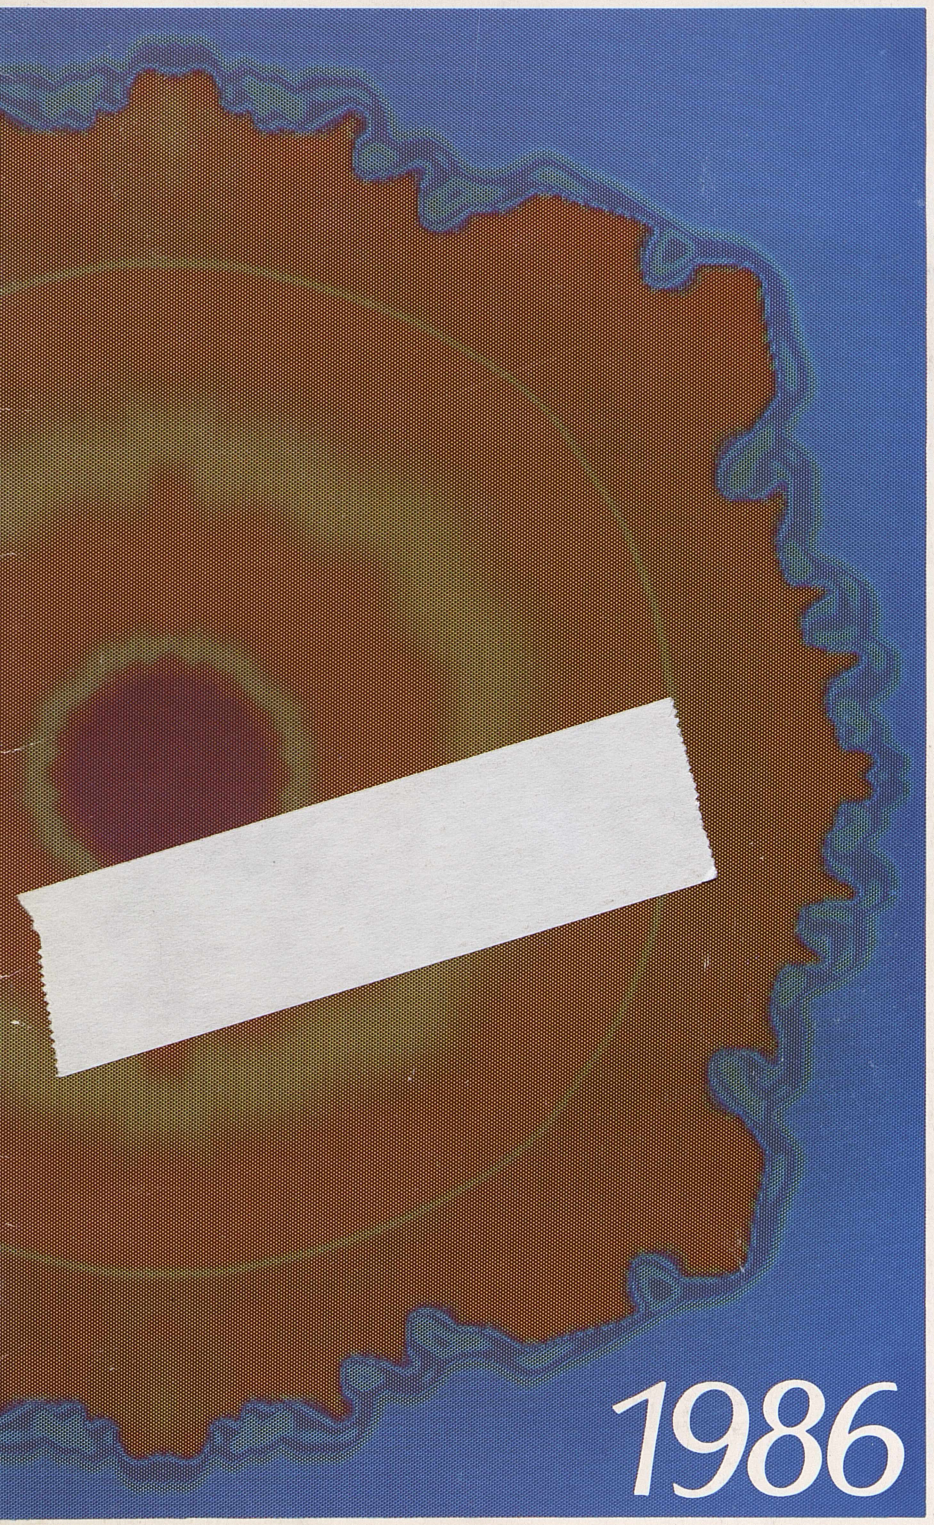

Projects in Scientific Computing

1986 1987

PITTSBURGH SUPERCOMPUTING CENTER

On the cover:

A study currently underway at the Pittsburgh Supercomputing Center is modeling the "wind" emitted by Seyfert galaxies. The red, high-temperature ($T \approx 10^8 K$) gas in the center travels outward at hypersonic speeds ($v \approx 1000$ km/s) into a dilute, cooler ($T \approx 10^4 K$) interstellar medium shown in blue. The high speed of the outflow creates a shockwave and turbulence at the boundary between the two media.

This study is being conducted by Allen V.R. Schiano and Arthur M. Wolfe of the University of Pittsburgh, Robert F. Carswell of the Institute of Astronomy, Cambridge University, and Ray J. Weymann of Mount Wilson and Las Campanas Observatories.

Projects in Scientific Computing

PITTSBURGH SUPERCOMPUTING CENTER

Advanced Computing for Engineering and the Sciences

Projects in Scientific Computing

Foreword

Supercomputers are the fastest available general-purpose scientific computers. They are needed whenever massive computation stands between what we know and what we want to know. For example, we know the equations describing the flow of air past an aircraft, but the computational task of wing design requires a supercomputer. Rational drug design based on our understanding of chemistry at the molecular level may soon be possible with the aid of supercomputers.

To think of supercomputers only in terms of increased speed is to miss a vital point. Their increased speed actually transforms the types of problems on which a scientist is willing to work. Decreasing computation time from weeks to hours enables a scientist or engineer to ask much bolder questions, to try more daring and novel approaches, to include more realistic complexities, and even to replace costly and time-consuming experiments by theoretical calculation.

The main mission of the Pittsburgh Supercomputing Center is to facilitate better science through supercomputing. This booklet contains just a sampling of the fascinating scientific projects being undertaken at our Center. The projects range from immediately practical strategies for air pollution control to abstract studies of elementary particles; from the small scale of atoms and particles to the immense scale of galaxy formation. Most of these studies would not have been undertaken without supercomputers. New questions are being asked which were impossible to tackle even a few years ago. The results will change intellectual disciplines and even society. Obtaining such diverse and important results even before the end of our first full year of operation is truly rewarding.

Michael J. Levine, Scientific Director

Ralph Roskies, Scientific Director

Table of Contents

5 PHYSICS

6 ASTRONOMY

The Pittsburgh Supercomputing Center

Beverly Clayton, Executive Director of the Pittsburgh Supercomputing Center, came to the PSC from Gulf Oil Corporation where she was Director of Systems Support. She has 21 years of experience in the data processing field, including management of technical personnel, coordination of major projects, technical planning and consulting, operating systems support, and scientific programming.

The Pittsburgh Supercomputing Center (PSC) is one of five national supercomputing centers established by the National Science Foundation in response to an acute need of national research communities for improved access to the best computing technology. The PSC is a unique collaboration of Carnegie Mellon University, a major private university; the University of Pittsburgh, a major state-related university; and Westinghouse Electric Corporation, a major corporation. It seeks to combine the innovative character of universities with the performance standards and support services of successful industrial computing centers to provide high performance computing to the national research community.

To take fullest advantage of the intellectual environment of both campuses, most of the staff is housed in the Mellon Institute Building, which is part of Carnegie Mellon University and located adjacent to the University of Pittsburgh campus. The beautifully-renovated office quarters of the PSC include a training center and facilities for visiting scientists.

