective," says Yalabik, "is to develop code that will accurately simulate very small semiconductor devices. Our initial project has been to model a GaAs field-effect transistor (FET)."

Modeling Quantum Dissipation

"Laboratories are being very creative," says Gunton, "in building smaller devices." As theoreticians, we're trying to describe electron motion in these devices in order to predict how they'll operate. One of the most difficult fundamental problems for modeling is electron-phonon interaction." An electron in a charged field, as in an FET, moves in short straight-line bursts interrupted by collisions with *phonons*—the term that describes the vibrating ions in the semiconductor material. At each collision, the electron changes direction and may dissipate energy as heat. "Recently a few groups have worked on solving this problem from first principles. The best theoretical approach is the Feynmann path integral, but computing requirements are prohibitive."

Gunton and Yalabik's alternative approach is to solve Schrödinger's equation for single-particle states. "Our main assumption," says Yalabik, "which is also routinely made in semi-classical approaches, is that carrier density is relatively low." In effect, the electrons are far enough apart that electron-electron interactions can be "averaged": "Each electron will be moving in an electrostatic field that results from the electron distribution at that instant in the device." To incorporate scattering and dissipation, Gunton and Yalabik apply perturbation theory to adjust the energy state of the electron at each time step. They also incorporate a random noise term to describe the effect of phonon collisions on the phase of the electron wave function.

The research group sees its approach to electron-phonon interaction as less than ideal. "Our treatment should give good results for systems weakly coupled to the phonon field," says Gunton. "But it's an approximation which can give only partial information for limited time and length scales. The main aim of this work at this stage, rather than detailed simulation, is to demonstrate feasibility of integrating the Schrödinger equation. It takes supercomputing to implement even this limited method. With advances in computing, we hope to use Feynmann path integrals." Gunton, Yalabik and their co-workers see the CRAY-3, currently in development as the first supercomputer to

use gallium arsenide technology, as likely to offer the computing speed necessary to use path integrals for realistic semiconductor devices.

A Fast Gate & Resonant Tunneling

The group applied their method to a GaAs FET with dimensions equivalent to laboratory fabricated submicron-size transistors: 127 nanometers (nm) along the current-flow axis (about .000005 inches). The three surface electrodes, gate, source and drain, had lengths of 26, 24 and 24 nm respectively. The dimension of principal interest is gate length. As the name implies, the gate sits between the source and drain and acts like a switch controlling electron flow across the transistor. For a transistor in a computer circuit, the significant speed characteristic, switching time, amounts to the average time an electron takes to travel the length of the gate.

James D. Gunton

The model used time steps reflecting the potential high-speed of these devices—.4 femto—seconds (fs) (.4 quadrillionth sec.) to simulate a time period corresponding to .01 picosecond (ps) (.01 trillionth sec.) of transistor operation. (A simulation to 1 ps, the ultimate objective, takes about 70 hours of CRAY X-MP processing.) Results show early stage formation of a *depletion layer* —a region of negative charge that blocks conduction—below the gate, indicating a very fast switching time. "We've represented qualitative features of what the experimentalists have seen," says Gunton, "but we haven't yet been able to give a quantitative explanation. The most interesting problem has yet to be solved."

In recent work, the group has studied *resonant tunneling* in alternating thin layers of GaAs and aluminum-GaAs (AlGaAs). "Experimentalists are making sandwiches of GaAs and AlGaAs repeated over and over," says Gunton, "and some interesting quantum effects occur." Through resonant tunneling, electrons at a "resonant voltage" can travel from GaAs across AlGaAs even though its conduction band is at a higher energy level. The group's most successful modeling to date involves these devices, in which resonant tunneling provides the possibility of very fast, precisely controlled switching. Their results—which so far have neglected the electron-phonon interaction—appear to be the first reported first principles simulation of the temporal behavior of these devices.

References:

M. C. Yalabik, G. Neofotistos, K. Diff, H. Guo & J. D. Gunton, "Quantum Mechanical Simulation of Charge Transport in Very Small Semiconductor Structures," *IEEE Transactions on Electron Devices* **36**, 1009 (1989),

G. Neofotistos, H. Guo, K. Diff & J. D. Gunton, "Resonant Tunneling in Double-Barrier Parabolic Well Structures," *IEEE Transactions on Electron Devices* **36**, 745 (1989).

This work was supported by U. S. Office of Naval Research grant N00014-83-K-0382.

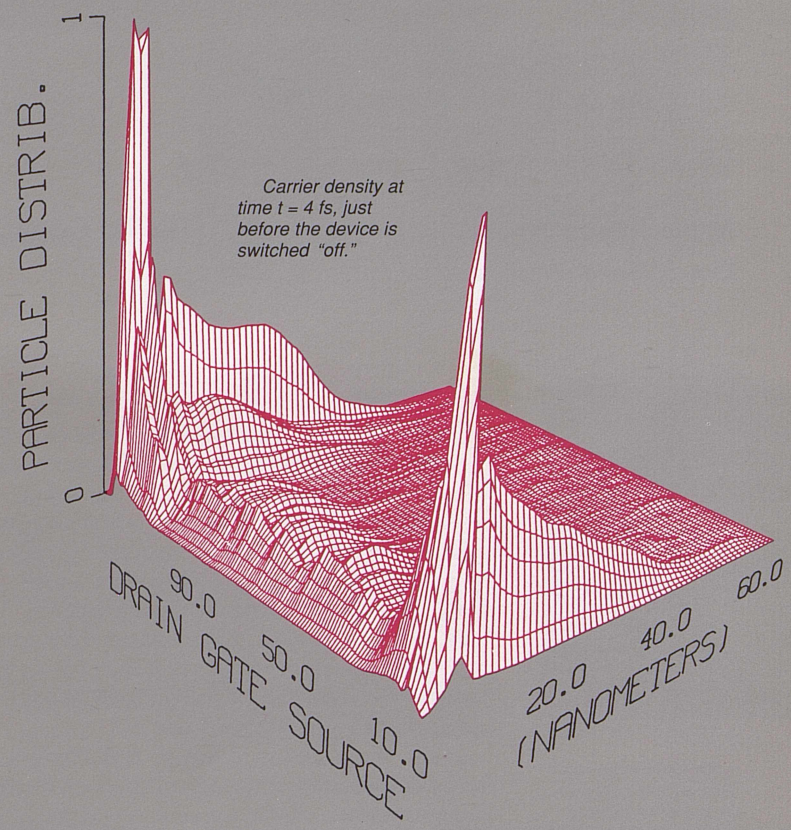

Carrier Density Distribution Before and After the Gate is "Off"

The simulation computed integrals over a 63 x 127 x 16 mesh with 1 nm spacing. It treated the gate as initially "on" and then instantaneously switched "off" after 4 fs.

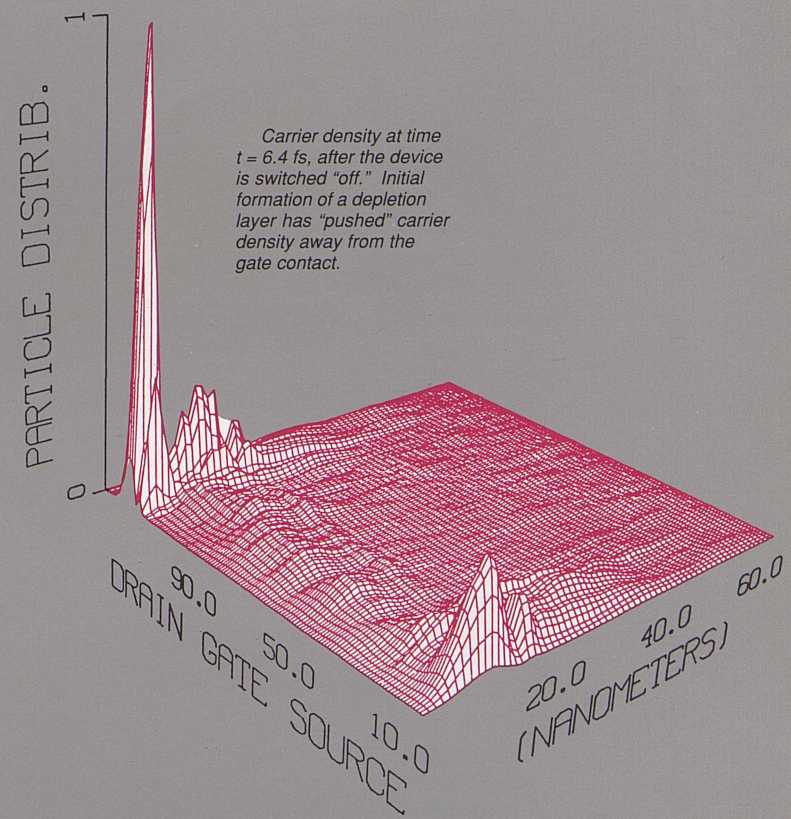

Death of a Hot Young Star

Modeling Gravitational Radiation from Stellar Core Collapse

Edward Seidel, Washington University, St. Louis

We had the sky, up there, all speckled with stars, and we used to lay on our backs and look up at them, and discuss about whether they was made, or only just happened.

Huckleberry Finn, *Mark Twain*

Supernovae

Things happen at their own pace in the cosmos and won't be hurried. It takes a million earth years for noticeable change in galaxy evolution, and even a short-lived star radiates for a million years before it expires. But then about every 600 years per galaxy on the average, a supernova comes along violently disrupting the normal pace. A large, extremely radiant and consequently short-lived star exhausts its nuclear fuel and evolves to the critical instant when it doesn't have enough energy to resist its own gravity.

In less than a second, it's over. A stellar mass ten times greater than the sun collapses on its core of atomic nuclei, which overcompresses and rebounds, sending a shock wave smashing outward against the inward rushing material of the star's outer layer. The resulting blast releases as much energy as the sun radiates in a million years, and for a brief moment the light is as bright as an entire galaxy. Thrown-off remnants of hot gas may be visible to astronomers thousands of years later. The core remains as a *neutron star*—a small, spinning clump of extremely dense matter.

Gravity Detectors & Supercomputing

According to Einstein's general theory of relativity, a supernova could be a strong source of gravitational radiation. "It's one of the most massive objects we know of and it's a very violent process," says Edward Seidel. Working with Eric Myra of SUNY Stony Brook and Thomas Moore of Pomona College, Seidel has undertaken to develop a computer model that can predict gravitational emission from core collapse events—supernovae and black holes. It's a relatively uncharted field—no experimental benchmarks and few precedents in modeling work.

"No one has detected gravitational radiation," explains Seidel. "It's one of the major untested predictions of General Relativity." Two developments suggest this will soon change. Several sensitive detection devices will begin operating over the next five-to-ten years. To use them effectively requires knowing about how much radiation to expect. This is where Seidel and the CRAY come in. "Existing predictions vary wildly—over several orders of magnitude. The theory is the same as it was 70 years ago, but it's so complicated that we haven't been able to solve a realistic enough system. With supercomputing, we can actually solve equations."

Modeling Gravity Waves

Einstein predicted that when objects move in certain ways they emit gravity waves, which travel at the speed of light. "It's a disturbance in the gravitational field. An electron bouncing back and forth emits electromagnetic radiation. Similarly, if you shake a mass you'll get gravitational radiation, but it has to be a more complicated motion involving so-called quadrupole moments." Among other things, this means a star must be non-spherical. "If a star has the most amazing explosion but remains perfectly spherical, you won't get any gravity waves. Basically, the more massive a star, the more non-spherical and the faster it moves, the more radiation it will emit."

Researchers haven't been especially interested in gravity, and most stellar collapse models have been spherical. Seidel's model draws from this work but says *A star isn't really spherical—almost maybe, but there's going to be an instability or rotation that makes it slightly non-spherical.* He incorporates these slight deviations as "perturbations" into a spherical background code. This lets him use linearized Einstein equations instead of the much more difficult non-linear, full-scale equations. The spherical code has a known solution, which means Seidel can get concrete results.

Seidel sees other advantages to the linearized equations. "The perturbation method may not apply to wildly non-spherical systems, but it should be far more efficient and accurate than non-linear codes for slightly non-spherical sources." The background, a sophisticated hydrodynamics code developed by Myra, has a detailed *equation of state*—an expression of how gas behaves. Work with the non-linear equations uses less sophisticated hydrodynamics. Myra also modeled neutrino emissions, which mark the critical point when the star lacks energy to resist the inward pull of the core. "We tried to include as much physics as the supernova community can provide."

Ringing Black Holes & Soft Explosions

Initially, Seidel ran several "pressureless collapse" simulations—black holes—and compared

them to detailed modeling done by Price, Cunningham and Moncrief in the late 70s. This is the theoretically predicted case in which gravity takes over and pulls the star to a single point of mass. Results closely reproduced the earlier work. "It's as if the black hole starts ringing like a bell, giving off this high amplitude radiation, stronger than from a core bounce." Detailed knowledge of this effect would be helpful in interpreting gravity detection data, and Seidel plans to study it further in the future.

In a rather unexpected result, Seidel found that "stiff" equations of state, which produce less violent explosions than "soft" equations, emit *more* gravitational radiation. If this picture is right, says Seidel, stars that don't explode because of stiffness—i.e., pressure builds sharply and the core bounces at lower density—are in some sense reacting more violently, emitting more energy than those that explode. This effect points to another potentially valuable path for future work.

References:

Edward Seidel, Eric S. Myra & Thomas Moore, "Gravitational Radiation from Type II Supernovae: The Effect of the High-Density Equation of State," *Phys Rev D* **38**, 2349 (1988).

Edward Seidel & Thomas Moore, "Gravitational Radiation from Perturbations of Stellar Core Collapse Models," in *Frontiers in Numerical Relativity* (N.Y: Cambridge U.P., 1989).

This research was supported by NSF grants PHY85-13953 and PHY85-03072 and NASA grant NAGW-122.

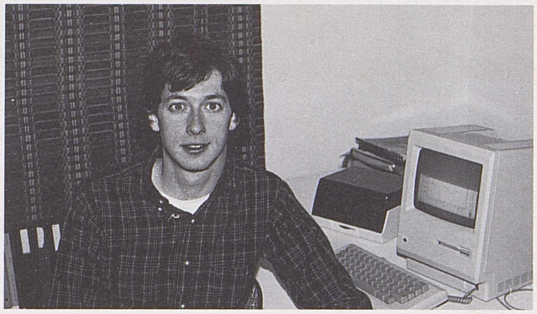

Edward Seidel

Hydrodynamic Evolution

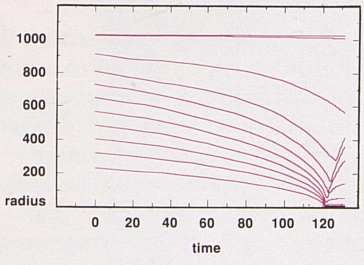

Gravitational Waveform

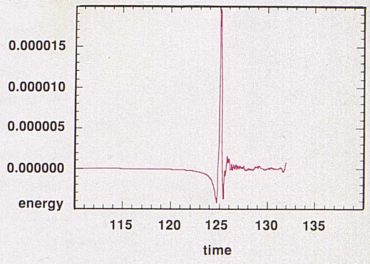

This typical hydrodynamic collapse model represents gas movement on a scale of kilometers versus time in milliseconds (ms). The collapse and core bounce at approximately 125 ms correspond to the gravitational wave peak. Emitted gravitational wave energy is shown as a decimal fraction of core mass.

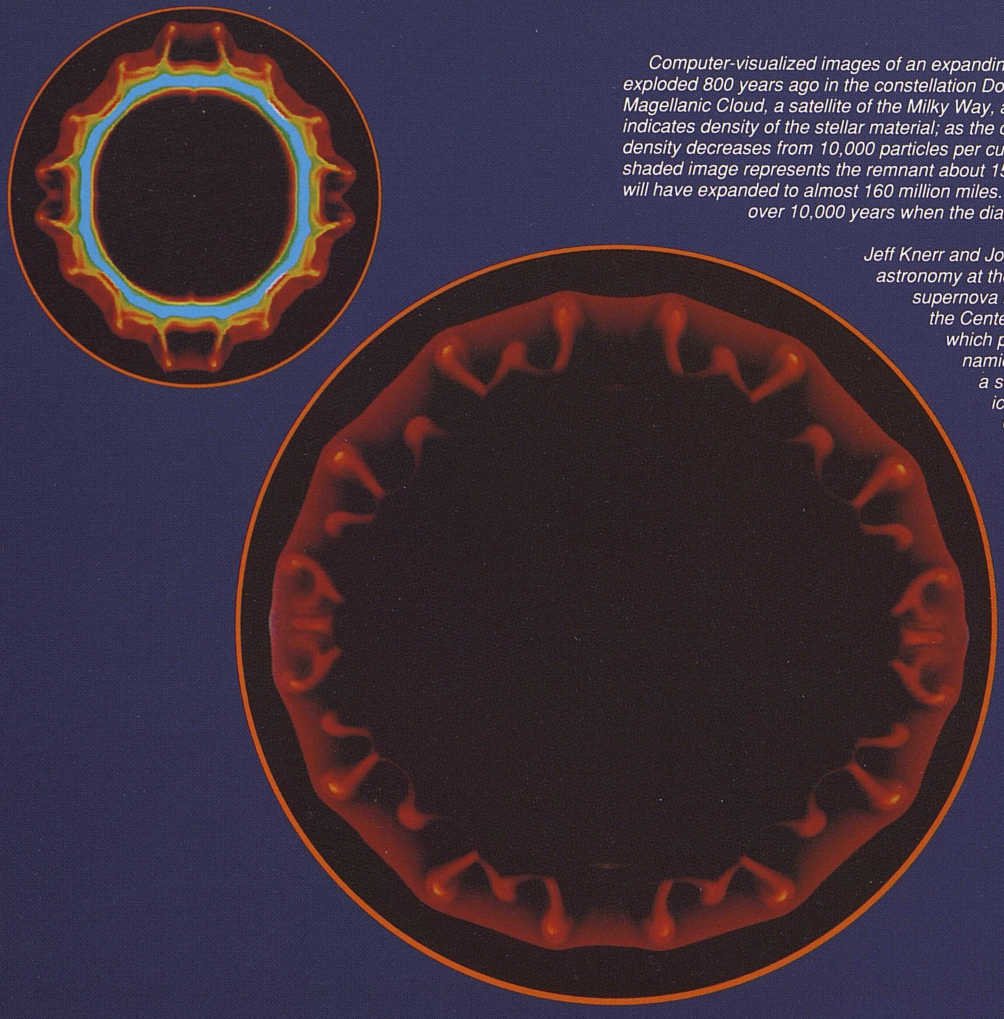

Computer-visualized images of an expanding, radiatively cooling remnant of a supernova that exploded 800 years ago in the constellation Dorado, the swordfish. Dorado is located in the Large Magellanic Cloud, a satellite of the Milky Way, and is visible in the Southern Hemisphere. Color indicates density of the stellar material; as the color shifts from white/blue through red to black, density decreases from 10,000 particles per cubic centimeter to empty space. The rainbow-shaded image represents the remnant about 1500 years after the supernova when its diameter will have expanded to almost 160 million miles. The all-red image shows the remnant at slightly over 10,000 years when the diameter has reached 620 million miles.

Jeff Knerr and Jon Morse, graduate students in physics and astronomy at the University of North Carolina, modeled this supernova remnant on Pittsburgh's CRAY X-MP. They used the Center's graphics capability to produce the visualization, which predicts gas fragmenting in accord with hydrodynamic theory. Knerr and Morse did this work as part of a seminar, "Hydrodynamic Processes in Astrophysics," led by Drs. Allen Schiano and Wayne Christiansen. The Pittsburgh Supercomputing Center encourages the use of its computing resources in courses like this, which offer students an opportunity to become acquainted with supercomputing techniques.

Heavyduty Lightweighting at ALCOA

Aluminum Can Design with Finite Element Methods

A. B. Trageser & R. B. Dick, ALCOA Laboratories

The Beverage Can Story

Remember beer and coke cans with seams? For Andy Trageser and Bob Dick, they tell an interesting story. In the 60s, nearly all beverage cans were a seamed wraparound of tin-plated steel. Then the aluminum industry developed a way to make seamless cans—drawing them from a flat sheet with a separate lid clamped on after filling. "That started us into lightweighting," says Trageser, who like Dick practices the art and science of engineering at ALCOA, "to save our can manufacturing customers' costs and to compete with steel." Last year, 95 percent of 88 billion beverage cans made in the U.S.—one per person per day—were aluminum.