The PSC has assembled a highly-qualified and enthusiastic staff, drawn from academia and from industry. The senior managers have years of industrial computing experience. Several staff members have substantial scientific research experience beyond the PhD level. Together, they form a dedicated team, assisting scientists and engineers in achieving high productivity in research through effective use of supercomputers.

The supercomputer itself and some of its ancillary equipment are located in a new 10,000 square foot machine room at the Westinghouse Energy Center in Monroeville, PA, 13 miles from the campuses. The Westinghouse site and the Mellon Institute offices are linked by a high speed electronic network. Supercomputer operations and facilities management are handled by the Westinghouse Electric Corporation, which has a distinguished history of providing supercomputing services.

In addition to funding from the National Science Foundation, the PSC is supported by grants from the Commonwealth of Pennsylvania and the Ben Franklin Partnership. It has formed research partnerships with Cray Research, Inc. to adapt and develop scientific applications for supercomputers and with Digital Equipment Corporation to improve the computing environments that support supercomputers. It continues to form industrial partnerships with major corporations to improve their competitiveness through supercomputing.

The PSC is a truly national center. Before the end of its first year of operation, it already had more than 800 users from more than 85 institutions in 33 states. A total of 25 major research universities have formed an affiliation with the PSC to advise it on policy matters, and to promote supercomputing within their own scientific and engineering communities. (See the Academic Affiliates listing on page 44).

Since most PSC users are not located in Pittsburgh, access to the supercomputer is primarily via national data networks. The PSC is linked to the other NSF national centers by a high-speed backbone network, NSFnet, and is in the process of developing a regional network, which will operate at even higher speeds. The PSC is also accessible through more well-established networks like ARPAnet and BITnet, and through dial-up facilities and GTE Telenet.

Local university users can exploit highly-developed campus networks to access the supercomputer conveniently from their offices. Out-of-town visitors are welcome to use the facilities at the Mellon Institute site or they may prefer to visit their campus colleagues and make use of the campus networks.

Education is an essential part of the PSC's mission. Consultants are available by telephone or electronic mail to answer questions. A monthly newsletter informs users of new developments at the PSC and provides helpful hints on how to use software most effectively. The PSC presents frequent workshops to academic and Industrial Affiliate users to aid them in making best use of the supercomputer's power.

The goal of all PSC activities is the advancement of science and engineering through supercomputing. By drawing on the resources of the national research community and the stimulating environments at the University of Pittsburgh and Carnegie Mellon University, the PSC strives to facilitate research achievements which will fulfill the vision of the National Science Foundation and of Congress in establishing the national supercomputing centers.

James Kasdorf, Engineering Computer Services Manager of the Energy Systems Division of Westinghouse Electric Corporation, is responsible for facilities management and the operation of the PSC Cray X-MP and its peripherals. He has been involved in large-scale scientific and engineering computing for the past eighteen years, including ten years in computer center management positions.

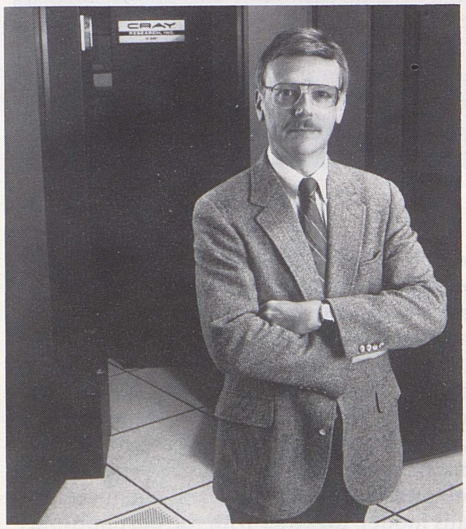

THE COMPUTING FACILITIES

The supercomputer at the PSC is a Cray X-MP/48, one of the world's fastest computers, which can perform 840 million arithmetic operations every second. Since the Cray operates so quickly on large amounts of data, it requires a large memory. The PSC's Cray has eight million words of extremely fast memory and is connected to sixteen DD-49's, the world's fastest disks. Often the computational bottleneck is in getting large datasets into and out of these disks. To circumvent this problem, the PSC's Cray is equipped with a 128 million word SSD (solid state storage device). The SSD can transfer data to the main processors 100 times faster than the disks, and can effectively expand the Cray's memory to 128 million words.

Supercomputer speeds are achieved by a clever combination of fast electronics, extensive use of parallelism, and ingenious packaging that allows the machine to be small while dissipating great amounts of heat. The PSC Cray has four independent processors, each of which has fourteen independent functional units, so that, for example, each processor can be adding and multiplying at the same time. Many of the functional units (like the adder or multiplier) are composed of subunits which themselves operate in parallel. Computer programs can be made to run most efficiently on the Cray through a technique called vectorization, which keeps many of these functional units and subunits calculating simultaneously.

Most researchers do not communicate directly with the supercomputer, but with workstations or "front ends": more conventional computers with friendlier user interfaces. The PSC's front ends are a pair of VAX 8650's running the VMS operating system, which is familiar to a large portion of the research community. VMS, which is also available on dozens of workstations at the Center, provides users with a powerful set of program development tools that are well-interfaced to the Cray. The PSC itself is expanding these tools, to facilitate conversion from the VAX to the Cray environment. UNIX, another operating system familiar in the research community, will soon be available on the PSC Cray. This addition will eliminate the need for mastering another operating system and will allow many researchers to use the same operating system from the workstation on their desks straight through to the Cray. The availability of UNIX will also improve the portability of computer codes from the workstation to the supercomputer.

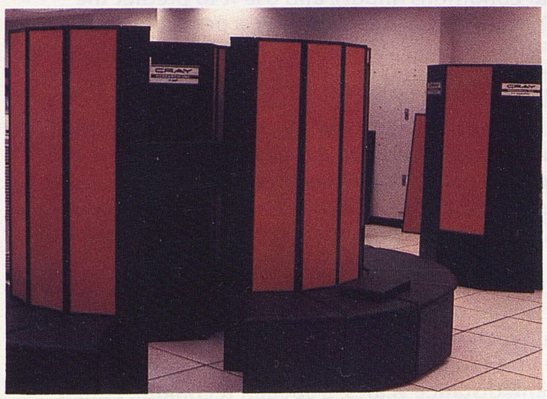

The PSC Cray X-MP/48, connected to the Solid State Storage Device (SSD) (in foreground). The I/O processor (in background) connects the Cray to tape drives, fast disk drives, and to the VAX 8650 front-ends. Black cooling units are located at the base of the Cray, the SSD, and the I/O processor.

1 ATMOSPHERIC SCIENCE

Designing Air Quality Control Strategies For Los Angeles

Gregory J. McRae, Armistead G. Russell, and Jana B. Milford
Carnegie Mellon University

BACKGROUND

Over 30% of the population of the United States is exposed to levels of ozone above federal limits. Nevertheless, current strategies for meeting air quality standards have not been very successful; at least 75 cities are not expected to be in compliance with the federal standards set for the 1987 deadline. Furthermore, implementation of present controls is expensive; the Los Angeles metropolitan area alone, which has one of the most severe air pollution problems in the nation, spends approximately 300 million dollars annually on air quality control. Computer studies that can simulate the outcome of new strategies before they are implemented offer an extremely cost-efficient method of identifying promising directions for air quality planning.

This study undertaken at the Pittsburgh Supercomputing Center has not only indicated why previous air quality control tactics may have failed but has also suggested an alternative strategy that offers the possibility of approaching EPA compliance within the Los Angeles area.

THE QUESTION

A variety of factors must be understood in order to simulate an air pollution problem on a computer. First, one must know the sources of emissions and how they vary with time. This step requires modeling traffic patterns. Second, one must understand the chemical processes that produce the pollution from these primary emissions and appreciate how these chemical processes depend on atmospheric conditions such as temperature, pressure, and intensity of sunlight. Finally, one must understand the meteorology responsible for the transport of the pollution across the affected area. In the computer, all of these factors are interrelated by a "mathematical model," a huge set of extremely complicated, but fairly well-understood, physical and chemical formulae. The computer must take this model, together with the best information available for the physical conditions (input), and use it to predict the level of pollution in various geographical locations over a period of time.