The story goes on. Aluminum costs more than steel, and the price has been rising. Steel "minimills" now have continuous casting processes that make sheet steel thin enough to form seamless cans. And there is competition from other materials as well. "We have to find ways to make cans lighter and lighter to keep fending off polymers, steel and glass. Lighter cans mean lower prices to the consumer, who's then more likely to buy cans off the grocery shelf instead of two-liter bottles or glass."

From the Lab to the Computer

Lightweighting means designing a can to use the thinnest aluminum that will still meet specs for strength and appearance. Since 1979, it has been a service that ALCOA provides to can manufacturers who use its aluminum sheet. By the early 80s, lightweighting had reduced sheet thickness from .015 inches to around .013. It began as a laboratory process involving considerable trial and error. "To come up with a design that looked encouraging," says Trageser, "one you could present to your customer and go to market, would take six months to a year. We wanted to cut that down and decided to create the first computer model in 1984." The first models used on-site computing (VAX 8650 and 8800) to simulate bottom and end deformations.

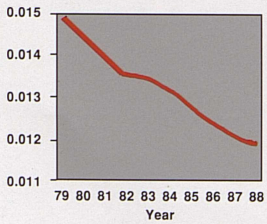

Progressive Lightweighting
In the 1970s the aluminum in beverage cans was nearly as thick as aluminum gutters, .015 inches. Lightweighting progress leveled-off in the early 80s, then resumed in 1984 due in part to computer modeling.

Because a can is symmetrical around its vertical axis, forces that push out the bottom or buckle the lid—internal pressures and axial loads—can be modeled in two-dimensions. "The basic assumption is that you can take one cross-section and revolve it 360° and you'll get the shape of the bulge," explains Dick. Normal internal pressure is about 60 lbs/in.^{2}, but it goes as high as 90, during beer pasteurization for instance, and the model for internal pressure determines *dome reversal pressure*—when the bottom bulges. The results matched experiment within 5% and showed that computers could be a valuable aid in can design. Nevertheless, Dick and Trageser soon realized that computational power limited what they could accomplish.

Modeling on the Supercomputer

In May 1987, ALCOA became the Pittsburgh Supercomputing Center's first industrial affiliate. "Our objective is that the creativity, innovation and imagination of our scientists shouldn't be limited by perceived limits on computational capability," says ALCOA's vice-president for research and development Peter Bridenbaugh. For Dick, it was what the doctor ordered. "The barriers we're running into in can design now can't be handled with easy two-dimensional models. Cans impact cans in the production process. They didn't dent before because they were so thick. Now they're getting thin enough that dents make a difference, and we're faced with these computationally intense three-dimensional, dynamic impact models."

Using the CRAY X-MP, Trageser and Dick have instituted three kinds of three-dimensional impact modeling. First, "dynamic snap-through" from internal pressure: unlike the 2-D model, this 3-D version includes a velocity component and shows change in shape and pressure over time (Fig. 1). They've also modeled the *drop test* (Fig. 2). "If cans are thrown off a truck, the impact may cause a bulge in the bottom," explains Trageser. "Though there's nothing wrong with the product, a consumer doesn't want a bulged or dented can."

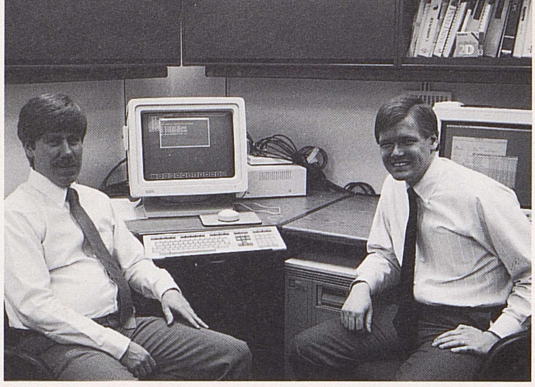

Andy Trageser and Bob Dick

Trageser and Dick have also begun to do the complex non-symmetric modeling required for dent analysis (Fig. 3). As cans are manufactured at rates approaching 2000 a minute, they move along a rigid guide rail and bang into it going around corners. "We're trying to determine the maximum velocity that those cans can travel before dents start to occur," says Dick. "This could lead to manufacturing changes like modifying the location or design of the rails." The model tracks the can sidewall as it dents on impact and then springs back to a residual shape.

Benefits of Modeling on the CRAY

Trageser estimates that to develop a new can design in the laboratory costs about $100,000 compared to about $2000 with computer modeling. "Typically, we'll use modeling to try out ideas; when we have something close, we go into experimental work—actually make tools and build cans. As far as major changes in shape, we can do that fairly accurately with modeling." Dick notes that supercomputing has also significantly reduced computing time as compared to the on-site VAX. "Those 50-hour jobs that used to take a week run in less than an hour. The models we're running now require a lot of computing time. For us, the CRAY is the way to go."

ALCOA has also begun using a Monte Carlo method to generate initial designs. "The CRAY generates random can shapes within a family or range," says Trageser, "and we look through them to see if any look promising." This method has shown itself to be a good marketing tool; ALCOA can present customers with a range of designs meeting the basic criteria rather than only one. In the future, Trageser and Dick see a need for more complex dynamic impact models—repeated drops, guide rail placement, and the effect of material strength, thickness and profile design on dents. Says Dick, "There's plenty to do, which is good because we're not ready to retire. It's all feasible because of the turn-around we get with supercomputing capability."

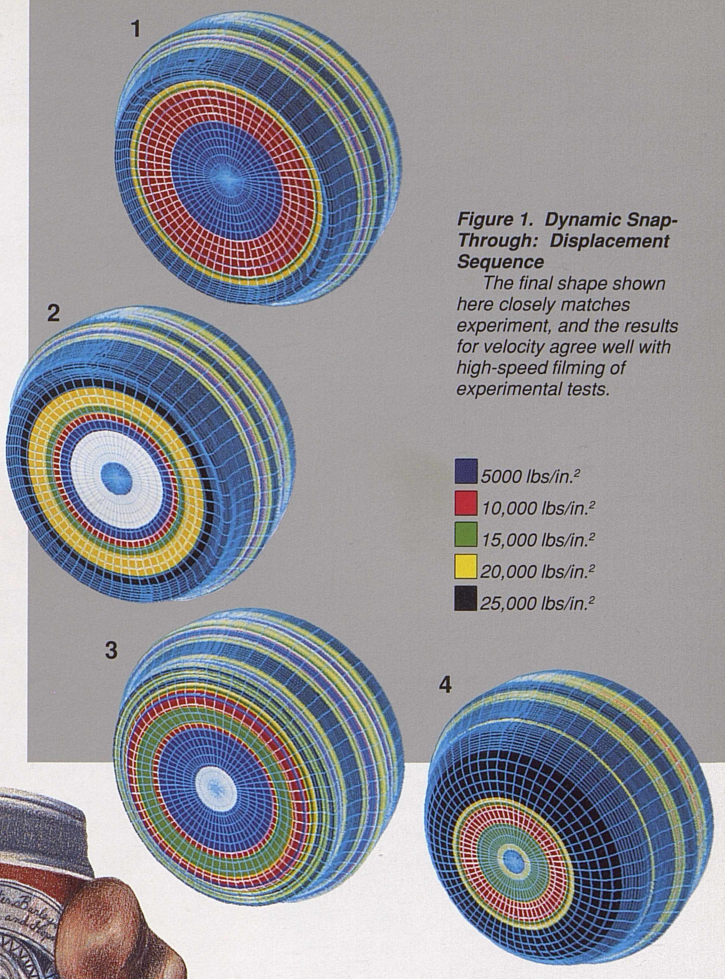

Figure 1. Dynamic Snap-Through: Displacement Sequence

The final shape shown here closely matches experiment, and the results for velocity agree well with high-speed filming of experimental tests.

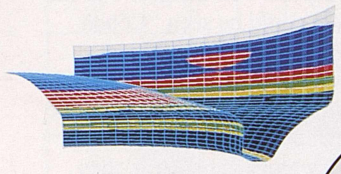

Figure 2: Drop Test: Partial Dome Reversal Sequence

Impact stresses on can bottom for a vertical drop of two feet. Colors as in Figure 1.

Figure 3: Dent Analysis

This sequence models an empty can colliding with a .5 in. guide-rail at 10 ft./sec. The first image (a) represents the finite-element mesh and the guard-rail. The color images show stress in lbs/in.² related to color: dark blue (6000 lbs/in.²), yellow (12,000) and green (20,000). Maximum dent of 0.35 in. occurred 2.1 millisec. after impact (b). The final residual dent of 0.2 in. occurs at 4.0 millisec. (c).

Reference:

A. B. Trageser & R. E. Dick, "Aluminum Can Design using Finite Element Methods," presented at the Society of Manufacturing Engineering Can Manufacturing Technology Symposium, Schaumburg, IL, Sept. 14-16, 1988.

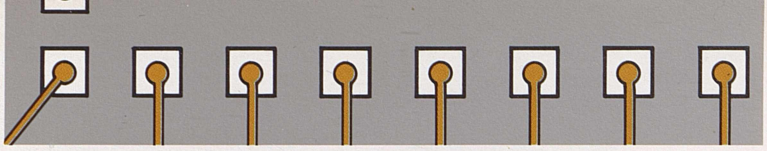

All that Blisters is not Gold

The Ball Formation Process in Aluminum Wire

Ira M. Cohen & P. S. Ayyaswamy, University of Pennsylvania

Wire Bonding of Semiconductor Chips

Gold in 1849 provoked a mad rush to California. In 1989 gold is still in California—at the terminals of silicon chips. Fine wires of nearly pure gold make the electrical connection between each finished chip and the "lead frame" on which it's mounted (Fig. 1). In a market as competitive as semiconductors, why use one of the most precious materials on earth? Because nothing else works. The best way to bond the wire involves discharging a spark across a gap at the wire's end, melting and forming it into a ball which is then pressed onto a "bond pad" at the chip's edge. At this late stage of production, only a very low failure rate is acceptable. Good bonds depend on well-formed balls, and gold—a good conductor of heat and electricity—forms defect-free balls more than 99.9% of the time.

The industry has tried aluminum, notes Ira Cohen, and remains interested because the price differential is enormous—over $400 an ounce versus a dollar or two a pound. "They couldn't make it work," says Cohen, who brings 25 years experience in "glows and arcs" to his project. "The reliability was around 90%, and that was intolerable. But industry efforts were cut and try." No systematic analysis had been done, and the situation presented Cohen and his colleague P. S. Ayyaswamy, who specializes in heat transfer, with if not a golden at least an aluminum opportunity.

Slow Motion Ball Formation

Cohen and Ayyaswamy set out to learn what caused aluminum to fail. The process is extremely fast—two milliseconds or less from start to finish, so they built a laboratory scale model that is 100 times larger, slowing the process enough for high-speed photography to capture still images. For the first time, ball formation could be seen. Cohen and Ayyaswamy were able to identify several problems with aluminum wire: it formed asymmetrical balls, it "ballooned" and it "necked"—the wire at the ball-wire junction became thinner and weaker. Armed with new physical insight, they formulated mathematical and numerical approaches to model the process computationally, ultimately aiming to specify the parameters of good ball formation.

International Outlook

Along with more work on aluminum, Cohen and Ayyaswamy have begun running simulations on copper, which Ayyaswamy sees as holding more potential in the long run than aluminum. "Copper balls are essentially void free if they're done properly. The Japanese have virtually gone into copper. I wrote four letters suggesting that we collaborate. They simply said they appreciate what we're doing, but at this moment it's not technologically the best thing for them to work with us."

Simulating Complex Heat Dynamics

Simulating ball formation poses some difficult numerical problems. Melting and re-solidification of the wire involve phase transitions from solid to liquid to vapor and back from liquid to solid, with moving boundaries between the phases. Unlike water, solid aluminum is denser than liquid, so the balls expand as they melt and leave voids as they re-solidify (Fig. 2). The model must precisely track the "solidification front," as Ayyaswamy terms it: "Whether distributed porosity or one hole, we have to know, and tracking the solidification front with time is a very complicated numerical problem."

For Ayyaswamy the most vexing numerical problem relates to the interface of the ball and unmelted rod, the point where necking can occur. "Think of it as a long rod, a pendulum at the end of which you have a bob. Where the bob joins the shaft, you have—strictly speaking—temperature continuity. But it's a very sharp turn; the wire's at one temperature, and the ball is molten. The heat fluxes become enormously high because you have a finite temperature gradient over a very small length. Computationally it's a nightmare."

To overcome these problems, Ayyaswamy, Cohen and Ph.D. students Lin-Jie Huang and K. Ramakrishna adapted a body-fitted coordinate method devised by J. F. Thompson. A fine net-like grid, mathematically fitted to the ball's surface, tracks the movements and distortions of ball formation. At each time step, an algorithm maps the curved coordinates onto a uniform rectangular grid. "We're able to get all the temperature profiles, details of flow structure, everything by doing this," says Ayyaswamy.

Supercomputing

Ayyaswamy believes that CRAY number-crunching power has given the project its viability. "Without the supercomputer," he continues, "this adaptive grid and these phase change problems would be virtually intractable." For typical conditions, grid sizes are 201 x 101 with 1000 or more time steps. The solidification phase alone requires about 35 minutes to run (compared to 25 hours on the University of Pennsylvania's Harris

HCX-9). "Resolution of these details—this is where we realize the benefit of the supercomputer. We can change our grid sizes and experiment to find efficient algorithms. This is inconceivable on standard machines. It would have taken forever."

The Keys to Flawless Ball Formation

Cohen and Ayyaswamy's ball formation model is to their knowledge the only one to simulate the entire process. The results closely replicate their scaled-up experiment and reproduce some defects observed experimentally and in industrial use of aluminum. Says Cohen, "I think we understand why defects arise in aluminum balls and how to prevent them." Necking and ballooning both result from overheating, though the phenomena differ: necking relates to surface tension breakdown from overmelting and ballooning to vaporization in the ball interior. Asymmetrical balls essentially disappear when heating is symmetrical. With alloyed aluminum, the balls solidify gradually over a temperature range rather than at a precise point. As a result, proper alloying eliminates large shrinkage voids.

What we've found that I don't think anyone has found before," says Cohen, "is what causes the defects in the ball. And that tells you what you have to do to eliminate them. But how far can you deviate from the parameters that give perfect balls and still get satisfactory balls? Is the precision required for non-precious metals practical? Or must you be so precise that random variations from shot-to-shot will prevent good balls?" Cohen and Ayyaswamy expect to answer these questions in future work at Pittsburgh, where access to the computing power of the CRAY Y-MP/832 will allow finer grid resolutions and testing of a large number of cases.

References:

I. M. Cohen and P. S. Ayyaswamy, "Ball Formation Processes in Aluminum Bonding Wire," *Solid State Technology* **28**, 89 (1985).

L. J. Huang, P. S. Ayyaswamy and I. M. Cohen, "Numerical Simulation of the Ball Formation Process in Wire Bonding Apparatus," Proc. 15th Ann. Conf., Production Research and Technology, U. of Calif. Berkeley, Jan. 9 - 13, 1989.

The following is a partial list of contributors: The National Science Foundation: DMC 8513128, DMC 8709537, plus REU Supplements; Commonwealth of Pennsylvania, Ben Franklin Partnership Program grants 06.500 NU, 07.510 RU; Research Fund of the University of Pennsylvania; International Society for Hybrid Microelectronics—Educational Foundation; Kulicke and Soffa Industries, Inc.

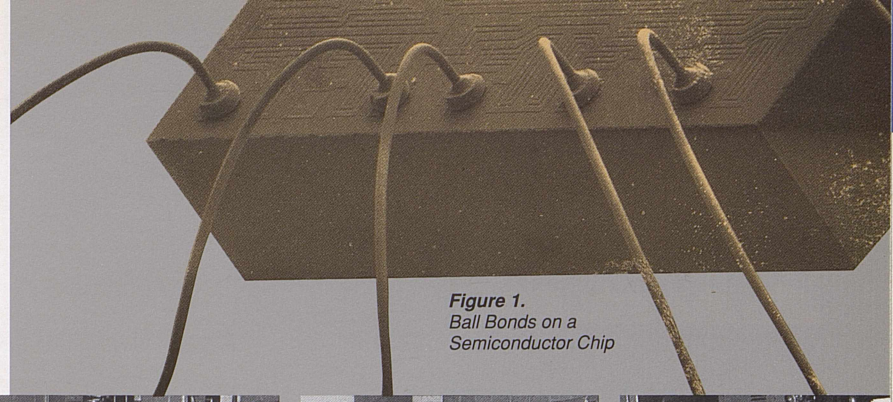

Figure 1.
Ball Bonds on a Semiconductor Chip

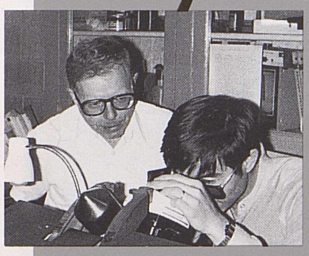

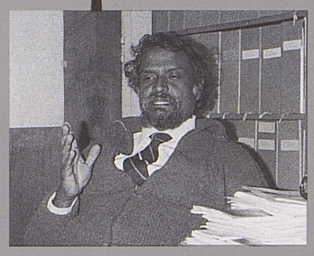

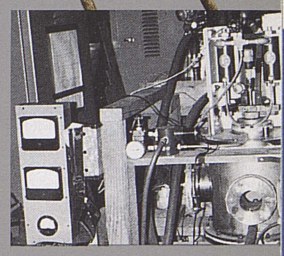

Ira Cohen looks on as Ph.D. candidate Lin-Jie Huang peers into a microscope. P.S. Ayyaswamy explains a principle of heat transfer. Their 100 times scaled-up ball-bonding apparatus awaits the next experiment.

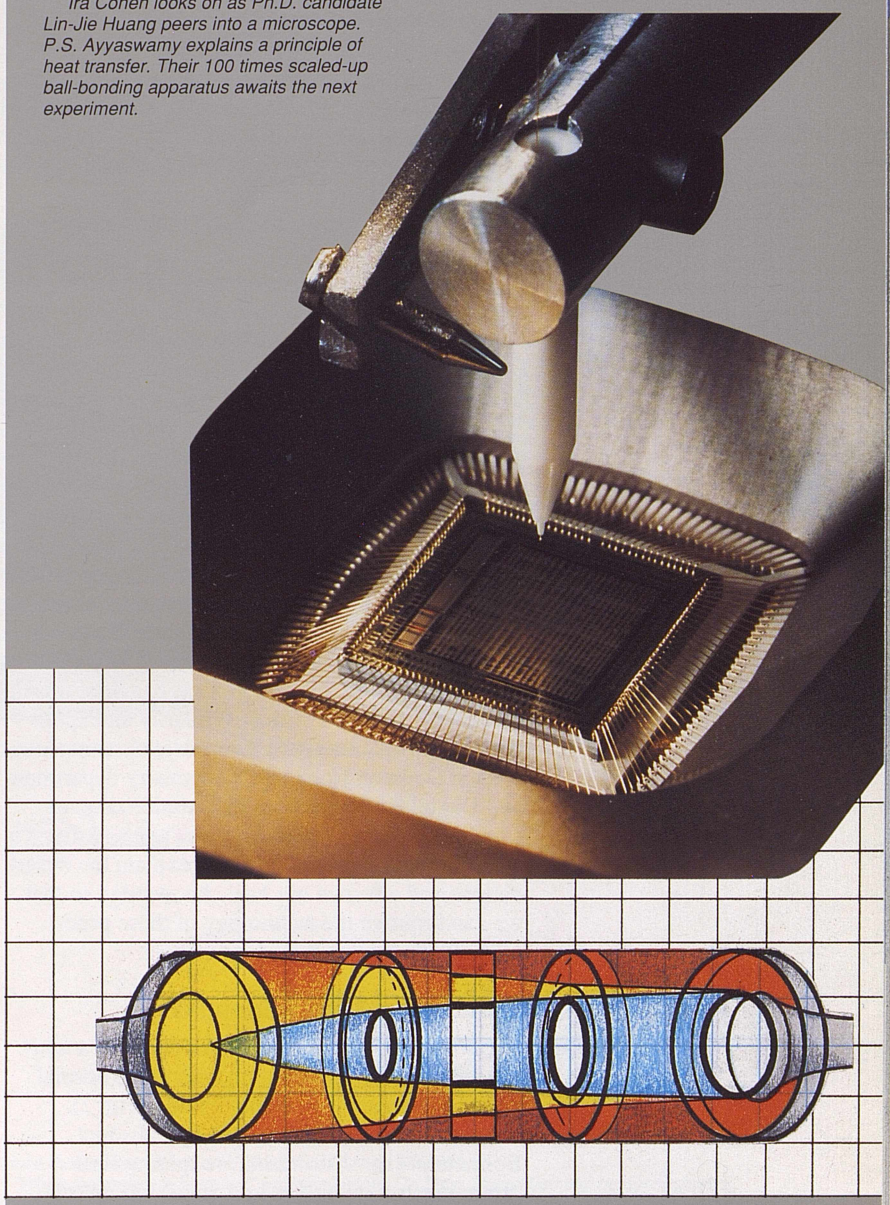

Figure 2. Solidification History for a Pure Aluminum Ball

When the arc is removed, the molten ball (2240° F) begins to solidify from the outside inward. Since solid aluminum is more dense than liquid, a vacuum forms in the ball's core. As the solid-liquid boundary moves inward, the vacuum-liquid boundary moves outward. When they meet, the ball is solid with a vacuum core.

Reactions at the State Border

Probing the Transition State of a Surface-Catalyzed Reaction

Daniel W. J. Kwong, N. De Leon & Gary L. Haller, Yale University

Catalysis at the Gas-Solid Interface

Does it work? Basically, that's the test of a heterogeneous catalytic reaction. Substances in different phases—a gas and a solid commonly—interact with the consequence that a molecular reaction in one of the phases, usually the gas, goes faster. Industry uses these processes extensively. In petroleum refining, platinum helps transform straight-chain hydrocarbons, which are low octane, to high-octane branch-chain hydrocarbons. Catalytic converters work by this process, and similar gas-solid interactions play a role in producing fertilizers, polymers, cosmetics, pharmaceutical chemicals and semiconductor chips.

Using molecular beam methods, chemical engineer Gary Haller has performed detailed experimental studies of carbon monoxide oxidation on a platinum surface, a paradigmatic catalytic conversion reaction. Oxygen and carbon monoxide *adsorb*—adhere chemically to the platinum surface, and through a sequence of several stages, the carbon monoxide and oxygen form carbon dioxide, which then *desorbs*—releases from the platinum as a free molecule (Fig. 1).

"The basic tenet is that the metal surface brings the gas molecules into a configuration favorable to reaction. Yet we know relatively little fundamentally about how this happens. Finding a catalyst for a particular production process has been mostly a matter of trial and error." Primarily an experimentalist, Haller enlisted theoretical chemist Nelson De Leon from Yale's chemistry department and Daniel Kwong, whose dissertation derives from the project. "The question," says Haller, "is what theoretical model will best explain the observations and advance our basic knowledge so that we can improve the technology of these processes?"

Is Statistical Theory Enough?

Haller's molecular beam work revealed a non-statistical energy distribution among the normal vibrational modes of carbon dioxide (Fig. 2). "Each of the three modes is fit by a distinct Boltzmann-like distribution," explains Haller. "Boltzmann-like" refers to a statistical distribution that can be characterized by a single temperature, and in a purely statistical distribution, all three modes would have the same rather than three different vibrational temperatures. "It's important to know the dynamics of the process," says De Leon. "If for some reason it's not statistical, then something's happening that has to be understood before you can develop an accurate theory." With this question motivating the work, Haller, Kwong and De Leon used Pittsburgh's CRAY X-MP for a molecular dynamical study.

Stochastic Classical Trajectories

Gas-surface interactions present unique problems that place them among the most difficult chemical processes to treat numerically. The calculation must account for the in-principle infinite extent of the solid-state lattice. The gas molecules interact strongly with several nearby surface atoms, but they also interact with a vast extent of others. To reduce this problem to solvable dimensions, Haller's group used the Generalized Langevin Equation approach (GLE). "The gas atoms directly couple to a few surface atoms designated as *primary*, in this case four." GLE treats the effects of the remaining solid atoms as a statistically-averaged field, with friction and random forces. For interactions with the primary atoms, GLE does a classical integration over the trajectory of the reaction, which is further averaged to give a final energy distribution.

A second problem in this case was how to model the full sequence of interactions. "It's still not feasible," notes Haller, "to construct an accurate *ab initio* potential energy surface for all the interaction forces." Given that a fully quantum approach is not numerically solvable, Haller looked for a workable approximation and deduced that events before carbon dioxide forms have slight influence on the energy distribution compared to the transition state or activated complex. This conclusion derives in part from the short surface residence of the activated complex, about .2 picosecond, two ten-billionth of a second. If it stays on the surface a long time, explains Kwong, it will begin to equilibrate with the platinum and "lose memory" of its formation energy.

Results and New Findings

Having concluded that the transition state is key to the reaction, the group's objective became to

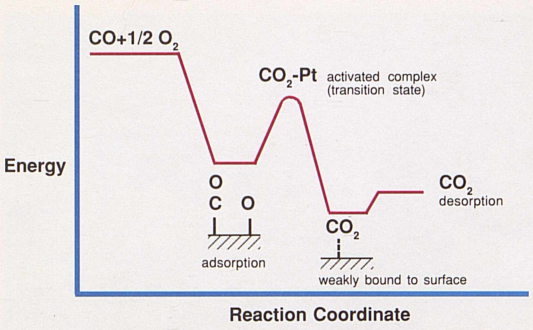

Figure 1. Potential Energy Diagram

find transition state initial conditions that would produce the experimental results. Their simulations systematically varied temperatures, energy distributions and configurations. The best agreement with experiment occurred with a "thermally hot" transition state in a vertically bent geometry (Fig. 3). The internal energies agreed closely with experiment, giving a non-statistical distribution—a finding supported by the identifiable transition state. "It's possible to understand certain gas phase experimental data from a purely dynamical perspective," says De Leon, referring to the non-statistical distribution. "And it's also possible that the dynamics can have a fairly large statistical component, such as within each vibrational mode."

Recent computations have revealed that the desorbing carbon dioxide splits into two groups. The majority, 80-90%, has a very short residence time. "The molecules just immediately pop off the surface," says Kwong, "and that takes about .4 picoseconds." The second group, "the trapping-desorption channel," tends to desorb 20 times more slowly—about 8 picoseconds. These findings raise interesting questions for ongoing work: "Why do two groups of carbon dioxide come from one reaction, and what conditions determine that a molecule will go through the trapping instead of the direct desorption channel? Nobody really has a good answer."

Supercomputing

Each trajectory of the GLE computation solved approximately 30 coupled first-order differential equations, and each simulation required 1000 to 1500 trajectories to get statistically valid results. Haller estimates that a MicroVax II would use 120 to 150 hours of processing for one simulation over a sufficiently broad range of trajectories. On the CRAY, it took about 50 minutes. Says De Leon, "The main thing is that the supercomputer made these calculations do-able."

References:

D. A. Mantell, K. Kunimori, S. B. Ryali, G. L. Haller & J. B. Fenn, *Surface Science* **172**, 281 (1986).

Daniel W. J. Kwong, N. De Leon & Gary L. Haller, "Desorption of Carbon Dioxide Molecules from a Pt(lll) Surface: A Stochastic Classical Trajectory Approach," *Chem. Phy. Lett.* **144**, 533 (1988).

This research was supported by NSF grant CBT8512228 and grants from Chevron Research Company and the Camille and Henry Dreyfus Foundation.

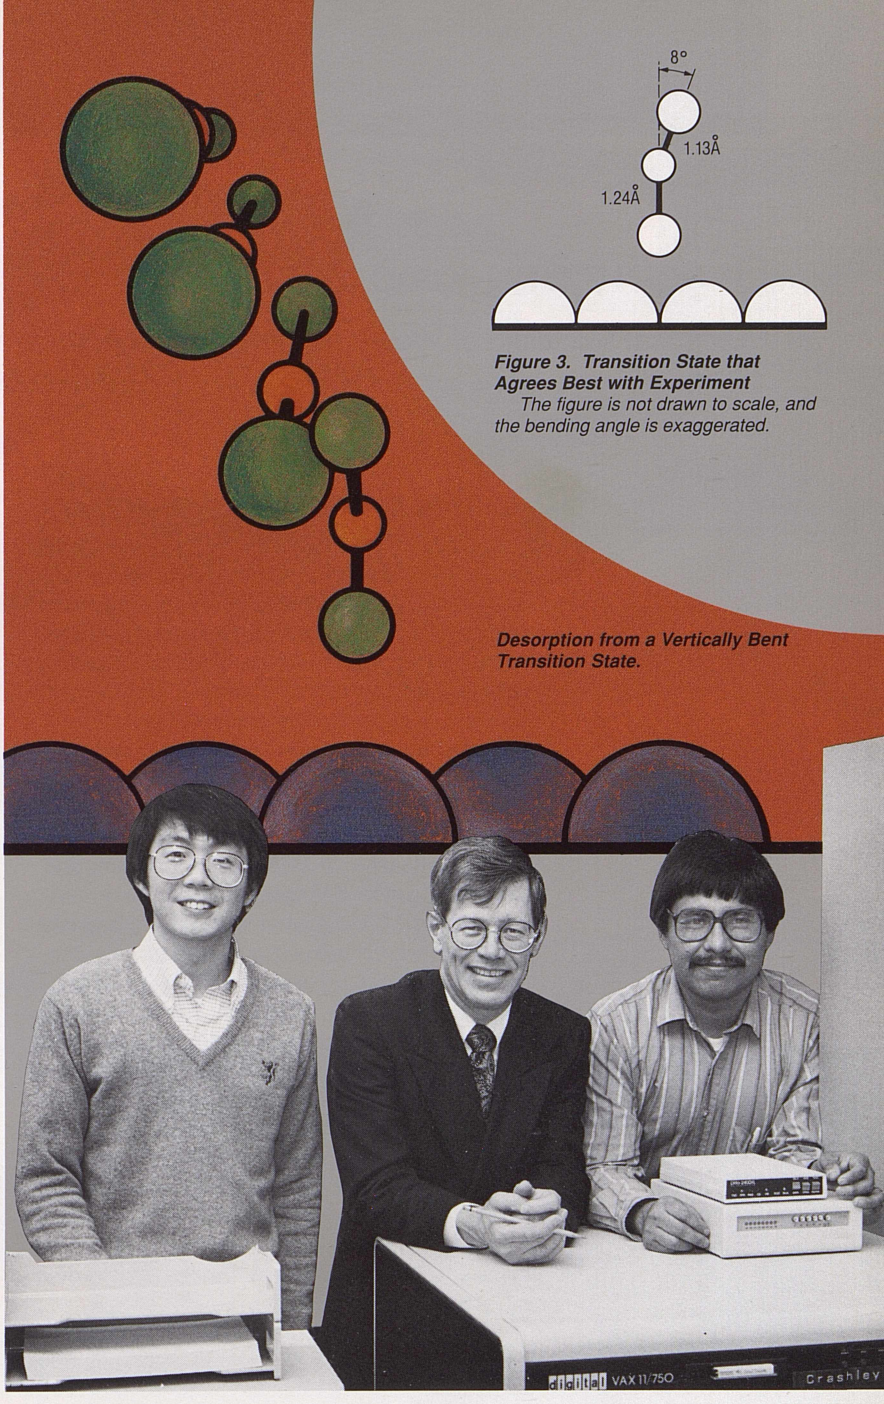

Figure 3. Transition State that Agrees Best with Experiment
The figure is not drawn to scale, and the bending angle is exaggerated.

Daniel Kwong, Gary Haller and Nelson De Leon

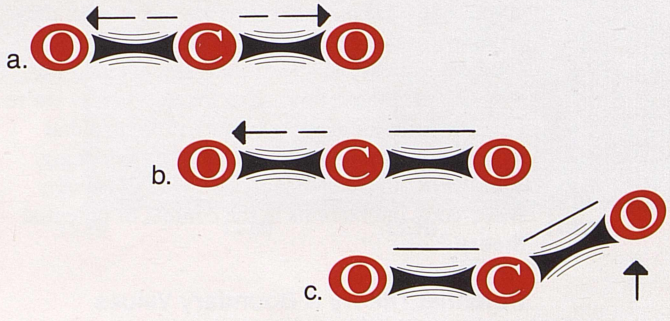

Figure 2. Normal Vibrational Modes of Carbon Dioxide
"When the molecule vibrates," explains Nelson De Leon, "its motion can be decomposed into what are known as harmonic normal modes. Carbon dioxide is a linear molecule; so two of the normal modes are symmetric and asymmetric stretches. Think of it as O-C-O and the oxygen atoms both stretching out from the carbon in-phase with one another. That's **symmetric normal mode** *(a). If you think of the two oxygens stretching out-of-phase with one another, you have the* **asymmetric normal mode** *(b). And the molecule can also bend, and that motion is called the* **bending normal mode** *(c)."*

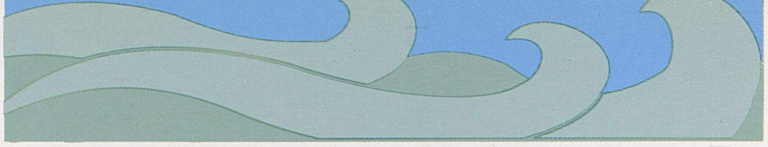

Over the Bounding Waves

Numerical Modeling of Wave Interactions with Marine Vessels

J. N. Newman, P. D. Sclavounos & D. K. P. Yue, Massachusetts Institute of Technology

Modeling Wave-Body Interactions

Breakers, swells, rollers, whitecaps—as the nautical lexicon recognizes, the tides and winds that move the oceans create waves in a myriad of shapes, sizes and velocities; the forces they exert on each other and on objects in the water are complex and difficult to analyze. "To effectively design a ship or offshore platform," says J. Nicholas Newman, "you need to have accurate predictions of the hydrodynamic forces acting between these structures and ocean waves." To provide these predictions, ocean engineer Newman and his colleagues at MIT, Paul Sclavounos and Dick Yue, work on modeling "wave-body interactions," their general term for the hydrodynamics of waves and ocean structures.

With ships and offshore structures, as with aircraft, engineering design has turned increasingly to computer modeling as a supplement and an alternative to traditional testing. "In this field, the typical experimental facilities are wave basins and towing tanks," notes Yue. "For many problems, we have developed very robust codes, and for some linear problems we're replacing the tank. For more difficult problems—very steep waves, breaking waves and others—we're not. We've made considerable progress, but we need to do a lot more."

Advances in computing hardware have spurred Newman's group to find more powerful numerical approaches to wave-body interactions. In the case of the *tension leg platforms* used in oil drilling, accurate modeling demands the power of supercomputing. "These platforms are expensive; they represent major investments of time and resources," notes Yue. "Traditionally, we had no way of doing these other than in tanks. Now we're creating generations of computer programs that require supercomputers, and these are replacing some of the important analyses. One of these is wave-body interactions in the context of potential theory."

Potential Theory & Boundary Values

Potential theory has been a staple in ship design for many years. The theory simplifies the Navier-Stokes equations, the basic equations for fluid flow, with several principal assumptions: *incompressible*, *irrotational* flow and absence of *viscosity* (or confinement in a small region). Making these mathematical terms concrete involves some inherent imprecision, but incompressibility basically means that density doesn't change with time; it amounts to saying that water particles don't expand or contract. Ignoring viscosity eliminates friction and separated flows. In an irrotational flow, vorticity vanishes—meaning roughly that certain whirlpool-like effects of the flow aren't important. "Inviscid, incompressible, irrotational flow—this is a tremendous approximation," says Yue, "but with large bodies, like the hull of a ship, it's useful."

"Potential theory describes the most important features of waves interacting with large vessels," says Newman, "but difficulties arise with complex geometries like the tension leg platforms and interactions with the *free surface*." Wave-body interactions involve two surfaces: One, the submerged surface, has fixed boundary conditions—the water contacts all portions and acts predictably acording to water depth and velocity. But on top of the water, the free surface, the boundary moves—a complicated set of velocity and pressure conditions apply. For waves of relatively low steepness, it's possible to treat the free surface linearly; and the problem is in effect solved in a fixed time-averaged physical space.

Green Functions

Newman's group has modeled these problems with Green functions, a powerful mathematical approach to complex boundary-value problems. "The source-potential, or Green function," says Newman, "provides the fundamental procedure for analyzing wave-body interactions." Roughly, it pretends that each point on the submerged body makes waves. In effect, it distributes free-surface wave sources over the submerged surface, which must be subdivided into a large number of panels to get accurate results. Computing Green functions and solving the resulting integral equations makes large demands on computing that have limited the approach's usefulness. Newman and his colleagues, however, have developed algorithms that significantly reduce processing and allow them to greatly increase the number of panels.

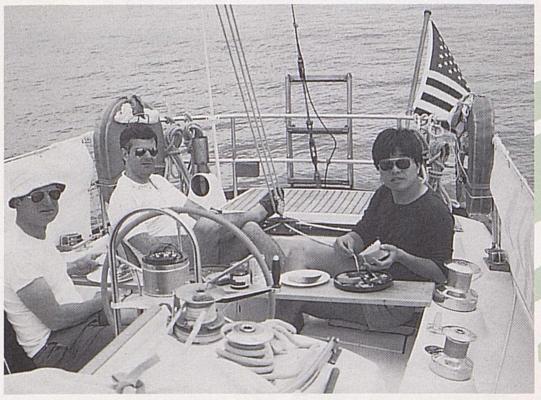

J. Nicholas Newman, Paul Sclavounos & Dick K. P. Yue on board "Katrina" during a sail off the west Netherlands coast, July 1985.

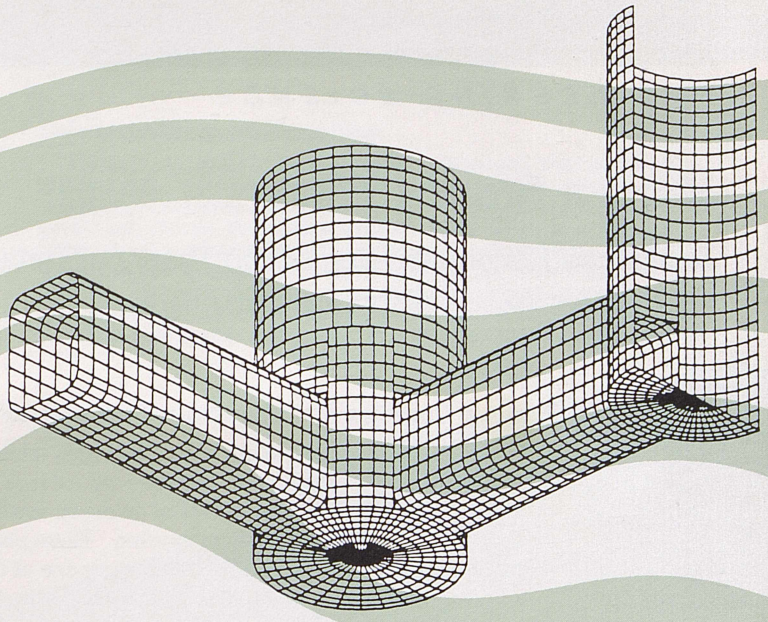

Tension Leg Platform

Submerged portion of one quadrant of a six-column tension-leg platform discretized with 3152 panels per quadrant—12,608 panels for the total structure. This structure floats with the free surface in the plane of the the uppermost panel edges.

Tension Leg Platform Results

Newman, Sclavounos, Yue and others in their group—including C. H. Lee and F. T. Korsmeyer—have implemented these algorithms (along with an accelerated iterative solver for linear equations) in a program called WAMIT (WAveMIT). When applied to tension leg platforms, WAMIT solved a configuration with 12,608 panels—3152 per symmetrical quadrant—a degree of fineness that, says Newman, "represents an order of magnitude advance over previous analyses for similar structures." Other published computations for this structure used between 240 and 992 panels with results that differed from WAMIT by 10-20% and more.