COMPUTATIONAL TECHNIQUES

In the simulations carried out using the Cray X-MP/48 at the PSC, more than 500,000 equations were required to characterize the evolution of ozone and 30 other chemicals produced in the air over Los Angeles. About 70 test cases were studied, each corresponding to different environmental conditions and each requiring approximately an hour of continuous running time on the Cray.

The equations that must be solved are nonlinear and stiff. Chemical reactions of importance occur on time scales of nanoseconds but the processes must be tracked over several hours: a dynamic range in time of 10^{12}. Similarly, the dynamic range in distance is also large, about 10^5. To simulate processes that involve such disparate time and distance scales, a special technique called "operator splitting" must be used. Using this technique, the whole study was completed in a few weeks on the Cray, whereas several years would have been required in a more conventional computing environment. This project would not have been undertaken without the availability of supercomputers.

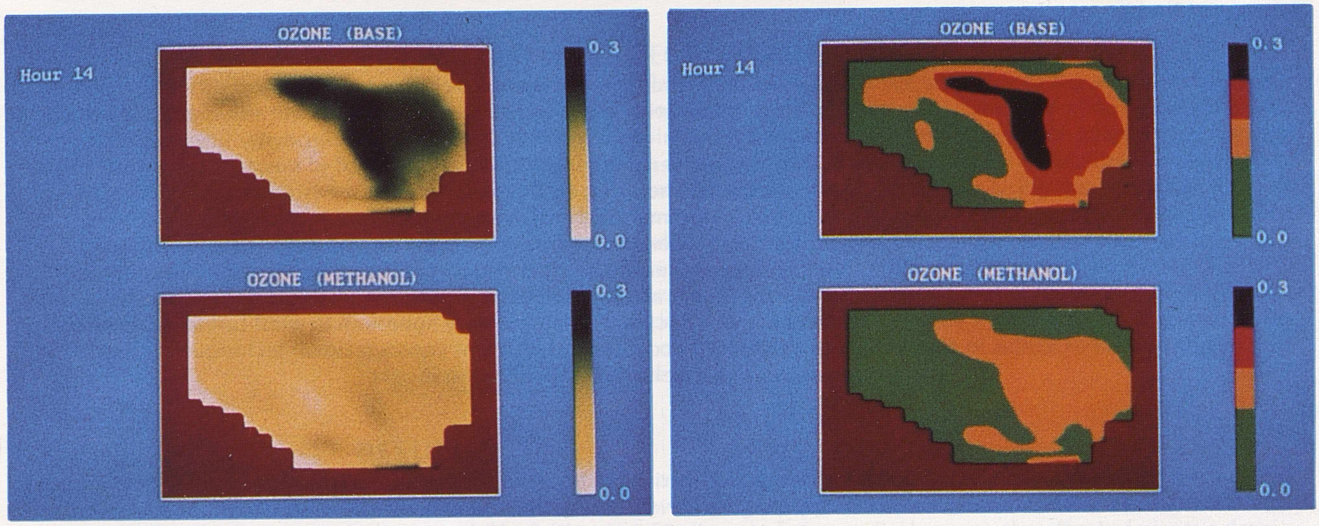

The predicted level of ozone (in parts per million) over the Los Angeles area is shown assuming the current use of fuels in the automotive fleet (top) and assuming a replacement of those fuels by methanol (bottom). The simulation was "frozen" at a time corresponding to 2:00 pm, the time at which the advantage afforded by the use of methanol is most striking. The set of diagrams on the left uses a continuous color scale to represent the ozone level; the set on the right uses a discrete color scale with green indicating areas that are within current EPA guidelines.

RESULTS

The results of the study were both startling and encouraging. To understand them, one must realize that the primary emissions from automobiles, refineries, and power plants include both hydrocarbons and nitrogen oxides. Together, these two types of emission interact with sunlight to produce photochemical oxidants, more commonly known as "smog." Ozone is one such photochemical oxidant. Current EPA regulations have focused primarily on hydrocarbon control as a means of reducing oxidant levels.

Three surprising conclusions were drawn from the study.

- Contrary to what had generally been assumed, it was found that the amount of ozone in the atmosphere does not necessarily decrease if particular emissions are reduced. The exact relationship between the amount of pollution and the level of emissions was shown to be so complicated that, in some cases, strategies for reducing particular pollutants could actually cause air quality to deteriorate in outlying areas. These results are generally applicable to any urban area.

- Because of the complicated interaction between nitrogen oxides and hydrocarbon emissions in the production of pollution, the EPA's concentration on hydrocarbon control alone is inadequate. To achieve air quality compliance, both types of emission must be regulated in a way that reflects their chemical interplay and their impact on other harmful but nonregulated pollutants.

- The PSC project offers hope for the future. If an alternate fuel like methanol were substituted for gasoline in a significant fraction of Los Angeles vehicles, substantial reductions in oxidant and particulate pollution could be achieved. The conversion to use of methanol alone could bring Los Angeles close to compliance with current EPA standards.

Aerodynamics of a Transonic Projectile

Chen-Chi Hsu, N.-H. Shiau and C.W. Reed
University of Florida

BACKGROUND

The ability to predict accurately aerodynamic forces and other flow characteristics is essential to more effective designs of flight vehicles and their control and propulsion systems. To study realistic aerodynamics problems, wind-tunnel experiments are traditionally performed to measure these flow characteristics. With the rising cost of experimental measurements, however, it is becoming extremely expensive to conduct parametric studies in a wind-tunnel. Moreover, each test facility has a limited range of application and, consequently, certain important flow conditions often cannot be simulated. With the recent advent of supercomputers, numerical simulations have become an effective complementary approach to wind-tunnel experiments. Nevertheless, accurate and efficient computational techniques to predict aerodynamic behavior are still being sought.

THE QUESTION

The main objectives of this supercomputing project are to study techniques in applying numerical simulation to complex aerodynamic problems and to assess the relative importance of the viscous drag compared to the pressure drag acting on a projectile traveling at nearly the speed of sound (transonic projectile). Different curvilinear grid networks representing the flow field were provided to computer programs that solve the governing equations for high-speed, compressible, turbulent flows. For the particular shape of the projectile studied here ("secant-ogive-cylinder-boattail" projectile shown below), experimental surface pressure measurements are available to test the accuracy of the numerical results. Previous simulations provided results that were not always in agreement with these measurements, especially over the cylindrical portion of the projectile. This study was able to show the importance of viscous flow computation for accurate aerodynamic force prediction.

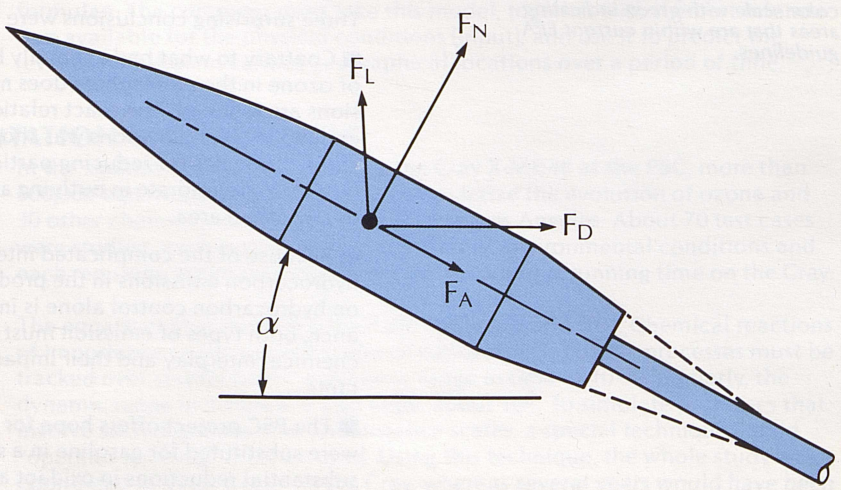

The secant-ogive-cylinder-boattail projectile is shown here at an attack angle of α. The cylinder-boattail juncture causes turbulence in the air flow.

COMPUTATIONAL TECHNIQUES

Numerical simulations indicated that the accuracy of the solutions depends on the grid network provided to the program. Since the choice of a good grid network is not trivial for a complex flow problem, a self-adaptive grid generation technique is now being developed to concentrate computer resources on the most complex portion of the turbulent flow. Using this technique, the program is able to recognize where the flow is becoming most complex and then respond by introducing a finer-meshed grid in that region to reflect more accurately the complex detail in the flow characteristics. The amount of computer time required for a simulation depends on the number of grid points used in the computation. As an example, even with the speed of a supercomputer, one study of a projectile at a 45° angle of attack required two full hours of Cray CPU time to obtain a converged solution.

RESULTS

The results of the PSC study showed that viscous drag is as important as pressure drag for flow past projectiles at zero angle of attack at speeds of 0.91 Mach (that is, 0.91 the speed of sound). However, the effect of viscous drag decreases with increasing Mach number and with increasing angle of attack. For a projectile model at a 10° angle of attack, the numerical results showed that viscous drag is negligible. The figure below shows, in cross section, a sampling of air particle trajectories a short distance downstream of the cylinder-boattail juncture for a projectile with a 45° angle of attack. Regions of "reversed flow," which serve to reduce the effect of drag, can be seen near the surface of the projectile.

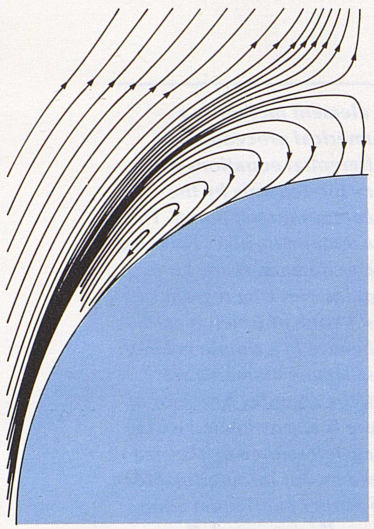

A cross section of the projectile downstream of the cylinder-boattail juncture. The direction of the surrounding air flow is indicated with arrows.

Access to the Cray X-MP/48 at the Pittsburgh Supercomputing Center has enabled this research to contribute significantly to the basic understanding of how to simulate complex transonic flow by numerically solving hydrodynamic equations.

Analysis of Sheet Metal Forming

R.H. Wagoner, et. al.
Ohio State University

BACKGROUND

Until recently, few analytic tools existed for analyzing sheet metal forming operations except for approximate analytic methods based on one-dimensional reductions of geometry. A numerical technique known as the finite element method (FEM) allows, in principle, the solution of problems of almost any complexity, with the only limitations being available computer accuracy and speed. For nonlinear, incremental problems, such as ones involving large-strain plasticity, these limitations restrict practical analysis on machines like the DEC VAX 11/780 to two-dimensional, isothermal, and time-insensitive (static) approaches. Even with these limitations, however, great strides have been made in the last ten years in interpreting the role of material properties on forming behavior. The research described here is aimed at removing the previous barriers by creating more complex models that are suitable for supercomputing.

The potential payoff of a practical sheet-forming program is enormous. As an example, the production of dies represents a large component of the lead time between design and production of a new car model. Reduction of this time and of the personnel time needed for experimental die testing would allow manufacturers to respond more quickly and economically to rapidly-changing world market conditions. This efficiency would, in turn, improve international competitiveness. In addition, the availability of such analytical capability would improve manufacturing quality through better and faster communication between designers and manufacturing engineers prior to production of final dies. With the availability of supercomputers, some of these benefits may be realized.

THE QUESTION

Toward these goals, investigations are underway using the Cray X-MP at the Pittsburgh Supercomputing Center to extend established practice in two critical areas. In the first investigation, two-dimensional FEM programs were written which couple structural and thermal aspects of sheet metal deformation. Preliminary experimental and analytical results indicate that significant temperature changes are generated during forming, necessitating solution of much more complex boundary-value problems. The resulting programs, in turn, necessitate supercomputation for reasonable solution times. The change from a DEC VAX 11/780 to the Cray X-MP increased computing speed by a factor of 30, and 60% vectorization of the code led to an additional factor of nearly three. The non-isothermal simulations now typically take five Cray CPU minutes, as contrasted to six or seven hours on the VAX. Results of these simulations have explained the basis of improved forming operations proposed by several researchers. In addition, material properties have been identified which bear on the contribution of changes in temperature to failure during forming.

Although the two-dimensional simulations are very helpful for qualitative understanding of forming processes and material behavior, a valid comparison with industrial practice must rely on even more general models. A large effort is currently underway to produce a program capable of simulating an arbitrary stamping operation in three dimensions. Such a program is quite general; nearly all of the structural and exterior parts of an automobile are manufactured by stamping. The three-dimensional nature of the simulations imposes the most severe time constraints yet attempted. Several simplifications are being made, including quasi-two-dimensional mechanics models ("membrane theory") and simple contact and friction boundary conditions ("Coulombic friction"). Best estimates for the CPU time required for the simulation of a simple but general mechanical part exceed one hour on a Cray X-MP. The more sophisticated simulations now in progress would not have been attempted without the aid of a supercomputer of comparable power.

The finite element method is a general numerical procedure for solving differential equations with an accuracy limited only by the available computational power. The core of the method involves breaking a structure or domain of any kind into a set of interlocking regions ("elements"), each of which is assumed to behave in a simple manner. When these simple elements are assembled, the complex behavior of the structure is approximated within an accuracy determined by the size of the elements. For infinitesimal elements, the proper continuum equations are usually recovered. Computationally, the discretization of space results in a large system of linear or nonlinear equations. In the nonlinear case, further discretization into time steps is done; and at each step the system of equations must be solved iteratively.

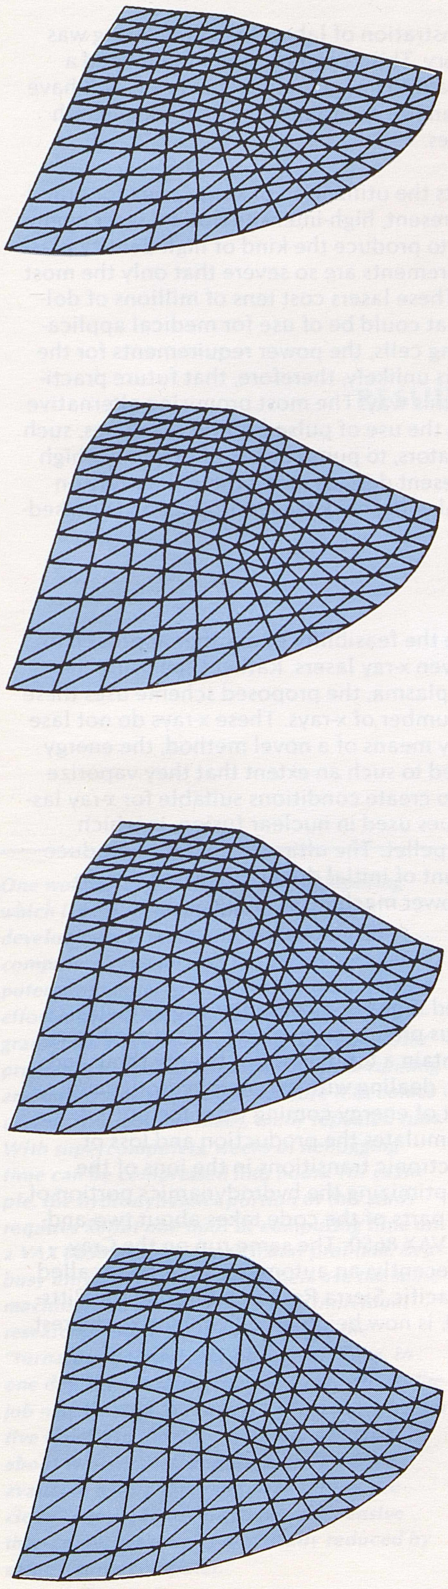

The nonisothermal, two-dimensional work described here was done in conjunction with Yong H. Kim, Muh-Ren Lin, Kavesary S. Raghavan, Yufei Gao, and Thana Ruangsilasingha. The three-dimensional sheet forming work is being undertaken together with Kwansoo Chung, Yves Germain, J.K. Lee, Anne B. Doucet, and Julie R. Knibloe.