Newman's group used several different panelizations and a time-domain analysis (solving in time as opposed to the frequency domain) that appears to be the first such calculation reported. Results among these approaches agree to within three significant digits, leading Newman to conclude that a robust panel code will converge to correct answers for the hydrodynamic forces and moments as the number of panels increases. "We were able to do a more complicated problem and analyze it better partly because other people couldn't add any more unknowns," says Yue. "With supercomputing capability, the challenge is to run even bigger problems and to deal with them in all their complexities. With our codes, supercomputing lets us crank it up and use five times, ten times more unknowns. We can make successively finer calculations and show that the answers don't change. This isn't routinely done, and what we have is more or less state-of-the-art."

Computed Profiles of a Plunging Breaking Wave

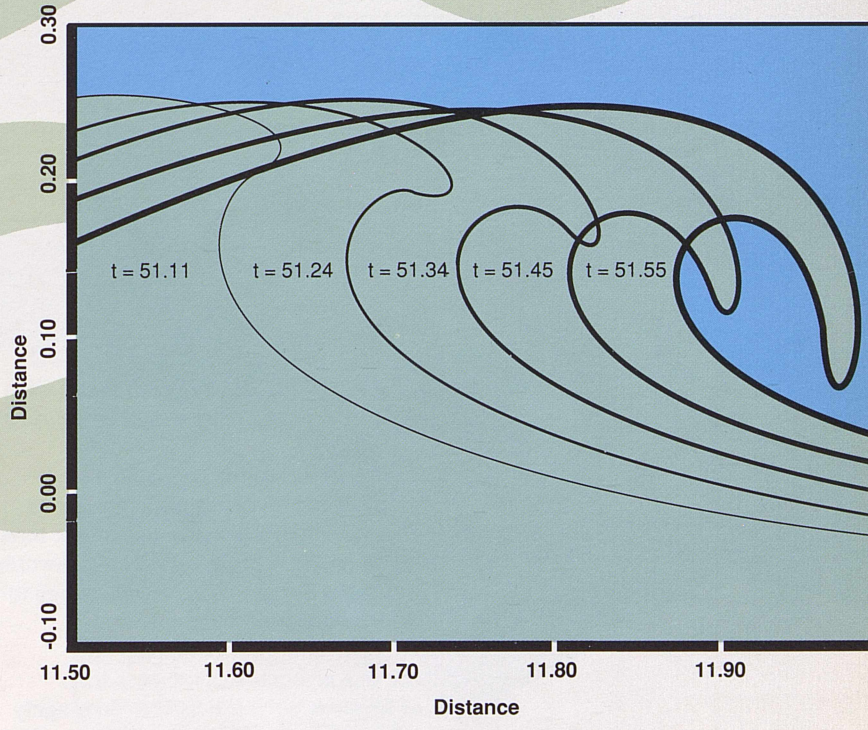

Newman's group has also developed nonlinear codes and applied them to several physical problems. For two-dimensional breaking waves, their numerical simulation agrees with waves created by a piston wavemaker in a wave tank within a few percent for the surface positions and fluid particle velocities. Yue believes that this agreement between theory and experiment up to reentry of the plunging wave shows the usefulness of fully-nonlinear potential-flow models. Their current research in nonlinear flows focuses on extending their codes to various three-dimensional wave-body problems.

References:

F. T. Korsmeyer, C. H. Lee, J. N. Newman & P. D. Sclavounos, "The Analysis of Wave Effects on Tension-Leg Platforms," invited paper, OMAE 1988 Conference, Houston, February 7-12, 1988.

D. G. Dommermuth, D. K. P. Yue, W. M. Lin, R. J. Rapp, E. S. Chan & W. K. Melville, "Deep-water plunging breakers: a comparison between potential theory and experiments," *Journal of Fluid Mechanics* **189**, 423 (1988).

This research is supported by NSF grant MSM-8514919 along with joint-industry and multiple grants from the Office of Naval Research.

With Ruysdael and Turner, if you look at the way they construct complicated water, it is clearly done in an iterative way. There's some level of stuff, and then stuff painted on top of that, and then corrections to that. Turbulent fluids for those painters is always something with a scale idea in it.

—Mitchell Feigenbaum

Many a Stormy Wind

Testing the Renormalization Group Theory of Turbulence

Victor Yakhot & Steven A. Orszag, Princeton University

A Complex and Ubiquitous Phenomenon

Turbulence, according to its Latin root, means a flow filled with turmoil. In nature we see it as dust devils, wind-whipped waves of an ocean squall, hot gases and smoke of combustion. These phenomena and others like them represent one of the important unsolved problems in physics. "Turbulence has been one of the major challenges in the physics of dynamical systems for over a century," says Steve Orszag, whose recent work on Pittsburgh's CRAY addresses the problem with a powerful theory and resourceful methodology. Progress in this area is essential because nearly all fluid engineering problems of any consequence involve turbulence: The diffusion of atmospheric pollutants, design of modern aircraft, ships, turbomachinery, combustion engines, pipelines for natural gas and oil and weather forecasting—for all of these and more, turbulence plays a critical role.

The Problem of Scale

The Navier-Stokes equations, the general equations that govern all fluid motions, have existed for over 100 years, but they remain unsolvable for all but the simplest flows. "The main difficulty," Orszag says, "is the huge range of scales for the motions in a turbulent flow. You have the large eddies where energy enters the flow, and you have very small eddies that also have to be calculated. Imagine a weather pattern occurring over thousands of kilometers and within that you have a little eddy along a wall. The range of scales from the largest to the smallest eddy would be a million to one."

To solve the equations for such a flow in realistic cases would involve a million separate variables (degrees of freedom) in each of the three spatial dimensions, 10^{18} in all—a billion gigawords of storage. And if this weren't already enough to prove the intractability of the problem . . . "Turbulent flows exhibit complicated time dependence that amounts to another million degrees of freedom, so you'll have to do roughly Avogadro's number (6×10^{23}) of operations. With today's computers, it would take you millions of years. If you figure computing speed has increased 1000 times about every 20 years, you could begin this job roughly in 2100."

Renormalization Group Theory

Physicist Ken Wilson, who until recently directed the Cornell National Supercomputer Facility, won a 1982 Nobel Prize for his work developing Renormalization Group Theory (RNG). Wilson quantified scale relationships among subatomic particles to solve phase transition problems. Orszag and his Princeton colleague Victor Yakhot have extrapolated from Wilson and others and have shifted the context from particle physics to turbulent flow. Their central idea is that small scales of turbulence are approximately universal from the viewpoint of a given large scale. "We can know collective behavior of a macroscopic object," says Orszag, "a table for instance; we can predict what will happen when we pound our fist on it without knowing every single detail of the molecular interactions."

They apply this notion through two basic hypotheses. Their *correspondence principle* generalizes the notion of a "heat bath" from statistical mechanics: an assumption equivalent to putting a basketball in a hot tub makes it possible mathematically to compute the kinetic energy of air molecules in a closed system—the basketball—as a statistical average rather than the sum of each molecule. Similarly, Yakhot and Orszag apply a "stirring force" to turbulent flow. By applying the force in a directionally random manner, they can treat the flow as in a state of equilibrium, no longer dependent on the particular velocity, pressure and other conditions, including shape, that got it started. They also choose the force so that it adds energy, realized as increased turbulence in the large flows, in an amount statistically equivalent to the energy of the small-scale flows.

This correspondence principle forms a basis for Orszag's RNG calculations. A small range of the least active small-scale flows is removed from the equations, which are then "rescaled" using a complex matching process known as perturbation theory. The operation is then repeated, removing more small-scales, to the limits of reasonable accuracy and solvable equations. "RNG gives you a theory of the small scales. Once you have a theory, you can eliminate them from the dynamics and get equations with not Avogadro's number but some modest number of degrees of freedom. And then you can solve problems."

Results: A Practical Solution

Numerical results from Orszag's simulations indicate that RNG theory properly approximates fully developed turbulence. In a wide range of tests, his large-eddy models have shown exceptional agreement with experiment—including results for boundary turbulence, the downfall of other models. These simulations, furthermore, use 100 times less storage than previous NASA studies. Orszag's large-eddy modeling appears to be, in his words, "well-suited for engineering design calculations of the most complex flows, including reacting and combusting flows."

Simulations testing RNG on complex smaller scale turbulence have also shown good results. RNG-predicted energy spectra agree closely with direct computation, and skewness, a measure of energy transfer among scales, matches well with experimental data. Of particular significance is Orszag's finding that the accepted value of the turbulent decay-rate parameter (n) is inaccurate and probably derives from an invalid assumption about flow uniformity. Along with their high accuracy, these simulations run 20 to 30 times more efficiently than previous complex flow simulations.

Supercomputing at the Limits of Capability

Orszag has used Pittsburgh's CRAY extensively to test his RNG methods. He's an experienced user of high performance computing, having literally written (edited) the book on *Supercomputers and Fluid Dynamics* (Springer, 1986). Even with RNG methods, his computations involve thousands of degrees of freedom; and with his state-of-the-art spectral codes, SHEAR3/BIGBOX, MEGABOX and others, Orszag typically uses 10 to 15 hours of CRAY time per run, hundreds or thousands of hours over the course of a year. Leading-edge turbulence research can progress no other way.

References:

V. Yakhot & S. Orszag, "Renormalization Group Analysis of Turbulence. I. Basic Theory," *J. of Scientific Computing* **1**, 3 (1986).

V. Yakhot, S. Orszag & R. Panda, "Computational Test of the Renormalization Group Theory of Turbulence," *J. of Scientific Computing* **3**, 139 (1988).

Mitchell Feigenbaum's comment on Turner is quoted in James Gleick, *Chaos* (N.Y.: Viking, 1987), pp. 186-87.

Title image: detail from "Steamer in a Snowstorm," J.M.W. Turner.

Crow Instability and Beltrami Flow

The two frames depict an RNG simulation on Pittsburgh's CRAY examining Crow instability. "It's a fundamental question of fluid mechanics having to do with whether the mathematics of the frictionless Navier-Stokes equations is well posed," says Orszag. The lift which keeps an airplane flying comes from two rotating flows or vortices produced by airflow over the wings. These trailing vortices, as they're called, align parallel to each other and rotate in opposite directions. Behind the plane, they usually merge and condense into what we see as a contrail. The two frames show three-dimensional representations of these vortices downwind where a perturbation has introduced highly complex turbulence.

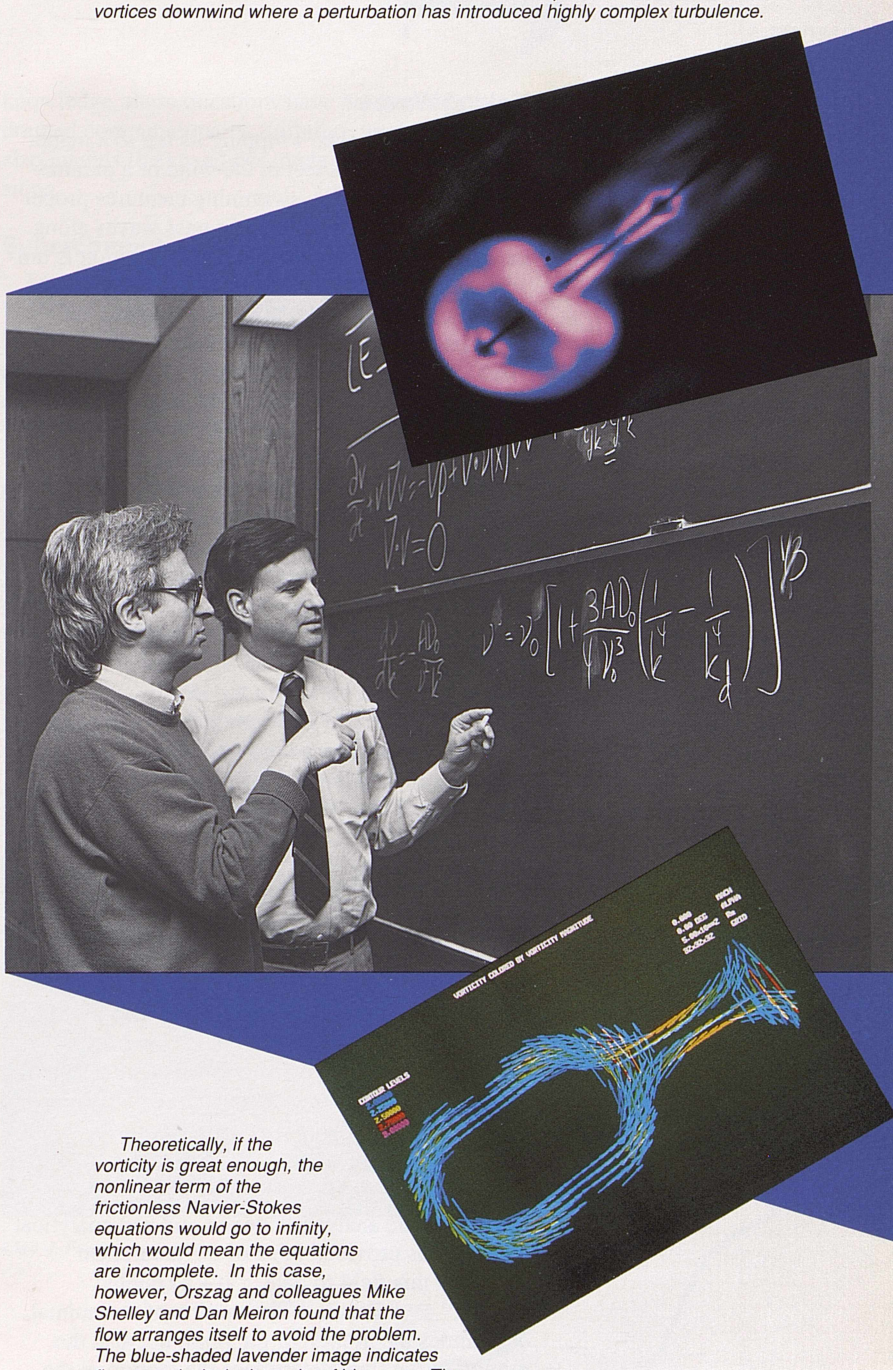

Theoretically, if the vorticity is great enough, the nonlinear term of the frictionless Navier-Stokes equations would go to infinity, which would mean the equations are incomplete. In this case, however, Orszag and colleagues Mike Shelley and Dan Meiron found that the flow arranges itself to avoid the problem. The blue-shaded lavender image indicates flow complexity by intensity of blueness. The region of greatest complexity is the small region roughly in the center of and crossing through the vortices. At the equivalent region on the other image, which shows vorticity vectors (with magnitude related to color), the vorticity direction nearly parallels the flow, or—to give its proper name—it's almost a Beltrami flow. *The term in question goes toward zero rather than infinity. "The flow is reacting to itself," explains Orszag. "It's trying to become more complicated, but then it says* no, I want to simplify myself as much as possible. *It's a very complicated flow, but there's no unnecessary complication."*

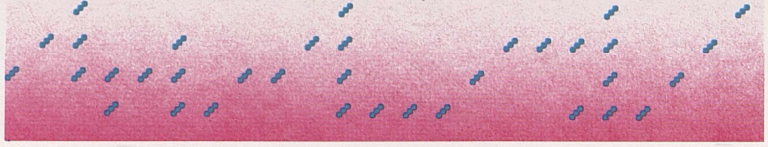

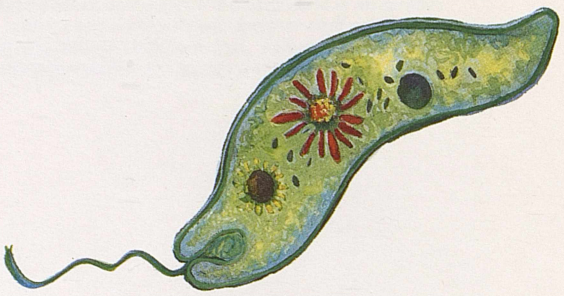

They Swam and They Swam

Mathematical Modeling of Aquatic Animal Locomotion

Lisa J. Fauci, Tulane University

Making Waves

From a watersnake whipping its tail to a manta ray flapping like a sheet in the wind or a paramecium rippling its cilia, swimming creatures propel themselves by undulations that pass waves along their body. "This kind of propulsive motion is universal all the way from microorganisms up to fish," says Lisa Fauci. "They all swim in the same geometry. And what's interesting is that the mechanisms are very different."

The difference, explains Fauci, has to do with the relation between viscosity and inertia. To what extent is the force necessary to swim determined by the resistance of the fluid, and to what extent by the tendency of the moving swimmer to keep moving—to glide with no additional effort? "If you could put yourself on the scale of a sperm, swimming would be like being in a pool of tar. You would move your arms very slowly, and that's really different from how an eel swims, kicking off the water and moving with the inertia."

These differences are important in understanding a fascinating phenomenon of aquatic animal locomotion, the human form of which can be observed in Esther Williams movies—synchronized swimming. "Say a single sperm is wiggling along and another swims up to it. After awhile they *phase-lock*. One's the copy of the next. But when fish swim next to each other, it's different. They *opposite phase-lock*—like a mirror image of each other. If one tail goes up, the neighboring tail comes down." Using Pittsburgh's CRAY X-MP in synchrony with an innovative and powerful method, Fauci has isolated previously undetected relationships in the speed and energy dissipation associated with synchronized swimming.

The Immersed Boundary Method

Asymptotic analysis, a pencil and pad analytical method, has provided most of the theoretical results in this field to date. "It's a coupled system," says Fauci. "The motion of the animal determines the motion of the fluid, but also the motion of the fluid determines the motion of the animal. Coupled systems are complicated, and the studies done in the past had to study very idealized cases." In order to solve the Navier-Stokes equations, the classic equations of fluid motion, asymptotic analysis is limited to small amplitude waves and zero or near-zero Reynolds number—a dimensionless quantity that represents the ratio of inertial to viscous effects.

To overcome these restrictions, Fauci collaborated with Charles Peskin and applied immersed boundary technique to aquatic locomotion. Peskin has developed this technique in stages over a period of years modeling the human heart. The force of the organism on the fluid is treated as if it acts only at points touching the organism—the immersed boundary. "If you aren't sitting at a point in the fluid that coincides with the organism," says Fauci, "you won't feel its presence. But if you're on the backbone of the organism, you're going to feel the force."

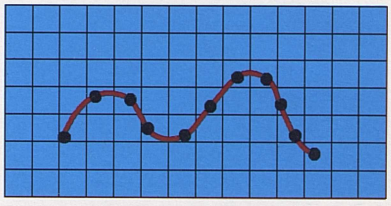

The Immersed Boundary Model

In the computational model, the fluid domain is represented by a square grid and the swimming organism by the immersed boundary, seen as a collection of moving points. In addition to more complicated wave motions, the Peskin-Fauci model can account for transient effects, such as a single, isolated whip of a fish's tail. So far as Peskin and Fauci know, their model is the only computational model that doesn't rely on steady state.

Simulating Synchronized Swimming

The immersed boundary model can also simulate multiple animals swimming near each other, a capability that Fauci tested on Pittsburgh's CRAY X-MP. She compared her model with asymptotic analysis results for two waving sheets (filaments of infinite extent) swimming closer than the wavelength of their undulation. Her results generally agreed with the prior studies, including the finding that sheets locked in opposite-phase used more energy but swam faster than sheets locked in-phase. "This is what you expect. In opposite-phase the sheets help each other swim. The motion is similar to peristalsis. They push the fluid between them in the same direction."

Fauci also ran two case studies examining what happens when animals swim near each other at different speeds. "You can start things out swimming not at all in synchrony, and for low

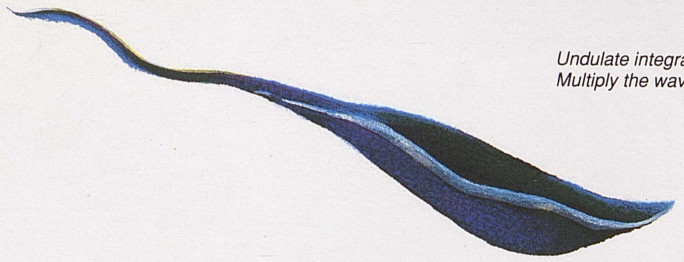

Undulate integrate from gills to tail I slither.
Multiply the waves do I to take me yon and hither.
—M. William Rosen
Biomedical Engineering,
University of California, Berkeley

Filament Pairs with Velocity Streamlines

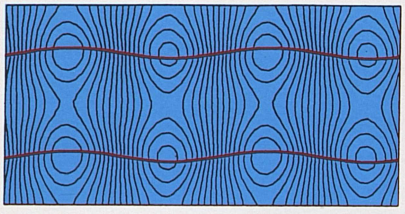

In-Phase

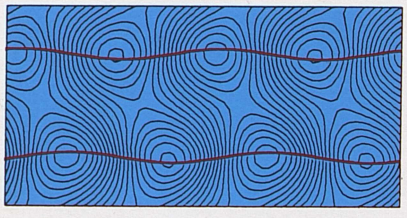

90° Out-of-Phase

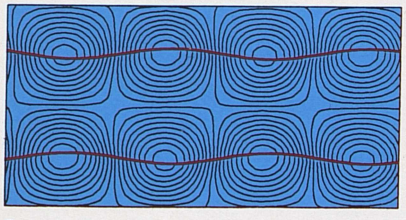

Opposite-Phase

Reynolds number and close proximity, they end up in-phase. The motion evolves into one which minimizes expended energy." For these initial conditions, the initial swimming speed of both filaments decreases until they synchronize. One might intuitively guess that this would always happen, but a second study showed otherwise. "If I make the Reynolds number higher and space the filaments farther apart, they end up swimming as mirror images." In this case, energy loss increases until they lock in opposite-phase with swimming speed maximized.

Fauci conjectures that the two cases represent equilibria, one stable and the other unstable depending in a particular case on Reynolds number and filament proximity. Her study showed the model's capability to follow the evolution of these nonlinear, dynamical systems, and she proposes in further research to determine more exactly the relationships among oscillating filaments. To her knowledge, these questions have not been addressed computationally. In future work, Fauci also intends to expand her model to three dimensions.

Supercomputing and Real Time Science

Fauci began her aquatic animal study on a VAX 11/780. "I'd submit my job overnight and come back the next morning only to find that I had a bug, and I'd have to wait another whole day to see anything. You can waste a week, really, with what could take a couple of hours now. But using the supercomputer is not just saving time for that program and finding the bug. You can do real time science. You can think of something and try it and see immediately whether it's giving you the output you want. It makes a big difference on how you plan. You can try out many ideas that you couldn't before. It's a very exciting transition."

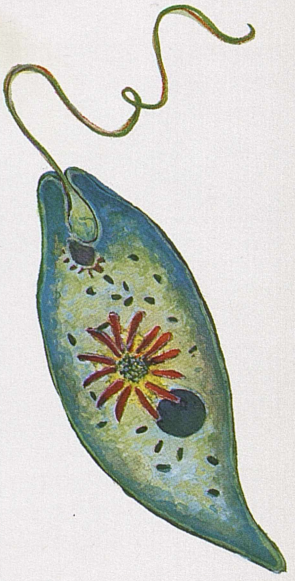

References:

Lisa J. Fauci & Charles S. Peskin, "A Computational Model of Aquatic Animal Locomotion," *J. Comp. Phys.* 77, 85 (1988).

Lisa J. Fauci, "Interaction of Oscillating Filaments—A Computational Study," *J. Comp. Phys.*, in press.

This project was funded by NSF Grant DMS8700858.

Lisa J. Fauci

Pictures Worth a Million Words

Computer Images and Orbits of 2-D Markov Chains

Marc A. Berger, Carnegie Mellon University and Michael F. Barnsley, Iterated Systems, Inc.

Scientists will (I am sure) be surprised and delighted to find that not a few shapes they had to call grainy, hydralike, in between, pimply, pocky, ramified, seaweedy, strange, tangled, tortuous, wiggly, wispy, wrinkled, *and the like, can henceforth be approached in rigorous and vigorous quantitative fashion.*
—Benoit Mandelbrot

Nature and the Complexity of Images

A mountain meadow in bloom, October foliage of the Alleghenies—words cast only vague hints at nature's infinite range of color and form. This complexity also pushes computer technology to its limits. Even simple images need vast reservoirs of memory, and existing image compression methods are ineffectual for all but basic geometric forms. "If you take a screen of 1000 x 1000 pixels, about the resolution of a high-quality graphics workstation," explains Marc Berger, "you have to store eight bits for each pixel to have a color range of 0 to 255. So one color picture takes at least eight-million bits, a megabyte of memory."

Even at high transmission rates, a few images can tie-up communication lines for hours, a factor that inhibits using computer images to interpret research. Compression is even more important in transmission and analysis of satellite images and flight simulation. Digitized landscape images are more realistic than conventional simulator images, but input/output and storage requirements defeat the purpose of pilot training. "They simply don't have a good way to store all the data," says Berger, "or to pump it in and out of the computer fast enough for the pilot to see it in real time. When he's actually flying, you know, they can't say *Stop, we have to change disks.*"

Self-Similarity via Simple Arithmetic

Berger has harnessed Pittsburgh's CRAY to a promising compression technique developed by Michael Barnsley at Georgia Tech, where Berger spent eight years. Barnsley discovered that randomness applied to simple linear equations can generate *fractals*, the self-similarity across scales that Benoit Mandelbrot identified as an element in nature's repertoire of complexity. The technique works through an algorithm that randomly chooses one among a predetermined set of mapping rules, called affine transformations, which map the image point-by-point.

You begin from an arbitrary point in a two-dimensional plane, and the rules then locate the next point, which becomes the new starting point. A set of two rules, for instance, would be like flipping a coin. Heads, go halfway toward the origin in the x-direction and half a unit up; tails, go a third of the way toward the origin in the y-direction. Each transformation must be *contractive*: it must move a set of points closer together. If this condition is met, the coin flips will reveal not just a random collection of dots; eventually a definite shape will emerge.

Compression and Supercomputing

Four transformations generate the maple leaf image shown at right. This means that 24 numbers—eight linear equations, each precisely identified by three numbers—describe the image. Conservatively, that's a compression of three to four orders of magnitude, hundreds versus millions of bytes. Furthermore, because affine arithmetic involves only addition and multiplication, the algorithm is very fast; decoding an image takes only seconds. And its randomness makes it ideal for vectorizing. Instead of a million sequenced iterations, the CRAY can run 64 independent sets for 16,000 iterations, an advantage Berger has used for animations. "As quickly as I can generate one image, the CRAY generates 64. In 15 to 20 seconds, we have the whole movie." Berger estimates it would take 100 times longer on a VAX.

Encoding and Parameter Space

As this work progresses, supercomputing will play a central role. Berger works with an interdisciplinary group at Carnegie Mellon—Bill Eddy in statistics and H. Meté Soner in mathematics along with Ph.D. students Mario Peruggia (statistics) and Jean-Philippe Vidal (computer science)—on various aspects of the encoding problem: Given an image, how do you find the transformations that represent it? This thorny question is the main obstacle to broad application of the affine method, and potential solutions involve computationally demanding fitting schemes. The group is using Pittsburgh's CRAY to test several encoding algorithms.

In related work, Barnsley and Berger have extended the basic theory by mixing images from more than one set of transformations. They have also produced a series of animations that probe "parameter space." By interpolating between the transformations that produce identifiable images, they produce new transformations, each with its own associated hybrid image. Decoding these transformations onto video tape creates the effect of one image flowing smoothly into another. The

A Walk Through Parameter Space

"You can take any of the images we have and walk from one to the other, and it will be a continuous deformation. We take the 24 numbers that correspond to the leaf and the numbers that correspond to the fern, and we interpolate to see what's between them. Every point in that 24 dimensional space is a picture. It's our way of probing into this parameter space. What are the variations? What does it look like when you go from one to the other?

"With the leaf-to-fern, I was surprised; there's a period in the middle where the leaf almost disappears. It shrinks down and then blossoms out as a fern. We had no idea what to expect."

viewer travels a newly opened visual pathway with surprises along the way even for those familiar with the technique—as if nature has revealed new facets of her complexity. These *walks in parameter space* are a valuable form of basic research in an emerging discipline.

References:

Michael F. Barnsley, *Fractals Everywhere* (N.Y: Academic P., 1988).

Michael F. Barnsley, Marc A. Berger & H. Meté Soner, "Mixing Markov chains and their images," *Prob. Eng. Inf. Sci.* **2**, 387 (1988).

Marc A. Berger, "Images Generated by Orbits of 2-D Markov Chains," *Chance* **2** (1989), in press.

This research is supported by AFOSR grant 87-0137.

Marc Berger and Michael Barnsley

What else, when chaos draws all forces inward
To shape a single leaf?
— Conrad Aiken, "The Room"

Maple Leaf

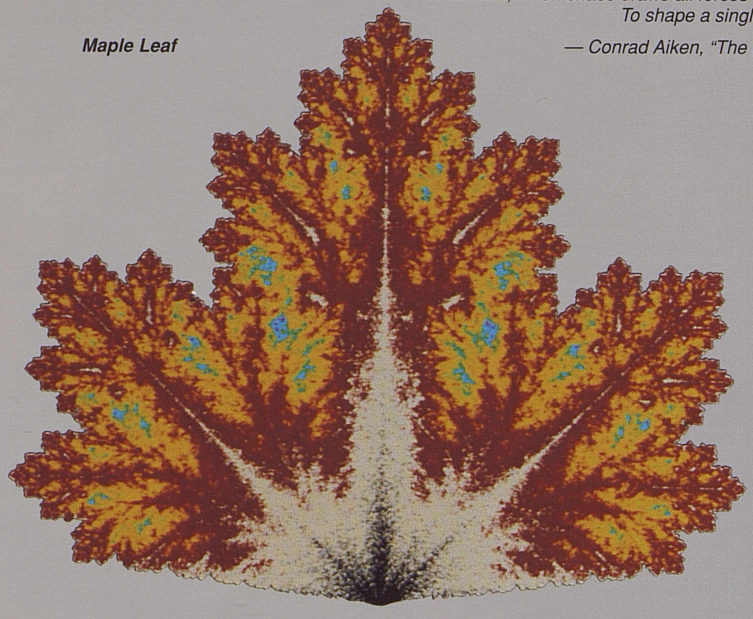

"It's a very complicated shape. The boundary of the leaf is a very intricate curve, and to keep track of the boundary is just as hard as keeping track of the image. Sometimes it creeps inside because there are little veins, and sometimes the boundary can almost fill up the image. Everywhere you have holes and veins, and the simple kind of compression techniques break down."

Image Generation Algorithm

Each affine transformation, T, is a pair of equations:

$$T(x, y): \quad x' = ax + by + c$$
$$y' = dx + ey + f$$

where x' and y' represent the shifted (transformed) coordinates of an arbitrary starting point (x, y) in the x-y plane. N is the number of transformations required to map a particular image. Each transformation has an associated probability, p, the total of which must equal 1. To color the image, count the number of times each pixel is plotted as a point, and use a color map to convert these frequencies to colors.

The process tends to be counter-intuitive; you expect a random collection of dots. Says Berger, "It leaves people asking, What's going on here?"

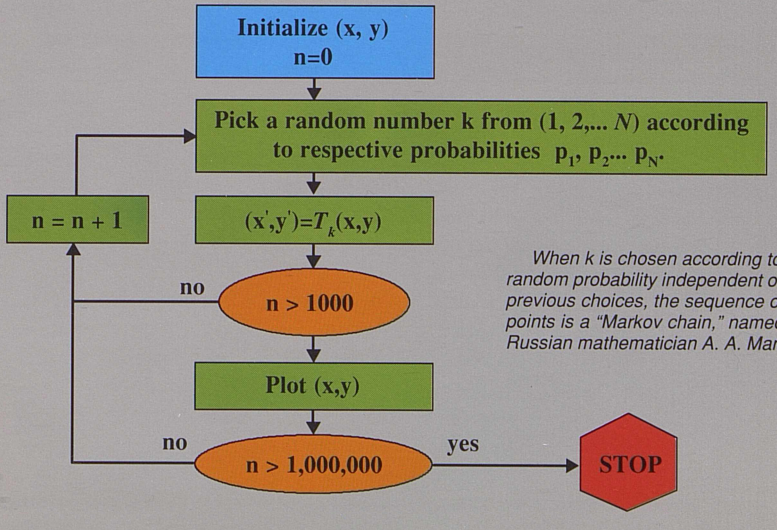

When k is chosen according to a random probability independent of previous choices, the sequence of points is a "Markov chain," named for Russian mathematician A. A. Markov.

When the Earth Moves

Investigating the Complexity of the Earthquake Rupturing Process

Shamita Das, Lamont-Doherty Geological Observatory of Columbia University
B. V. Kostrov, Institute of Physics of the Earth, Academy of Sciences, Moscow, USSR

Earthquake Prediction in the Age of Science

By compiled data, earthquakes are among the deadliest of natural catastrophes—an average toll of 20,000 a year in the twentieth-century. Furthermore, in their potential for devastation on a grand scale, earthquakes are probably the deadliest event, short of nuclear holocaust, that humans can anticipate. The Armenian quake, for instance, was only moderate in force — 6.9 on the Richter scale. Not even enough, explains Shamita Das, to trigger all the distant seismic instruments which record the data used in earthquake study. The quake that struck Tangshan, China, July 27, 1976, was ten times more powerful, 8.0 Richter. Again there was no warning. It took 650,000 lives.

Despite the data and the danger, despite occasional successful predictions (and just as many false alarms), earthquake study has barely hatched from its shell of primitive superstition and myth. Not until the 70s, when government support began to materialize, did American scientists turn to full-time research in earthquake prediction. Still, a 1982 survey of the field states that no reliable prediction system exists and that at best some limited assured capability may be viable in a decade.

Karl V. Steinbrugge, Earthquakes, Volcanoes, and Tsunamis, *pp. 33-35.*

Where was Science?

December 7, 1988, 11:41 a.m.—the clock tower in Leninakan's central square froze the moment. The earth rumbled and 40,000 people lost their lives. A natural tragedy half-a-world away has seldom reached us with such force as the earthquake in Armenia—scenes of devastation, shocked survivors, futile outrage at their government: *Why was there no warning?* they asked. *Where was your great God science?*

Reliable earthquake prediction is the long-term goal that Shamita Das and B. V. Kostrov work toward. In Das' particular arena, earthquake source mechanics, research revolves around such questions as: What characteristics of a fault—what stresses and what strengths or weaknesses of the fault material—signal that it is "ripe," that the earth will soon release built-up stress? And what characteristics signal enough stability to allow the inverse prediction — that an earthquake will not occur? Source mechanics also has the more immediate aim of providing the information needed to build structures that will withstand earthquake shock. "That's actually the more important aspect of my work," says Das. "How much will the ground shake when an earthquake *does* occur? Basically, with this information, engineers can build buildings that won't fall apart. And this saves many lives."

The Question: Asperity or Barrier?

Das and Kostrov's research jumps directly into the fray of two conflicting theories, the barrier model and the asperity model. "It's a long story actually," says Das. "Earthquakes happen on faults in the earth, and it is believed that a fault may be weak everywhere except some regions." How to interpret these regions generates the controversy. The asperity model treats them as asperities, areas of anomalously high-strength that will give way only under very heavy loading and consequently generate extremes of stress release—a major earthquake.

On the contrary, says the barrier model, the severity of an earthquake doesn't depend on the strong parts of the fault; it depends on the weaker regions, where the fault is far more likely to give way. The strong regions, the barriers, block dislocation along the fault and fragment it into separate quakes—in all likelihood reducing peak severity. "Whether you identify the stronger regions as barriers or asperities," says Das, "makes a big difference for the next earthquake. It's a difference in the physics of the problem. And we don't really know which is correct. What's dangerous is to predict earthquakes based on one model or the other."

Arguments supporting either model must rely on correlating the form of an earthquake's seismic wave, as recorded at seismic monitors, with field studies of the fault that produced the quake. In the asperity model, peaks in the wave form correlate with broken asperities. In the barrier model, however, peaks indicate rupturing between barriers. "In our opinion, part of the conflict arises because people haven't analyzed the data properly, and better methods might resolve some of the important issues."

Thirty Earthquakes via Supercomputing

As their test case, Das and Kostrov used a square fault with embedded barriers at regular intervals and one or a few asperities (Fig. 1). "We just wanted the most schematic examples," says Das. "When you want to disprove a theory, you have to find only one model that doesn't fit. We didn't want to be misled by complexities we didn't understand." They ran more than thirty simulations, an unprecedented amount, that allowed them to vary and constrain fault parameters one-at-a-time. The simulations progressed systematically through treatments of barrier density, faults with friction, a fault with barriers and friction, and multiple asperities on a fault with distributed barriers and friction.

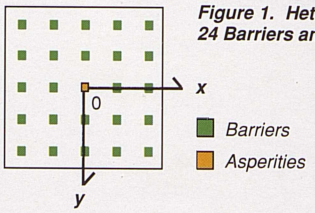

Figure 1. Heterogeneous Fault with 24 Barriers and a Central Asperity

To help keep the computations manageable, Das and Kostrov applied their boundary integral equation method. With this numerical approach, as opposed to the more commonly used finite element

and finite difference methods, the dimensions of the problem reduced significantly. "You have to look only at the crack and the plane of the crack; you don't need to keep track of the disturbance in the entire medium." By recent improvements, the grid is applied to the slipping part of the fault plane only, instead of the complete stressed area, a refinement that Das has found can reduce computations as much as two orders of magnitude.

"Even one simulation would be hard to do without the supercomputer." Because of the size of the fault and the resolution of sampling on the fault, roughly 1000 grid-points, Das' matrices exceeded the eight-megaword capacity of the CRAY X-MP. As a result, she used the solid-state disk to buffer the calculations. "The main advantage of the supercomputer for me is the speed. Nobody has the speed of the CRAY; that's the big difference." Das has extensively vectorized her code, and she estimates an overall speedup factor of 10,000 over conventional computing.

Results: A Healthy Dose of Complexity

"We showed that the wave form can have peaks without asperities. They can come from friction and from other properties of the fault." Das sees this as the study's central finding. The simulations showed that three separate factors contribute to the seismic wave form: asperities, which register high stress drop when they rupture; barriers, which have an effect that depends on the density of their distribution along the fault; and friction, causing jerkiness in the fault rupture process. In conclusion, say Das and Kostrov: "None of these three causes of waveform complexity can be considered more important than the others."

References:

Shamita Das & B. V. Kostrov, "An Investigation of the Complexity of the Earthquake Source Time Function Using Dynamic Faulting Models," *J. of Geophysical Research* **93**, 8035 (1988).

B. V. Kostrov & Shamita Das, *Principles of Earthquake Source Mechanics* (N. Y.: Cambridge Univ. P., 1988).

This work was supported by the National Science Foundation under grant RII-85-03099 and by the U.S. Geological Survey under grant 14-08-0001-G-1389.

The Walls Come Down and Science Prospers

B. V. Kostrov heads the seismology group at Moscow's Institute of Physics of the Earth, the world's largest center for earth science research. The collaboration with Shamita Das ensued from a U.S.-U.S.S.R. agreement in environmental protection. Working jointly since 1982, they have together refined Das' boundary integral equation method and produced a series of papers and a recent book. When the Pittsburgh Supercomputing Center opened in 1986, Das applied for time, and supercomputing became a key partner in this international collaboration. "We work out the mathematics and the mechanics in Moscow," explains Das, "and then I use the CRAY for the demanding computations. The collaboration has been extremely productive."