Four stages in the three-dimensional deformation of an finite element (FEM) mesh. Since the specimen is symmetric, only one-quarter of the total mesh shown here need be considered in the computation. The problem is to predict deformation patterns of a circular sheet of metal which is indented by a square punch.

Designing More Economical X-ray Lasers

Bernie M. Penetrante
Weber Research Institute
Polytechnic University

BACKGROUND

In 1984, the first unambiguous demonstration of laboratory x-ray lasing was made at Livermore National Laboratory. This accomplishment launched a new era in high technology. It is not hard to imagine that x-ray lasers will have as much, if not more, impact on human life as ordinary, optical wavelength lasers have had in the past few decades.

The one major drawback that prevents the utilization of x-ray lasers for practical applications is their expense. At present, high-intensity optical laser beams are used to pump the power required to produce the kind of high density plasma suitable for x-ray lasing. The requirements are so severe that only the most powerful optical lasers can be used. These lasers cost tens of millions of dollars apiece. To produce x-ray lasing that could be of use for medical applications, such as x-ray holography of living cells, the power requirements for the laser-driver would be even greater. It is unlikely, therefore, that future practical x-ray lasers will be constructed in this way. The most promising alternative to the use of expensive laser-drivers is the use of pulsed-power machines, such as capacitor banks or explosive generators, to pump bursts of extremely high energy directly into the x-ray laser. Present designs along these lines rely on the compression of the lasant material with shock waves produced via pulsed-power.

THE QUESTION

The purpose of this study is to explore the feasibility of another approach to the development of pulsed-power-driven x-ray lasers. Rather than using the shock waves directly to compress the plasma, the proposed scheme uses these shock waves to first produce a large number of x-rays. These x-rays do not lase since they are incoherent. However, by means of a novel method, the energy density of these x-rays can be amplified to such an extent that they vaporize and ionize the lasing target material to create conditions suitable for x-ray lasing. This concept is similar to techniques used in nuclear fusion, in which x-rays are used to compress the target pellet. The ultimate goal is to produce an x-ray laser that uses the least amount of initial driver energy, that is, to produce the least expensive pulsed-power machine possible.

COMPUTATIONAL TECHNIQUES

To study the feasibility of the proposed scheme, a computer program has been developed to simulate the various processes involved. There are three major parts to the program, which contain a total of about twenty thousand lines of FORTRAN code. The first part, dealing with magnetohydrodynamics, simulates the effect of the initial burst of energy coming from the pulsed-power machine. The second section simulates the production and loss of radiation. The third deals with the electronic transitions in the ions of the plasma. Much effort was devoted to optimizing the hydrodynamics portion of the code. A typical run using all three parts of the code takes about two-and-one-half hours of computer time on a VAX 8650. The same run on the Cray X-MP requires about seven minutes. Recently, an automatic vectorizer called FORGE has been made available by Pacific Sierra Research through the Pittsburgh Supercomputing Center. FORGE is now being used to vectorize the rest of the code.

The size of the program made it necessary and convenient to make extensive use of the UPDATE and BUILD utilities of the Cray operating system to manage the numerous subroutines of the program. To facilitate the use of these utilities and provide readable FORTRAN code, interactive and intelligent PC-based software was developed which automatically writes Cray Job Control files, reformats FORTRAN codes, and writes documentation and alphabetical indexes of subroutines and variables. The software, called FORSEE (Fortran Structuring and Editing Environment), was written in Turbo Prolog, an artificial intelligence language for personal computers. With "pop-up" menus and help files, it becomes easy not only to write well-structured and well-documented code but also to use the Cray operating system while logged-in remotely from a PC.

RESULTS

In x-ray lasing projects, experiments are very expensive, and one is forced to rely on computer simulations to an extent not usually encountered elsewhere. The field is still in its infancy; and it is yet to be determined which approach is the best, the cheapest, the simplest, the most efficient or which approach can yield the highest transition energies. This supercomputing project represents an initial attempt in the design of what may constitute the next generation of x-ray lasers. Preliminary results indicate that pulsed-power-driven x-ray lasing can be achieved with presently available pulsed-power machines. Before 1984, one often heard: "X-ray lasers would be exciting if they could be made." Now, one hears: "X-ray lasers would be useful if they could be made better and more economical." Although many more parameter combinations will have to be studied, it is encouraging to see that x-ray lasing experiments could someday be affordable even in universities, where research budgets are not as large as those of national laboratories.

One wonderful fact about supercomputing which is seldom mentioned is that code development is more efficient with supercomputers. Anyone who has ever done computer programming knows that a lot of human effort is spent, not so much in writing the program itself but in debugging it. The computer programmer spends considerable time running and debugging the program before it becomes useful. This requires many, many repeated runs. With supercomputers, weeks of debugging time can be compressed into hours. For example, the hydrodynamics portion of this study requires about one hour of computing time on a VAX 8650 to study a particular problem. In a busy university, many researchers use the same machine, and the time before an individual researcher may actually see results, the "turnaround time," may be several hours. In one day, then, a researcher who really loves the job may be able to run and update the program five times. On the Cray, each run will take about two minutes, and the results will be available within three minutes. Hence, precious "people time," often more expensive than computer time, is drastically reduced by using a supercomputer.

Mechanics of Ductile Rupture

A. Needleman and R. Becker
Brown University

BACKGROUND

A fundamental understanding of fracture mechanisms is crucial for assessing the safety and reliability of structures and mechanical components and for laying ground work for the design of stronger and tougher materials. Computer simulations are being used to study ductile fracture processes to relate macroscopic toughness and ductility to measureable (and controllable) features of the material's microstructure. Ductility refers to a material's ability to deform without tearing; toughness characterizes the resistance of material to crack growth.

THE QUESTION

Two types of calculations are being performed at the Pittsburgh Supercomputing Center. One class of calculations is aimed at developing theories that can accurately describe the process of void nucleation, growth and coalescence and their effect on the overall mechanical response of a material. For example, structural metals, such as steels and aluminum alloys, generally contain inclusions or voids of various sizes distributed randomly throughout the material. These inhomogeneities play a dominant role in the failure of ductile metals.

In collaboration with Professor Subra Suresh of Brown University and Dr. Viggo Tvergaard of the Technical University of Denmark, a second set of calculations is being carried out to investigate progressive failure at crack tips in an advanced aluminum-lithium alloy. The limited toughness of such alloys hinders their usefulness. Calculations performed to date have provided remarkably accurate quantitative predictions of fracture toughness in terms of material properties and the density and distribution of defects.

The calculations are based on constitutive relations incorporating physically-based models of the micro-rupture process. In other words, models of fracture and of the nucleation, growth, and coalescence of microvoids arise as a natural outcome of the deformation process. To analyze the effects of microstructure on fracture behavior, the complete loss of stress-carrying capacity and the associated creation of new free surfaces are incorporated into the constitutive description of the material. This formulation leads to a class of novel and extremely challenging computational problems.

COMPUTATIONAL TECHNIQUES

The effects of the nonuniformity of the defect distribution are investigated by discretizing a volume of the material. Based on experimental observations of porosity in sintered iron specimens (iron that has previously been raised to molten temperatures), each subregion is assigned a defect density. Deformation of the region through the entire failure process is analyzed by a novel finite element method.

The processes analyzed here, although quasi-static, are strongly history dependent and must be calculated in a stepwise, incremental manner. The nonlinearity of the equations, coupled with the unstable nature of material failure, necessitates the use of many time steps. Although no single increment requires much computer time, complete analyses have taken from 30 minutes to 30 hours of CPU time on the Cray X-MP; the type of investigations carried out here require many such extensive parameter studies. For comparison purposes, the computer programs used have also been run on a VAX-11/780 computer and have been found to run between a factor of 120 to 180 times faster on the Cray X-MP than on the VAX-11/780. Clearly, only the computational speed of a supercomputer such as the Cray X-MP makes these investigations feasible.