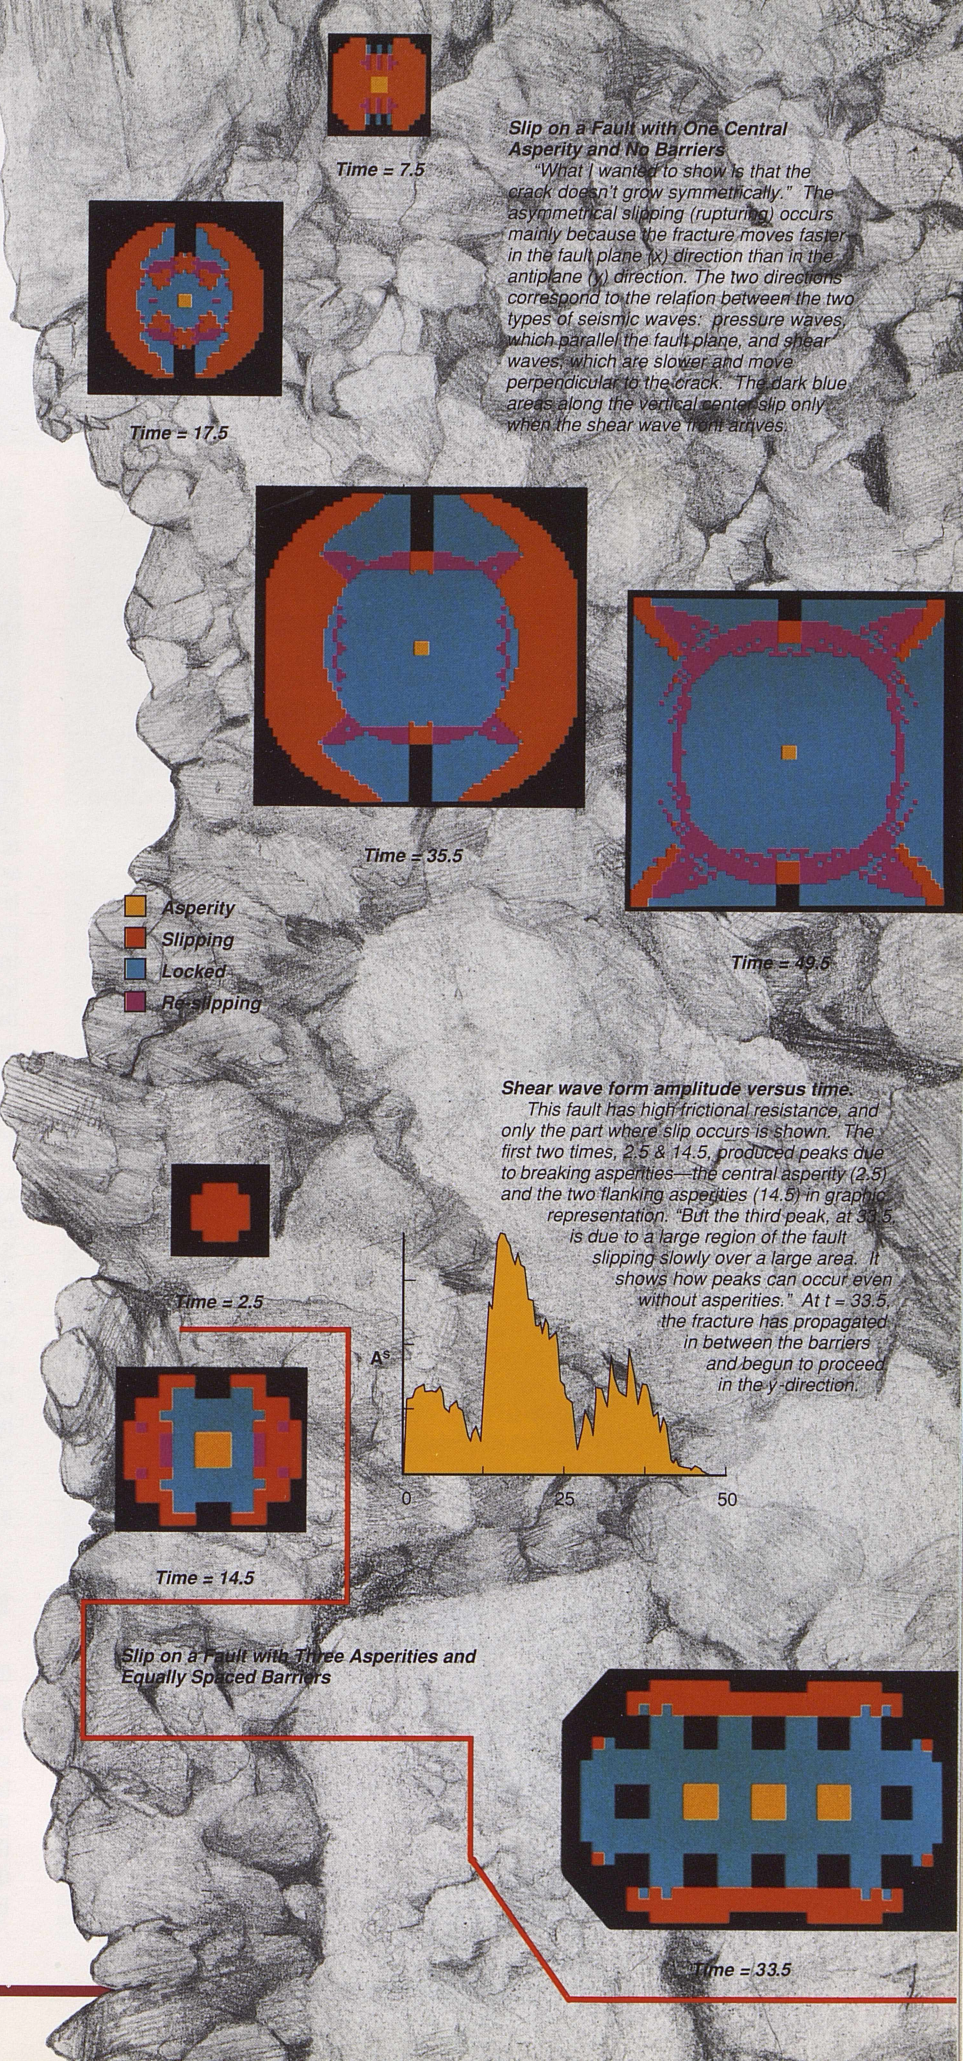

Slip on a Fault with One Central Asperity and No Barriers

"What I wanted to show is that the crack doesn't grow symmetrically." The asymmetrical slipping (rupturing) occurs mainly because the fracture moves faster in the fault plane (x) direction than in the antiplane (y) direction. The two directions correspond to the relation between the two types of seismic waves: pressure waves, which parallel the fault plane, and shear waves, which are slower and move perpendicular to the crack. The dark blue areas along the vertical center slip only when the shear wave front arrives.

Shear wave form amplitude versus time.

This fault has high frictional resistance, and only the part where slip occurs is shown. The first two times, 2.5 & 14.5, produced peaks due to breaking asperities—the central asperity (2.5) and the two flanking asperities (14.5) in graphic representation. "But the third peak, at 33.5, is due to a large region of the fault slipping slowly over a large area. It shows how peaks can occur even without asperities." At t = 33.5, the fracture has propagated in between the barriers and begun to proceed in the y-direction.

Slip on a Fault with Three Asperities and Equally Spaced Barriers

Clearing the Fog from Rainfall Models

Atmospheric Stabilization and the Distribution of Cumulus Clouds

Jorge A. Ramirez & Rafael L. Bras, Massachusetts Institute of Technology

From Magic to Supercomputing

Weather forecasting has come a long way since Joseph interpreted the Pharaoh's dream of fat and lean cows. Science and computing have brought high reliability to a practice steeped for centuries in omens and oracular wisdom. Yet last summer's drought and the ordeal of Hurricane Joan painfully accentuate the need for even more accurate, timely forecasting, both for long-range trends and hour-by-hour changes. With this need in mind, MIT hydrologists Rafael Bras and Jorge Ramirez have worked at improving hydrological models for predicting rainfall in river basins.

Bridge Over Troubled Models

Hydrological precipitation models are predominantly stochastic. "We describe space-time precipitation as a random field," says Ramirez, "and we try to increase knowledge by looking at past data." Collection points within a watershed area provide the data. Using probability models developed from this data, hydrologists can project precipitation throughout the watershed. These models have become fairly successful at reproducing observed rainfall, and their predictions have the added benefit of an attached reliability measure. If the forecast says the chance of rain today is 20%, for example, a farmer can rationally assess the risk of planting.

These methods differ from meteorological models of climate and weather on a global scale. At these large scales, the models include mathematical expressions for the physical processes involved—jet streams, solar radiation and other phenomena. The stochastic models, say Bras and Ramirez, only "conceptualize" physical processes. A prime example is assuming that cumulus clouds are clustered, a conceptualization, notes Bras, that can have significant impact: "You're concentrating the convective phenomena in a smaller area, and you'll have higher rainfall intensity." The problem is that the assumption lacks a rigorous physical basis.

The gap between the two model types defines Bras and Ramirez' ultimate goal: *physico-stochastic* precipitation models. "We're trying to come to a middle point," says Ramirez, "a model that resembles the hydrologists' in the sense that it's readily usable for forecasting at river basin scales and that also has some solid physics in the parameters, even if not the whole thing."

Stabilization Function and the Inhibition Hypothesis

To represent the physics of cloud convection, Bras and Ramirez have developed a comprehensive mathematical expression known as a *stabilization function.* The idea, explains Bras, is to capture all the effects of a small area of convection on the surrounding environment. This small area, usually referred to as a convective cell, can be thought of as a warm bubble of air that rises from ground level to form a cloud. As the bubble ascends, its thermal instability—its warmth relative to surrounding air—dissipates through condensation and other interactions. Ultimately, it loses buoyancy and becomes thermally stable. These dynamics result in a stabilization function that controls the rate of change in what Bras and Ramirez call *convective available potential energy* (CAPE) in and near the convective cell.

Convection inhibits convection. From analysis and modeling with their stabilization function, Bras and Ramirez have drawn this basic insight. Single cloud numerical simulations on the CRAY X-MP confirm that cloud convection reduces CAPE. It follows that the probability of clouds occurring is reduced near an existing cloud, a conclusion which Bras and Ramirez call the *inhibition hypothesis,* which means that the distribution of the resulting cloud-field must be regular, not clustered and not random. Says Ramirez, "Convection *by itself* will never produce clustering, it can only produce regular fields."

Testing the Inhibition Hypothesis

Hydrologists and meteorologists have not thundered their approval of the inhibition hypothesis. "People have believed in clustering for so many years that reluctance has been the rule," says Ramirez. Nevertheless, the MIT researchers believe not only that their conclusion is correct but that it should not be startling when viewed within the bounds of their study; they consider only thermodynamics and exclude external influences like wind and topography.

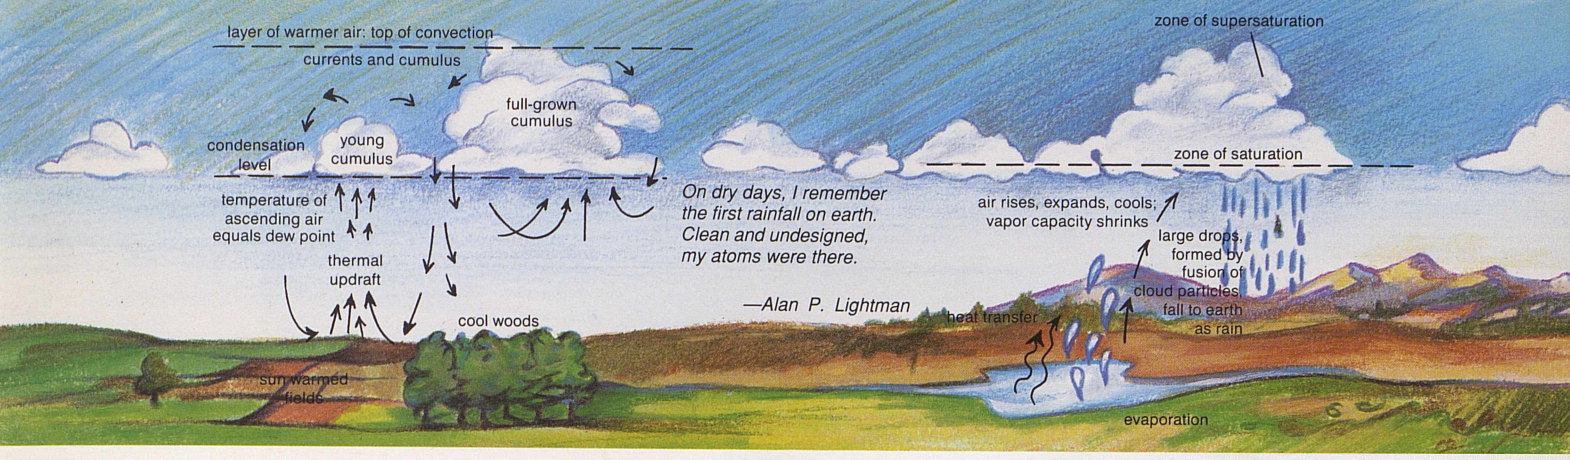

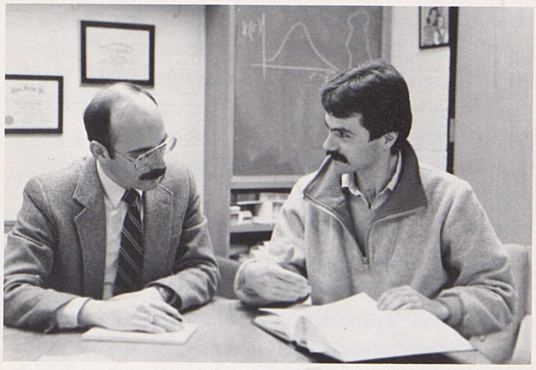

Rafael Bras and Jorge Ramirez

To test their findings, Bras and Ramirez digitized a series of SKYLAB cloud-field photos and analyzed them numerically on the CRAY. By objective statistical criteria, the cloud-fields were regular. Bras and Ramirez also ran multiple-cloud simulations using a model from the National Center for Atmospheric Research (NCAR). They initialized the model randomly so that the triggering mechanisms for convection had no built-in clustering or regularity. In each case, the resulting cloud-field was regular. "The thermodynamics altered the initializing pattern," explains Bras. "If the thermodynamics had no effect, you'd expect a cloud wherever there was a heating anomaly."

For practical purposes, says Bras, observational cloud data with sufficient resolution to infer thermodynamics doesn't exist. Access to Pittsburgh's CRAY has enabled them to test theoretical analyses which otherwise would amount to interesting speculation. "The CRAY allowed us to simulate convection under a host of different initial and boundary conditions," says Ramirez. "And supercomputing has made it possible to develop and run complex models like the NCAR model we used here. These models produce what hydrologists call *surrogate observational data* which we can then use to propose and verify new theories. Supercomputing has enabled us to do research which will have a great impact on the way hydrometeorologists view precipitation fields."

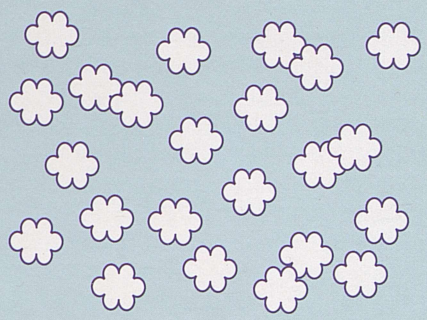

Random

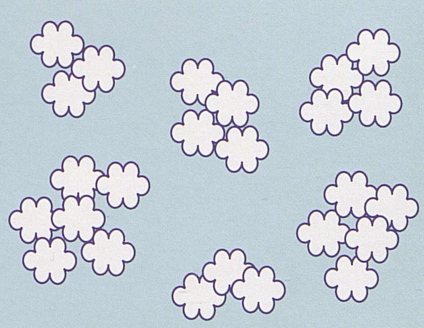

Clustered

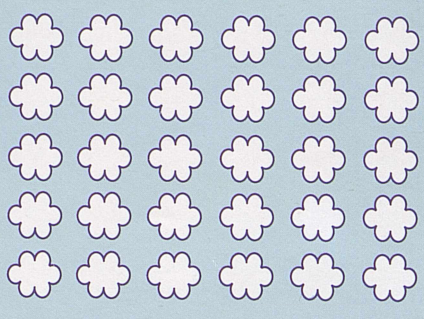

Regular

Random, Clustered or Regular?

I wandered lonely as a cloud. *Despite Wordsworth's classic line, clouds are seldom lonely. The convective processes that trigger cloud formation almost always occur at scales larger than a single cloud, and like birds of a feather, clouds flock. But how do they flock? Does convection create random, clustered or regular cloud distribution? It's a question hydrologists have considered settled in favor of clustering. Nevertheless, say Bras and Ramirez, cumulus cloud convection produces regular distribution. To illustrate the difference among the three patterns, Rafael Bras draws an analogy to throwing darts.*

"Close your eyes and start throwing darts at a blanket. Within a square piece of the blanket, chances are you'll have a fairly uniform distribution of the distances between darts. If there are small distances and big distances more or less evenly distributed, then the pattern is ***random****. If you open your eyes and try to throw the darts in the same square and you're a good shot, the darts will be* ***clustered****, there'll be many small distances between them. And a* ***regular*** *pattern would be where you want to throw a grid, or as close to a grid as possible—every dart equally separated from every other."*

Bras and Ramirez limit their hypothesis of regular cloud distribution within bounds determined by its derivation. Forcing mechanisms not intrinsic to cumulus convection—wind shear, for instance, or a mountain ridge—will affect cloud patterns and, in reality, will often disturb the regular distribution that occurs from convection alone. They also argue that cloud patterns depend on the scale of observation. The clustering of a large field, when observed at scales larger than the scale of external forcing mechanisms, will be regular at the scale of several cloud radii.

References:

J. Ramirez, "Cumulus Clouds: The Relationship Between Their Atmospheric Stabilization and Their Spatial Distribution," Ph. D. Thesis, MIT, Cambridge, MA. (1987).

J. Ramirez & R. Bras, "Clustered or Regular Cumulus Cloud Fields: The Statistical Character of Observed and Simulated Cloud Fields," *J. Geophys. Res.*, submitted (1988).

The researchers wish to acknowledge the collaboration of Professor Kerry Emanuel of MIT. Funding for this project provided by National Science Foundation and NASA Joint Grant 8611458-ATM, and the National Weather Service through a Cooperative Agreement between the Office of Hydrology and MIT's Ralph M. Parsons Laboratory for Water Resources and Hydrodynamics .

It brings the water from the sea to shore,
Transports the seed to rooting place.

. . . restless mind has yet to touch
The secret of the wind, the ceaseless wind.

—J. A. Dutton

Strangely Attracted

Quantifying the Predictability of a Coupled Ocean-Atmosphere Model

Jon M. Nese & John A. Dutton, Pennsylvania State University

Predictable Unpredictability?

In the heady days of the 50s and 60s, when computers were new and determinism still firmly held sway in many fields of science, computer modeling seemed to promise that an exact science of weather forecasting was on its way, just around the next bend in the road of progress in computing technology. Since then dynamical systems theory has been telling a much more complicated story. The weather is fundamentally unpredictable. Yet atmospheric scientists like John Dutton and Jon Nese have explored areas of predictability within the unpredictability that can give more reliability to weather forecasts.

Dutton heads the College of Earth and Mineral Sciences at Penn State and Nese is polishing his dissertation on their research together. *Predictability*—that's the key word, says Nese. "How good are the forecasts? What changes in the model affect their usefulness?" With the CRAY X-MP across the state at Pittsburgh, they applied dynamical systems theory to a coupled ocean-atmosphere weather model. Their project's roots in dynamical systems begin with work in the 60s by Edward Lorenz at MIT, and their results extend this work to suggest a useful new measure of short time-span predictability.

The Butterfly Effect & Phase Space

The reliability of any model depends in part on the data it starts with—its initial conditions. Before Lorenz's work, explains Dutton, this was seen as the main obstacle to better forecasting. "It was assumed that better data and more computer power would improve predictions without limit." Lorenz showed something different—sensitive dependence on initial conditions or *the butterfly effect*. "We're talking about a system of differential equations. In principle, there should be a single solution. But suppose you start with slightly different initial conditions." With a simple, three-variable model, Lorenz showed that a small initial difference can magnify greatly over time. In effect, a butterfly dusting pollen on a flower in China could cause a cyclone off Brazil.

Such a system is *chaotic*; its solutions don't follow a regular pattern. To get at the question of predictability for such a system, dynamical systems theory borrows the idea of *phase space* from thermodynamics. Instead of the traditional plot of solutions versus time, a phase space map consolidates all the system's variables into a single point that represents the instantaneous "state" or phase of the system. The space is defined by a set of coordinates, one dimension for each variable, and the point's movement with time charts the evolution of the system.