RESULTS

Numerical results show the strong role played by micro-inhomogeneity in limiting ductility; in the circumstances that were studied, the nonuniformity of the defect distribution reduced ductility by a factor of two. The first figure shows an actual void distribution near a fracture surface in a sintered iron specimen. The second figure shows the predicted porosity (computed from this study) just prior to fracture. The last figure shows a predicted fracture surface formed from the coalescence of voids. Note that in the first two figures there is a large void near the fracture surface that does not take part in the final fracture. These simulations provide a tool for understanding how voids coalesce to form fracture surfaces. Eventually, of course, one would like to know how to delay this process.

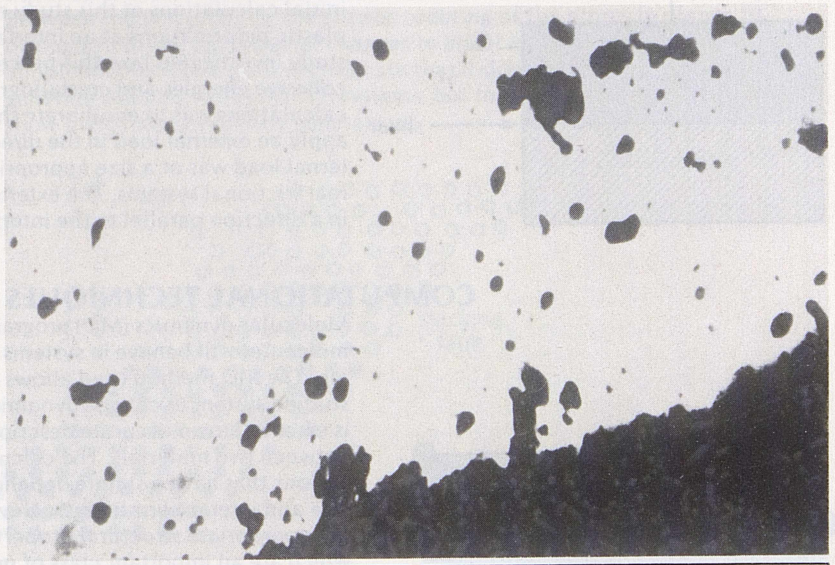

Observed porosity distribution near the fracture surface in a sintered iron specimen. The fracture surface is the large, blue area to the lower right.

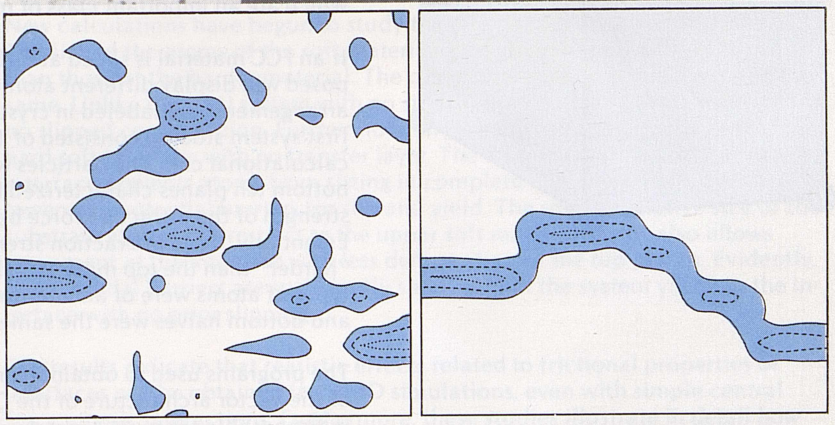

Left
Computed porosity contours prior to fracture.

Right
The predicted fracture surface created by the coalescence of voids.

FUTURE DIRECTIONS

Further work will use this information as a basis for understanding the factors that limit toughness in these alloys with the aim of discovering ways to improve their ductility. In addition, crack growth, as determined from this microstructurally-based model, will be related in future studies to general phenomenological descriptions so that the factors responsible for crack growth may be identified.

Wear and Tear at the Atomic Level

M.W. Ribarsky and Uzi Landman
Georgia Institute of Technology

BACKGROUND

Tribology, the study of the causes and consequences of friction between two surfaces, is of primary importance in industrial engineering applications to study the wear of moving parts. Research carried out at the Pittsburgh Supercomputing Center has resulted in the development of molecular dynamics programs that examine for the first time the detailed microscopic dynamics of friction and wear caused by deformations.

THE QUESTION

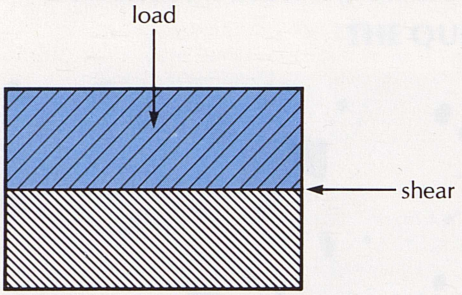

When placed under stress, solids may undergo elastic or plastic deformations. In an elastic deformation, the system will return to its original shape when the external stresses are removed; plastic deformations produce a new shape. The initial calculations of this study explored how elastic deformations yield to plastic deformations at an interface between two materials. In particular, the study investigated how this process depends on system parameters, such as cohesive energies and crystallographic orientation. The approach in all the calculations was to equilibrate the system at a low temperature and th apply an external load in the direction perpendicular to the interface. The external load was of a size appropriate to the loads on the areas of contact of real frictional systems. The external shear stress was then gradually increased in a direction parallel to the interface until the system yielded.

COMPUTATIONAL TECHNIQUES

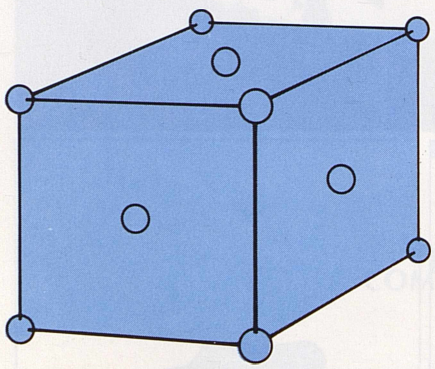

Molecular dynamics (MD) programs were developed to study how individual molecules will behave in systems under these conditions of loading and shearing. The MD method used allows the volume, shape, and structure of the studied system to change dynamically under an applied external stress, which is essential to an accurate description of the deformation at an interface between two materials. The calculations were performed using atomic interactions that have a simple dependence on parameters that control atomic size and interaction strength. Nevertheless, these simple interactions exhibit the appropriate structural properties of face-centered cubic (FCC) materials, which are an important class of materials in tribological applications.

If an FCC material is sliced at different angles, the surfaces that will be exposed will display different atomic arrangements. Some of the most common arrangements are labeled in crystallographic notation by (001) and (111). The first system studied consisted of 1,200 particles in each periodically-repeated calculational cell. The particles were arranged in 20 (001) planes with the bottom ten planes characterized by an interatomic attraction twice the strength of the attractive force between atoms of the top half of the system. Essentially, these interaction strengths imply that the bottom material is "harder" than the top material. The forces at the interface between these two types of atoms were of an intermediate strength. The atomic sizes in the top and bottom halves were the same.

The programs used to obtain the results described below take great advantage of the vector architecture of the Cray X-MP. Access to the Cray at the PSC was indispensable for quickly producing the results of this work. Similar programs took approximately ten times longer to run on a Cyber 885 and three times longer on a Cyber 205 supercomputer.

RESULTS

An inspection of particle trajectories after a slip occurred in this system revealed that a transfer layer of two or three atomic planes from the top half adhered to the bottom half. The trajectories also showed that slip did not occur in the (001) planes during yield. Instead, successive planes in the soft material

showed that one new row at a time joined the flow. A geometrical analysis of this slip revealed that it involved planes that intersected the (001) planes parallel to the interface. Thus, even though the system was sheared along the (001) plane, the slip occurred in (111)-type planes.

To study the effects of crystallographic orientation, a system with a (111) interface was then considered. The interaction strength and atomic size parameters were the same as in the (001) system. After loading and equilibration, the (111) system showed slip to a stable new structure. A close look at the atomic trajectories showed that this new structure corresponded to the formation of a "stacking fault" in the soft material. The figure below shows a central slice of the (111) system before and after the creation of the stacking fault. Between the two pictures, the *abcabc* ordering of the atoms in planes 10 through 13 of the top material changes to *ababca*, forming a stacking fault. Simple theoretical arguments indicate that it does not take much energy to generate stacking faults in (111) systems, but for the first time these results show the dynamics of their generation.

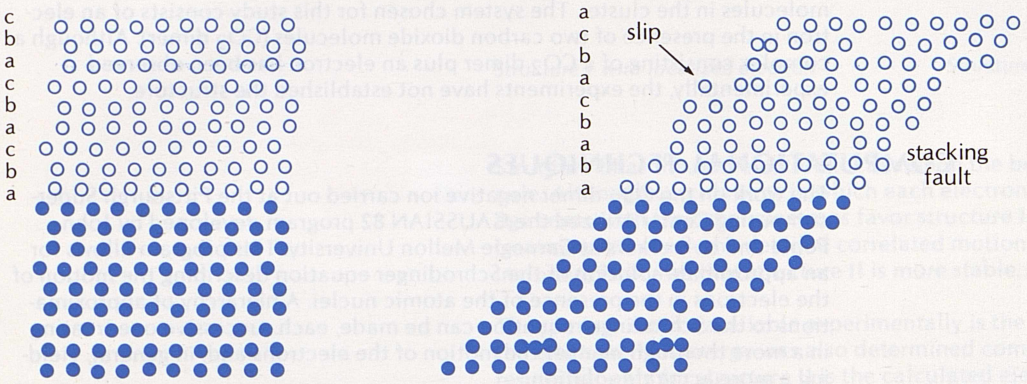

Central slice of a (111) system before (left) and after (right) generation of a stacking fault. The sequencing of the atoms above the interface is different before and after.

New calculations have begun to study the effects of atomic size mismatch. The size of the atoms of the soft material is assumed to be 1.5 times larger than those of the harder material. The other system parameters remained the same. Unlike the (111) system with no size mismatch, this system developed no slipped configuration. Furthermore, the system yielded directly at the hard-soft interface with no transfer layer. The atomic size mismatch at the interface prevents atoms from sitting in complete alignment there and thus lowers the potential barriers for slip and yield. The smaller relative size of the substrate atoms with respect to the upper soft material atoms also allows movement at the interface with less deformation of the top planes. Evidently, the potential barriers are sufficiently small so that the system yields at the interface with no prior slipping.

The results indicate that realistic effects related to frictional properties of interfaces can be obtained from MD simulations, even with simple central force atomic interactions. Furthermore, these studies illustrate in detail how atomic parameters such as interaction strength and atomic size can significantly affect the nature of the yield process at the interface of two different materials.

3 CHEMISTRY

Quantum Chemistry of Molecular Clusters

Kenneth D. Jordan
University of Pittsburgh

BACKGROUND

The properties of molecular clusters are particularly interesting because clusters play an important role in a variety of processes, including nucleation, and aerosol and droplet formation. Clusters have also attracted considerable interest due to their utility in studying the dynamics of intermolecular and intramolecular energy transfer.

THE QUESTION

A fundamental question in the study of charged molecular clusters is whether charge is localized on one molecule or whether it is spread out over several molecules in the cluster. The system chosen for this study consists of an electron in the presence of two carbon dioxide molecules (CO_2 dimer). Although a complex consisting of a CO_2 dimer plus an electron has been observed experimentally, the experiments have not established the structure.

COMPUTATIONAL TECHNIQUES

The work on the CO_2 dimer negative ion carried out at the Pittsburgh Supercomputing Center utilized the GAUSSIAN 82 program developed by John Pople and co-workers at Carnegie Mellon University. This program allows for an approximate solution of the Schrodinger equation describing the motion of the electrons in the presence of the atomic nuclei. A hierarchy of approximations to the Schrodinger equation can be made, each successive one treating in a more thorough manner the motion of the electrons and, in general, yielding a more accurate solution.

In the case of the CO_2 dimer negative ion, the Schrodinger equation describes the motion of 45 electrons in the presence of six nuclei, making an accurate solution a formidable computational task.

Since the energies of the stable configurations of this system are quite close, it was essential to use a "high-order" approximation requiring approximately five CPU hours of Cray X-MP time per structure. This code runs 120 times faster on the Cray than on a VAX-780. The largest runs use between 2 and 3 Megawords of memory. Without the availability of the Cray, higher order approximations would have been unmanageable.

RESULTS

For the CO_2 dimer negative ion, two stable arrangements of atoms were found, and are depicted below.

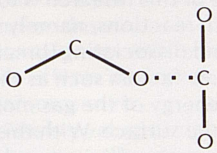

Structure I with localized electron

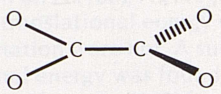

Structure II with delocalized electron

Structure I has the odd electron localized on the bent (leftmost) CO_2 molecule. Although calculations in which each electron interacts only in an average sense with the other electrons favor structure I, a more sophisticated treatment (in which the detailed correlated motions of the electrons are included) indicates that structure II is more stable.

The primary datum available experimentally is the energy required to eject the electron. This energy was also determined computationally, with the result that only for structure II is the calculated ejection energy consistent with experiment.

FUTURE DIRECTIONS

Due to the increase in complexity of the problem with the number of atoms, a straightforward extension of the procedure used for the CO_2 dimer to larger clusters will not be possible, even with the computing power of the Cray X-MP. New programs are being developed, therefore, which will build on the information obtained from the dimer calculations to create models of larger clusters. With these techniques, it will be possible to extend investigations to systems with up to several hundred atoms.

Chemical Reactions Between Gases and Solid Surfaces

Andrew E. DePristo and Chyuan-Yih Lee
Iowa State University

BACKGROUND

A fundamental problem in theoretical chemistry is to provide a detailed understanding of the factors that influence chemical reactions that occur at gas-solid interfaces. Such reactions are of immense practical importance since they occur in many industrial applications, including petroleum refining, microelectronic fabrication, and production of ammonia from nitrogen and hydrogen. They are also among the most difficult reactions to treat by theoretical techniques since they involve a strong interaction among a small number of atoms (e.g., a gas molecule and a few nearby solid atoms), which in turn interact less strongly with a vast number of atoms (e.g., the remainder of the solid).

THE QUESTION

The objective of this research is to elucidate the features of the most fundamental of such reactions, namely those involving a single gas molecule interacting with and dissociating (breaking apart) on a metal surface. The goal is to understand how factors such as the structure of the solid surface and the velocity and energy of the gas molecule influence the dissociation of the molecule on the surface. With this knowledge, it may be possible to make such reactions more efficient and selective via modification of either the surface or the gas molecule.

COMPUTATIONAL TECHNIQUES

In the calculations in progress at the Pittsburgh Supercomputing Center, all molecules in the gas and all atoms in the solid are assumed to obey classical mechanics. An individual calculation is referred to as a trajectory and requires the numerical solution of approximately 45 coupled, second-order differential equations. Random forces and frictional forces are included in order to realistically describe energy exchange between a small number of solid atoms near the surface and the remainder of the bulk solid. Solution of these equations is accomplished using standard algorithms, specially adapted to the vector architecture of the Cray X-MP. In particular, evaluation of all the forces in the system is the most time-consuming step, and a great deal of effort has been spent in the vectorization of this part of the code. Each calculation must be repeated many times and the results averaged to yield physically meaningful values. As an illustration of the computational speed advantage offered by the Cray X-MP, a run of 100 trajectories requires about five hours of CPU time on a FPS 164, but only eight minutes on the X-MP.