Attractors & Lyapunov Exponents

The phase space map of a chaotic system shows that unpredictability itself is predictable within limits. "The characteristics of the equations," says Dutton, "require that the solutions be confined asymptotically within bounded regions." The regions are called *attractors*, and by definition a chaotic system has a chaotic or *strange attractor*: Its trajectories approach each other but never intersect; a point is never repeated. Though never intersecting, adjacent trajectories will tend to converge or diverge, and this tendency for an attractor as a whole is measured by its *Lyapunov exponents*.

"It's a way of totally summarizing predictability over the entire phase space—the larger the Lyapunov exponents, the worse the predictability," says Nese. For a chaotic system, the largest Lyapunov exponent gives an error-doubling time for the attractor and essentially tells when the forecast has no value. "A computer forecast can be pretty good for two or three days, OK at five, at ten maybe some information is usable and at twenty it has absolutely no value," says Dutton. "There's a predictability limit, and no matter how much you invest or what you do, you can't get past that."

Results: Activating the Ocean

Dutton and Nese focus their study on a modified moist general circulation model (GCM) developed by Lorenz—a truncated spectral system of 28 coefficients, in effect 28 dimensions in phase space. Complex enough, says Nese, to produce a realistic atmospheric flow and moisture pattern, yet small enough that the CRAY X-MP can calculate Lyapunov exponents over long integration times. It bridges the gap between GCMs having thousands or even millions of variables and the lowest order GCMs, like Lorenz's three-variable model.

Jon Nese confers with John Dutton

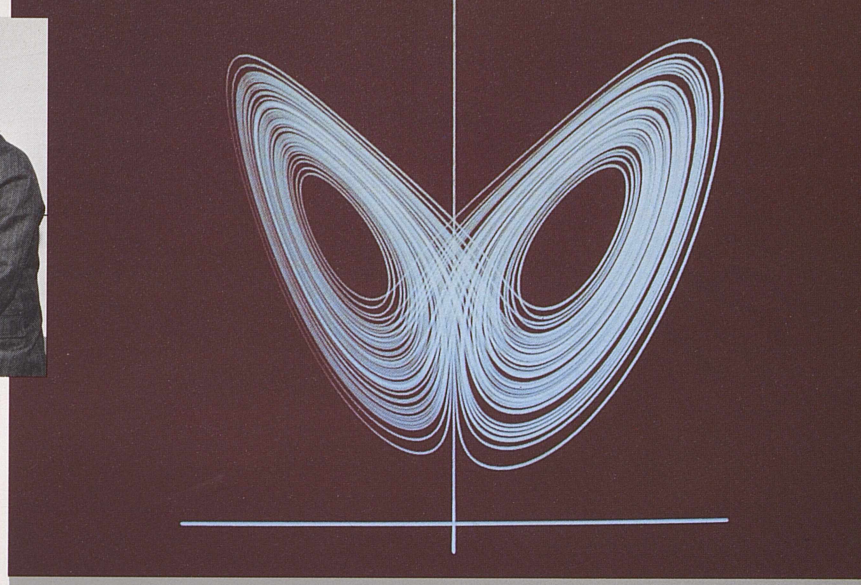

Phase Space: The Three-Variable Lorenz Model

This strange attractor, the Lorenz attractor, resides in a three-dimensional phase space and represents a system of three differential equations describing atmospheric convection.

Developing the statistics and dynamical systems characteristics of this model already taxes the CRAY, says Dutton. "We couldn't do 2800 variables, for example; 28 is reasonable with today's computational power."

One of their main objectives was to couple the ocean to their GCM and investigate how ocean-atmosphere interactions, essentially wind-driven currents, affect predictability. "You expect," says Dutton, "that the ocean, because it has more mass and is more of a thermal reservoir than the atmosphere, would have inertia and make the system somewhat more predictable. In effect, that's what it does." Coupling added seven dimensions to the model, significantly complicating the computation, and it decreased the largest Lyapunov exponent ten to fifteen percent. The direction of this effect isn't surprising, but as far as Dutton and Nese can tell, it hasn't before been quantified. This study appears also to be the first to calculate Lyapunov exponents for other than very low-dimensional systems.

Their work on overall predictability led Dutton and Nese into a study of local divergence rates, a way of quantifying variations in predictability at different locations on an attractor. "In physical terms," says Dutton, "this means that the weather is more predictable at some times than others. We can say when we forecast from a particular point on an attractor, *Well, this forecast is good for 12 days*. And somewhere else on the attractor, we say *Uh-oh, dangerous territory; this one's good for only two days*. So there's a whole new sub-strata of predictability. This is very useful information, and it confirms what we've suspected for a long time."

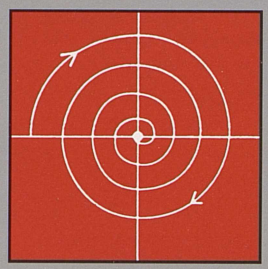

A Single-Point, Steady-State Attractor

This phase space map represents a pendulum slowing down as it loses energy to friction. The trajectories show the state of the system in two dimensions: position left- or right-of-center on the x-axis and velocity on the y-axis. They spiral inward toward an attractor at the origin: a single point that represents steady-state—the pendulum stops.

Local Predictability of the Lorenz Attractor

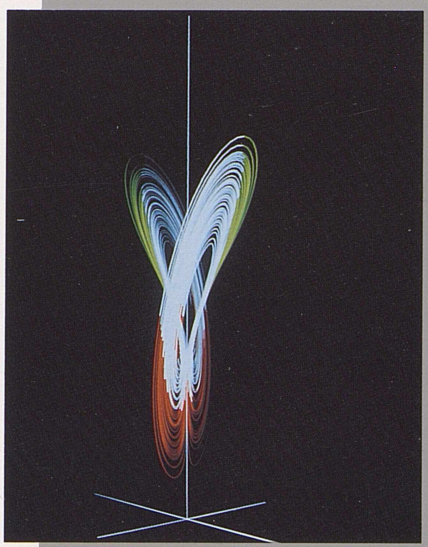

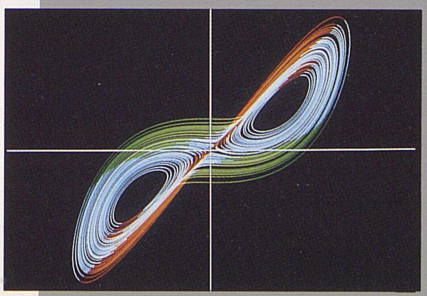

Color-coding indicates variations in local predictability. Green areas, negative divergence rates, have the greatest predictability. Red areas, positive divergence rates, indicate unpredictability more or less equivalent in magnitude to the predictability of the green areas. The white areas have divergence rates that range between the green and red values, and are not dramatically predictable or unpredictable.

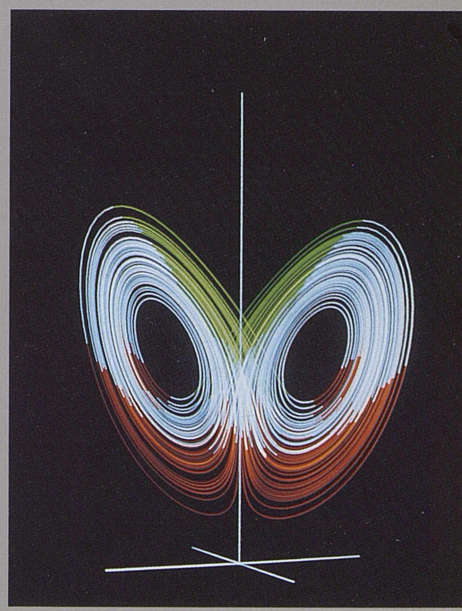

References:

J. M. Nese, "Quantifying Local Predictability in Phase Space," *Physica D* (April 1989), in press.

J. M. Nese & J. A. Dutton, "Calculated attractor dimensions for low-order spectral models," *J. Atmos. Sci.* **44**, 1950 (1987).

An NSF Graduate Fellowship and Pittsburgh Supercomputing Center grant supported this research.

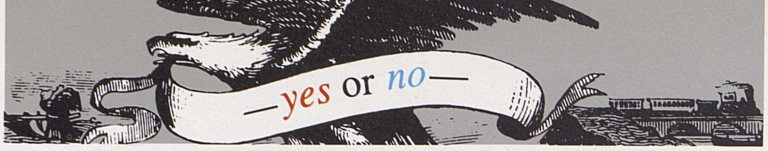

Unrolling 200 Years of Congress

Issues, Agendas and Constituency in Congressional Roll-Call Voting

Keith T. Poole & Howard Rosenthal, Carnegie Mellon University

A Dynamic Map of History

To capture the results of 11 million voting decisions over 200 years of Congress in a concrete, visually accessible format, creating a dynamic map of major issues in American history—this is what political economists Keith Poole and Howard Rosenthal have set out to do with video animations they are producing at the Pittsburgh Supercomputing Center. Their raw material is the complete roll-call voting records of Congress from 1789 to 1989. "Think of each member of Congress as a point on a map," says Rosenthal, "and the purpose of the map is to study historical change. You have to have some way to link these points over the course of 200 years."

Poole and Rosenthal have found the linkage in a relatively simple, two-dimensional spatial model. "Individuals come to Congress with a position that will predict how they look at a whole host of issues, whether defense policy, gender policy, income tax or whatever." The horizontal dimension classifies voting by general perception of "liberals" versus "conservatives" on predominantly economic issues—conservatives right of center, liberals left. The second dimension, projected on the vertical axis, reflects liberal versus conservative stances on social issues—conservatives above and liberals below center. With only these two dimensions, the model correctly represents over 85% of the 11 million individual votes involved.

"We believe the first dimension captures basic philosophical differences between the two major parties," says Poole. "In recent times especially, party differences have tended to revolve around economic issues." Farthest "right" are Republicans who have opposed government intervention in the economy and income redistribution to the poor and organized labor. "The second dimension seems to pick regional differences," continues Poole, "and those translate, prior to the Civil War, to conflicts over race. After World War II, again they're civil rights conflicts. This second dimension also picks up differences on issues like abortion and Vietnam."

How It Works: Policy Space & Glue

"Basically the model represents the two choices on each vote—*yes* or *no*—as points in a kind of policy space," says Poole. "Individuals are points in the same policy space, and for each policy outcome, they vote according to the closest choice to them. It's literally a geometric model." A *maximum likelihood* algorithm picks the policy and individual points to maximize the probability of the observed voting pattern. "The glue that holds it together over 200 years," explains Rosenthal, "is overlapping terms of service. It's like a moving picture with a series of images that change slowly. Few people now in the Senate were there in 1962 when Ted Kennedy was first elected, yet Kennedy has hardly changed position. He's part of the glue for senators before and after him. If we replaced Congress completely every two years, this whole project would be meaningless."

Seeing is Learning: Video as History

"We have about 30 to 40 minutes of video now," says Rosenthal. "I showed it to an undergraduate class, and I'm sure students will tell you it was the most interesting class of the semester. On monetary history in the 19th century, for example, I was able to show directly how votes on free silver cut across parties—dividing along the vertical dimension in a very clean way." The video also shows that pre-Civil War political parties divide more or less as usual on economic issues while on the social dimension the split is literally north and south; party loses meaning.

"You can show that basically there's always an economic redistribution conflict coexisting with whatever is a critical disturbing issue at the time." Until the Civil War, the disturbing issue was slavery. From then until the Depression, it was a conflict over egalitarianism in industrial areas. After the Depression, explains Rosenthal, the coalition Roosevelt formed to resolve that issue became destabilizing itself because it reintroduced race as an issue. In the 50s and 60s, the vertical dimension became civil rights. "Actually, we've found something quite different from the usual historical story that the Depression produced a tremendous political realignment."

Keith Poole & Howard Rosenthal

Video as Research Tool

"With the video," continues Rosenthal, "we can see that the Depression didn't produce a realignment so much as it just shifted the majority." Despite a tremendous influx of Democrats, party voting patterns stayed basically the same. In the 40s, however, votes on voting rights for black members of the armed forces show southern Democrats moving up on the vertical scale. "You start to get this three-party system with southern Democrats as a separate party, and a definite realignment in policy becomes apparent. It was a consequence of the Depression but didn't happen until 15 years later."

The video enabled Poole and Rosenthal to make other observations obscured by the massiveness of the data. "Once a person is elected," says Poole, "he or she doesn't change. He's carried feet first out of the chamber with his ideological boots on." This is one reason why sweeping changes in policy are likely to occur only by massive electoral replacement. The video also shows a gradual implosion over time in the degree of policy difference among legislators. Before 1890, individual points tend to fill the computer screen, indicating a prevalent lack of consensus. In the 20th century, the positions draw inward, and the parties move closer together.

Poole and Rosenthal's animations also represent voting by regional analysis, a view that first enabled them to interpret the second dimension as social issues. They have recently included images that show political action committee campaign donations related to legislators on the policy map. In future research, they plan to study the effect of agenda-setting procedures on vote outcome. One of their ultimate goals is to refine their videos with advanced graphics and voice-over and to make them available as an educational supplement.

References:

K. Poole & H. Rosenthal, "Patterns of Congressional Voting," GSIA Working Paper # 88-89-07, Pittsburgh: Carnegie Mellon University.

K. Poole & H. Rosenthal, "The Spatial Stability of Congressional Voting: 1789-1985," 1988 Annual Meeting of the American Political Science Association, Washington, DC, September 1-4, 1988.

This research was supported by NSF grant SES-8310390. Coordinate estimates that formed the basis for the videos were computed at the John Von Neumann Center at Princeton University.

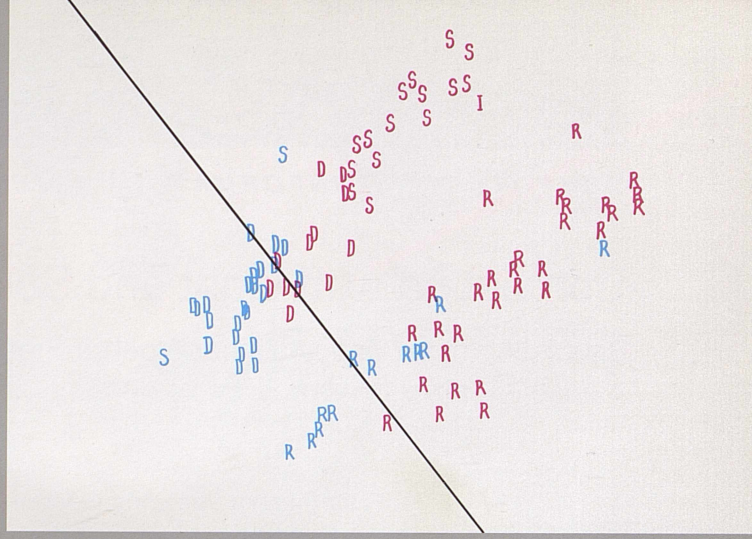

This "map" represents a single roll-call vote, the critical vote in the Senate ratifying SALT I—called the "Jackson Amendment" after its proponent, Senator Henry Jackson (D-WA). D represents a non-southern Democrat, S a southern Democrat and R a Republican. Red indicates a "yea" and blue a "nay." The roll-call "cutting line" separates "yeas" from "nays" according to the maximum likelihood that senators on the "yea" side would vote "yea" and vice-versa. Voting "errors" are senators on the opposite side of the cutting line from their vote. Errors should generally occur close to the cutting line.

"A coalition of moderate and conservative Democrats and Republicans who wanted to go on record as being strong on Soviet policy supported the Amendment," says Poole. "There are few errors and those that occur are close to the cutting line. This shows that the model can account for foreign policy hawks and doves as a combination of the economic and social dimensions. Doves occupy the southwest quadrant."

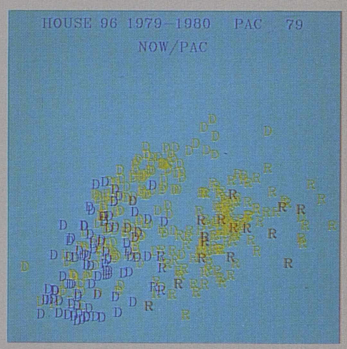

These maps represent contributions by three political action committees (PACs) during the 1979-80 House election campaign. ("NOW" is the National Organization of Women.) D, R and S represent party: Democrat, Republican and southern Democrat. Blue means the PAC gave money to the representative. Red means the PAC gave money to the representative's opponent. Pale yellow means the PAC made no donation.

"PAC contributions tend to be spatially based," says Poole. "Labor unions give to liberals and corporations to conservatives. Most PACs support incumbents and don't spend on challengers; corporate PACs, for instance, often give to incumbents widely scattered in both dimensions except for very liberal Democrats and generally don't spend money to defeat anyone. Labor unions, on the other hand, tend to spend for liberal incumbents and against conservatives."

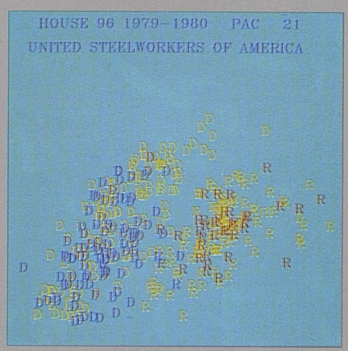

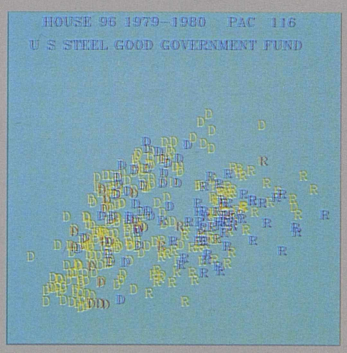

Computing with Light and Sound

Simulation of an Optical Linear Algebra Processor

Bradley K. Taylor, Caroline J. Perlee & David P. Casasent, Carnegie Mellon University

Faster than a Speeding Electron

Slow, clumsy electrons—that's how we'll think about electronics when optical computing makes the leap to realized technology. Because of their speed and focused energy, laser beams have the potential to process data 100 times faster than the fastest semiconducting devices now in use. An optical transistor could switch states in about a picosecond, a trillionth of a second. Since about the mid-70s, researchers—including Bell Labs—have been investigating this potential as a logical step beyond current technology.

Overcoming the technical difficulties of making components for an all-optical digital computer is likely to require many years and research dollars. "You have to reduce the size and power dissipation each by a factor of about a thousand, and you'll be a grandfather by then," says David Casasent, director of Carnegie Mellon's Center for Excellence in Optical Data Processing. Rather than an all-optical system, Casasent has set his sights on a reliable Optical Linear Algebra Processor (OLAP), a project well within the limits of current technology.

Off the Drawing Board

Casasent's group has built two laboratory OLAP systems and has extensively tested and evaluated the second one. "We wanted to prove that you could really do it in the lab instead of just saying you could." The literature abounds with proposals for various optical processors, but the OLAP is to Casasent's knowledge "the only acousto-optic OLAP ever built and tested." Laboratory results and simulations on the CRAY X-MP show that an advanced version of Casasent's OLAP architecture can do more than a billion matrix-vector operations per second, equivalent to state-of-the-art supercomputing. "And it would not cost 10 million dollars," says Casasent, understating the highly competitive economics of optical processing (Fig. 1).

The OLAP is not a general-purpose system, but as Casasent notes, if you needed to choose a "core processor" for scientific and engineering applications, it would be a linear algebra processor. Nearly all computations involved in solving differential equations use matrix-vector techniques. He views the OLAP as a peripheral that would enhance matrix-vector operations on a general-purpose system. "The OLAP is a special path," says Caroline Perlee, who with Casasent advising wrote her thesis on it. "Every time you do a matrix-vector operation, you route the data to the OLAP."

Parallel Light and Acousto-Optics

The OLAP relies on the unique parallelism of laser light as well as its speed. Multiple parallel focused beams won't scatter and mix like electron currents that pass too near each other in a semiconducting material. Economically, this is one of the main advantages of optical technology. The OLAP uses it to do arithmetic in the multiple-bit *word* format with basically one computing cycle per word operand, rather than serially bit-by-bit.

The OLAP accomplishes this word parallelism through the interaction of sound and light in an acousto-optic cell (AO). Each of several AO *channels* has a lithium niobate crystal for each bit-position of a word. The AO converts electronic input to sound waves and transmits them to the crystals. The sound vibrations alter each crystal's light refracting properties, and parallel beams from a laser diode then illuminate the channel. "By varying the strength of the input signal," says Perlee, "we change the refractive index of the cell and vary the amount of light out the other end, and in short it really does a multiply—the product of the light intensity from the diode and the signal strength input."

Supercomputer Simulations

Casasent, Perlee and Brad Taylor, who also did his thesis on the OLAP, have used the CRAY X-MP to simulate OLAP performance. They ran several case studies—computational fluid dynamics and finite element earthquake calculations to compare simulated results against the lab prototype. These studies also evaluated performance with several algorithms for solving systems of linear algebraic equations. The simulations modeled nine potential OLAP error sources, such as circuit noise and fluctuations of light output from the laser diodes.

Supercomputing capability enabled Casasent's group to run over 100 different simulations to complete their error source studies. Processing time of 24 to 36 hours per run on a VAX 11/750

was reduced to a matter of minutes. By verifying simulation results against prototype tests, Casasent established that the error source models accurately represent OLAP performance. The simulations identified problems in the detector, which he subsequently corrected. "With the CRAY," says Casasent, "we were able to do enough of the modeling to do a good job. It's sort of unusual—using a supercomputer to simulate an optical one, but it's a very nice marriage."

References:

Caroline J. Perlee & David P. Casasent, "Optical Matrix-Vector Processing for Computational Fluid Dynamics," Proc., Soc. Photo-Optical Instrumentation Engineers 936 (1988).

B. K. Taylor & David P. Casasent, "Optical Laboratory Solution and Error Model Simulation of a Linear Time-Varying Finite Element Equation," Proc., Soc. Photo-Optical Instrumentation Engineers 977 (1988).

NSF grant ECS-8500830, NASA Langley Research Center grant NAG-1-575 and AFOSR grant 84-0239 supported this research.

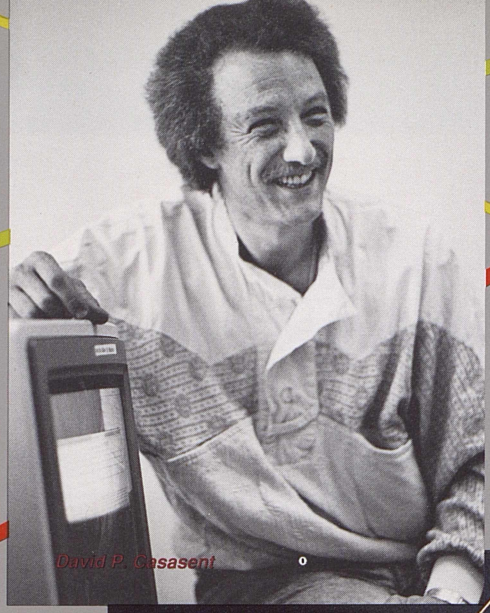

David P. Casasent

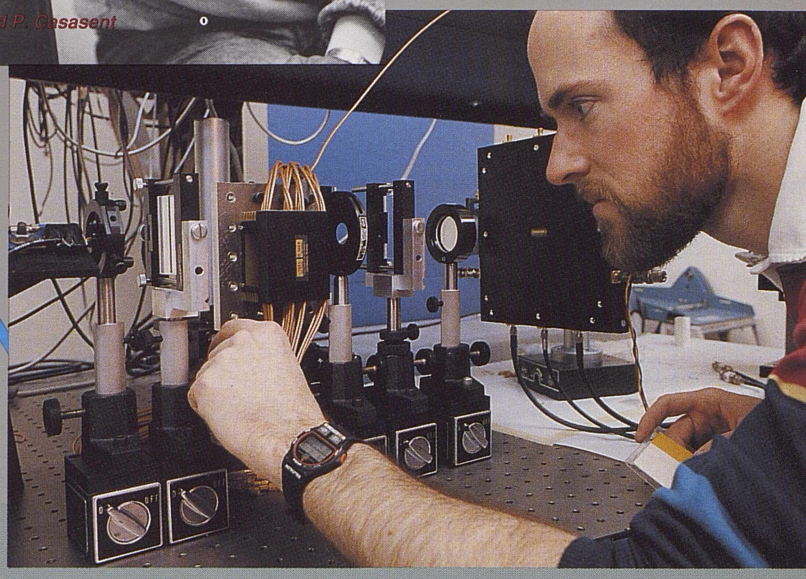

Ed Baranoski, a Ph.D. candidate in Electrical and Computer Engineering, aligns the OLAP.

Figure 1. Optical versus Digital Performance-Ratio as a Function of Component Speed

The graph compares optical to digital performance for 12-bit devices performing 12, 24 and 36 bit multiplications as a function of individual component speed in megahertz (mHz), millions of cycles per second. The comparison is based on the A/D converter in the optical processor versus a digital multiplier chip. As component speed increases linearly, the parallelism of optical processing increases optical performance exponentially.

"The more important criteria, rather than sheer speed, should be operations per dollar or per watt. Ed Baranoski and I compared digital technology to optical and matched the components in terms of what digital could do and what optics could do. So you say where's the crossover point? Basically it's when digital on a multiply gets down to a microsecond, which is where they are now, about 1.2 mHz. So it's 50-50, both are the same, and nobody's going to use this new thing if it's the same as the old. As a general rule, you want it to be 10 times better. Digital is talking about going up to a couple hundred mHz. And when they're there, given any way you want to do the operations, optics is orders of magnitude better."

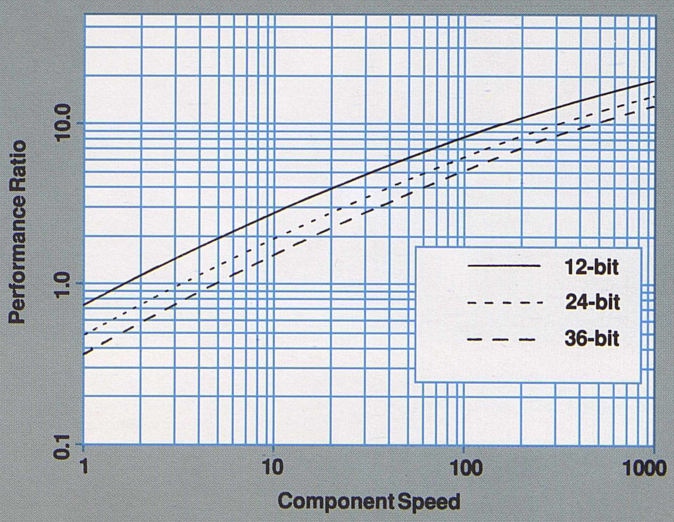

The Academic Affiliates of the Pittsburgh Supercomputing Center

Carnegie Mellon University
Case Western Reserve University
Drexel University
Duke University
University of Florida
Georgia Institute of Technology
Johns Hopkins University
Lehigh University
University of Maryland
University of Michigan
Michigan State University
North Carolina State University
University of North Carolina
Northwestern University
Ohio State University
University of Pennsylvania
Pennsylvania State University
University of Pittsburgh
Rensselaer Polytechnic Institute
University of Rochester
Temple University
University of Tennessee
University of Virginia
Washington University
West Virginia University
University of Wisconsin
Yale University

Credits

Editor:
Mike Schneider
Pittsburgh Supercomputing Center

Printing:
Broudy Printing, Inc.
223 Auburn St.
Pittsburgh, PA 15206

Design/Layout:
Steve Palmer
Dawn Kostilnik
Tim Kostilnik
Pittsburgh

Production Coordination:
Vivian Benton
Pittsburgh Supercomputing Center

Writing:
Mike Schneider
Liz Johnson (12, 22), Kelly Spencer (32, 42), Janet Vietmeier (44)
Thespine Kavoulakis (16, 18)

Illustrations:
Steve Palmer

Fred Carlson, fine arts professor at Carnegie Mellon, incorporated this project into his Winter 1989 Advanced Illustration course, facilitating an interchange between science and art that resulted in the following illustrations: Jen Hebert (32), Norma Pribadi (20), Carin Reeve (16), Brian Seifirt (22) & others (Shana Chandler, Seifirt, Trish Wilde) that will appear in a future edition of this book.

Ellen Elmes, detail from "Autumnweave" (34).

Photographs:
Ken Andreyo, Carnegie Mellon (3, 4, 5, 6, 45)
David Aschkenas, Pittsburgh (5, 6, 14)
Melissa Farlow, Pittsburgh (7)
Dave Mengle, Penn State (41)
Larry Rippel, Westinghouse (7)
Jim Schafer, Pittsburgh (1)
Jerry Voron, Westinghouse (5)
Kulicke and Soffa Industries, Inc. (24)
Motorola MOS Integrated Circuits Division (16)
IBM Research Division (18)

Production Assistance:
David Brock, Kip Fiebig, Thespine Kavoulakis

Computer Images:
Joel Welling & Matt Tolbert
Pittsburgh Supercomputing Center
Mario Peruggia (34), Carnegie Mellon, Statistics Dept.
Scott Berger (34), Cornell University
Richard Miles & Greg Russell (30), Princeton University

Science Consulting:
Martin Head-Gordon & Teresa Head-Gordon,
Carnegie Mellon Department of Chemistry
Phil Andrews, David Deerfield, Doug Fox, Mike Kopko, Joel Welling
Pittsburgh Supercomputing Center

Thanks to the Mellon Institute Library staff for help with research.

We extend thanks and deep appreciation to the scientists and engineers without whose work this publication would not be possible.

The Pittsburgh Supercomputing Center, a joint project of Carnegie Mellon University and the University of Pittsburgh, together with Westinghouse Electric Corporation, receives significant support from:

The National Science Foundation
The Commonwealth of Pennsylvania
The National Institutes of Health
The Ben Franklin Technology Center of Western Pennsylvania
Cray Research, Incorporated
Digital Equipment Corporation
International Business Machines Corporation
Ardent Computer Corporation

Mike Levine, Beverly Clayton, Jim Kasdorf & Ralph Roskies enjoy the dedication ceremonies for the Pittsburgh Supercomputing Center's CRAY Y-MP/832. Along with being the world's fastest computer, the Y is a serviceable bench.

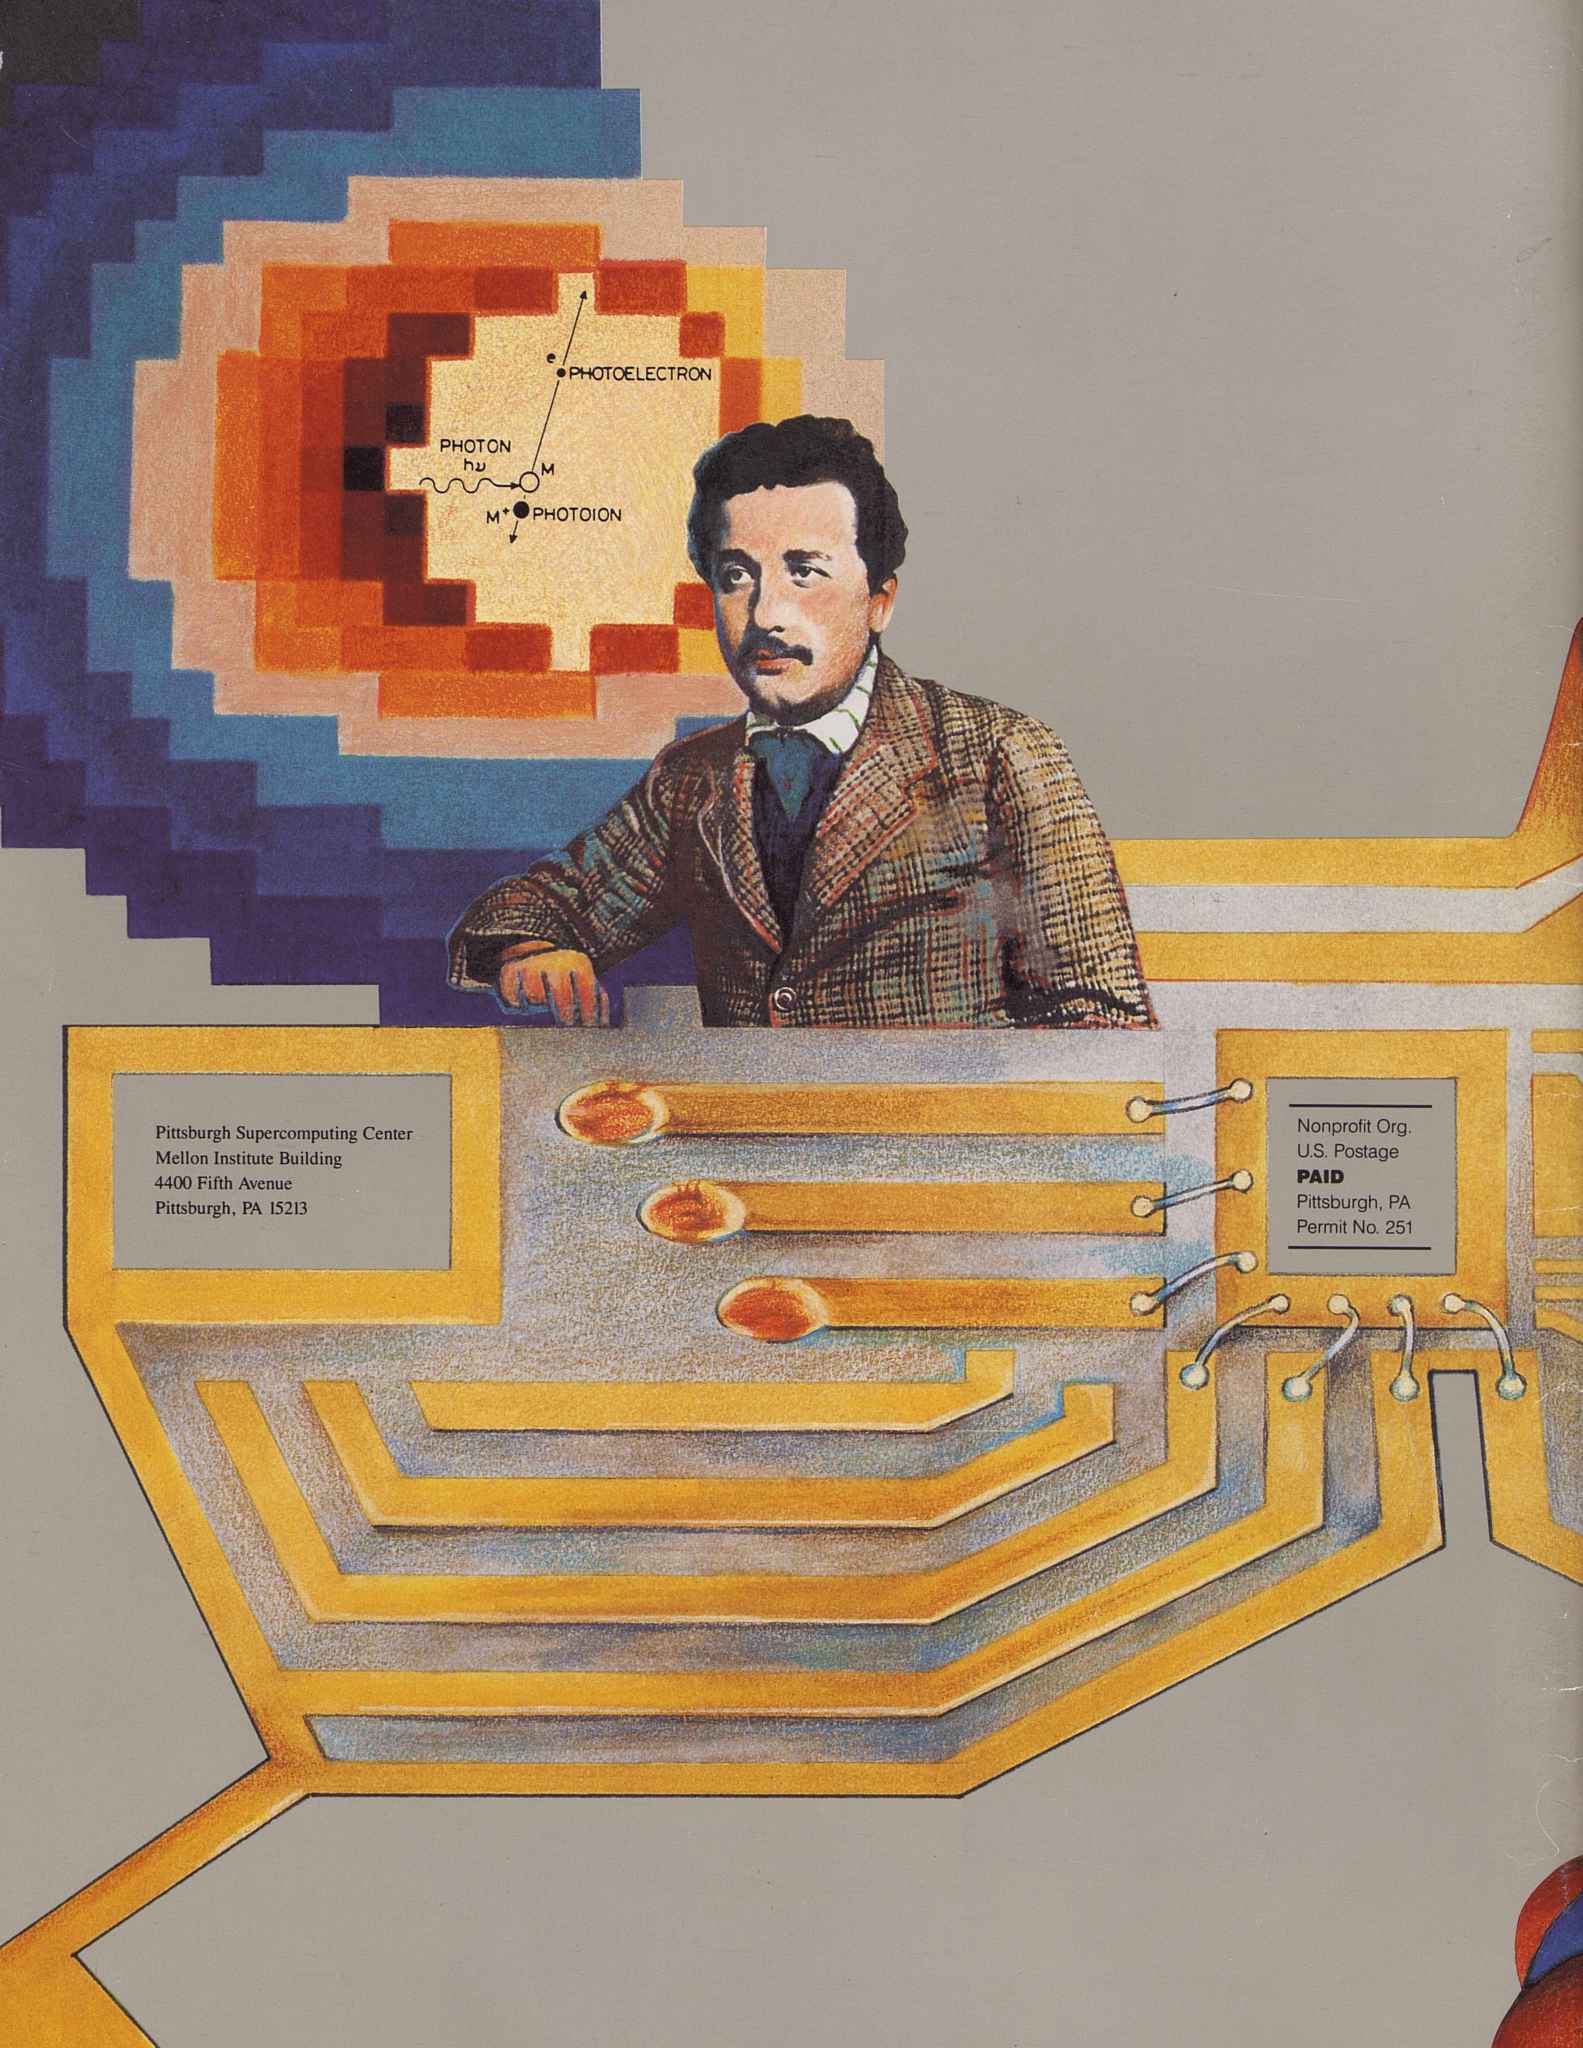

Pittsburgh Supercomputing Center
Mellon Institute Building
4400 Fifth Avenue
Pittsburgh, PA 15213