To understand the importance of a supercomputer to this research, one must realize that the forces responsible for the chemical reaction are not known to the accuracy needed to predict the attributes of the reaction reliably. Thus, the forces between the molecule and the solid surface must first be determined theoretically from dynamic simulations. This is accomplished with an iteration procedure as follows:

1. *An interaction potential which yields the forces is assumed.*
2. *The dynamical simulation is performed.*
3. *The results are compared with experimental data.*
4. *The interaction potential is adjusted.*

Steps 2 through 4 are repeated until satisfactory agreement is found. This process may require many iterations before a good description of the forces is discovered. Each iteration requires human intervention to interpret the results and to adjust the interaction. For complicated systems such as molecules on solid surfaces, the simulations could not be performed fast enough to allow meaningful results without a supercomputer.

Once accurate interaction potentials have been determined with the above procedure, further simulations are then performed to determine all the factors influencing the chemical dynamics at the surface. These factors are interpret-

ed in terms of characteristics of the interactions. Since experimental data do not yield information about all of the important factors, more insight is provided by these theoretical studies than by experimental results alone.

RESULTS

As an illustration of the type of information which this study is providing, consider the dissociation of hydrogen gas (H_2) on nickel (Ni) surfaces. If crystals are cut at different angles to expose different surfaces, or "faces," then the chemistry at each surface will also be different. Commonly-studied faces of crystal are labelled in crystallographic notation by: (100), (110), and (111). In the PSC study, the relationship between the probability of dissociation, S_o, and the translational, vibrational, and rotational energy of the hydrogen molecules was calculated for H_2 dissociation on Ni(100), Ni(110), and Ni(111) crystal faces. While the variation of S_o with translational energy was known, the effects of rotational and vibrational excitation were not. A substantial enhancement of S_o with increasing vibrational energy was found on the more closely-packed (111) and (100) faces but not on the open (110) face. By contrast, it was found that on all of these surfaces, rotational excitation of the H_2 inhibits dissociation by preventing orientations of the H_2 which lead to dissociation. This finding implies that the structure of the surface and excitation of the gas molecule influence the reaction in an interrelated manner, and that no simple rule about variation of S_o should be expected.

FUTURE DIRECTIONS

Further work planned and already in progress will build on these encouraging results to study a variety of gas-solid surface systems, including nitrogen gas dissociation on tungsten and iron, and carbon dioxide dissociation on nickel.

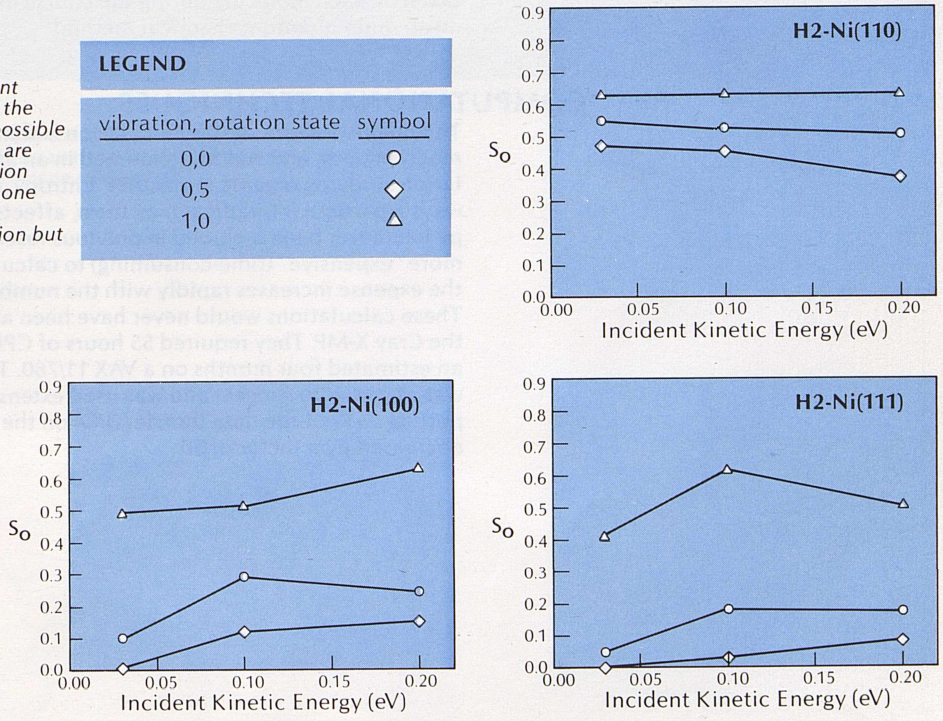

The probability, S_o, that a hydrogen gas molecule will dissociate on three different nickel surfaces is shown as a function of the incident energy of the molecule. Three possible kinetic states of the hydrogen molecules are indicated: one with no vibration or rotation about the hydrogen bond ((0,0) ○ state), one with no vibration but some rotation ((0,5) ◇ state), and one with some vibration but no rotation ((1,0) △ state).

Theoretical Studies of Biradical Dynamics

Charles Doubleday, Jr.
Columbia University

BACKGROUND

Simple organic reactions in the gas phase at high temperatures (above 300°C) are important in industrial applications, flame chemistry, and atmospheric chemistry. These reactions are believed to proceed in two stages. In the first stage, the starting material forms an intermediate molecule that exists only for a very short time. The particular reactions studied here give rise to intermediates called "biradicals." In the second stage, the biradical decays to form any of several products of the reaction. Much experimental research has been directed toward understanding the nature of biradicals since their properties determine the product distribution and the rates of product formation. Unfortunately, the biradicals produced in these reactions are difficult to study because they exist for less than a billionth of a second. Firm experimental data on the properties of biradicals are therefore rare, and several important questions in biradical chemistry are unanswered by the existing data. This is an area in which theory can be of great value, both by suggesting interpretations for existing data and by driving new experiments.

THE QUESTION

The goal of this research is to understand the mechanism by which products are formed in these high temperature reactions. The reactions studied must be simple enough to be computationally tractable; but they must, in some sense, be a prototype of a large class of reactions so that conclusions drawn from specific cases have wide applicability. The figures on the facing page show the two reactions chosen for this study. The first proceeds via a trimethylene (three carbon) biradical, and the second involves the tetramethylene (four carbon) biradical.

To be a true intermediate, a biradical must exist in a local free energy minimum. The question posed was: over what temperature range do trimethylene and tetramethylene exist as local free energy minima (true intermediates), and what is the free energy required? To answer this, the minimum energy path taken by each molecule during the course of its reaction was calculated, using an *ab initio* quantum chemical method.

COMPUTATIONAL TECHNIQUES

The unique feature of this calculation is that entropy was included along the reaction path, and was implemented in an *ab initio* quantum chemistry calculation studying organic molecules. Entropy, which measures the number of ways a particular biradical may form, affects the free energy required. Entropy had never been included in previous research, primarily because it is much more "expensive" (time-consuming) to calculate than potential energy, and the expense increases rapidly with the number of atoms in the molecule. These calculations would never have been attempted without the speed of the Cray X-MP. They required 55 hours of CPU time on the Cray, compared to an estimated four months on a VAX 11/780. The vector capability of the Cray was essential to success and was used extensively in the computer code. By putting 95% of the data transfer (I/O) on the SSD, the I/O wait time was decreased by a factor of 50.

RESULTS

The results show that if entropy is ignored, trimethylene does not correspond to a potential energy minimum, and tetramethylene corresponds to a minimum too shallow to account for its existence as an intermediate. When entropy is included along the reaction path, the resulting free energy curve does show an appreciable minimum, provided that the temperature is high enough for entropy to dominate the changes in free energy. Trimethylene exists as an intermediate only when the temperature is greater than 1500K, but tetramethylene exists above 300K. Below these temperatures, the biradicals do not exist as intermediates (that is, their lifetimes are shorter than a molecular vibration). These entropy-locked biradicals are predicted to be extraordinary molecules: they exist only at high temperature, and disappear when the temperature is lowered.

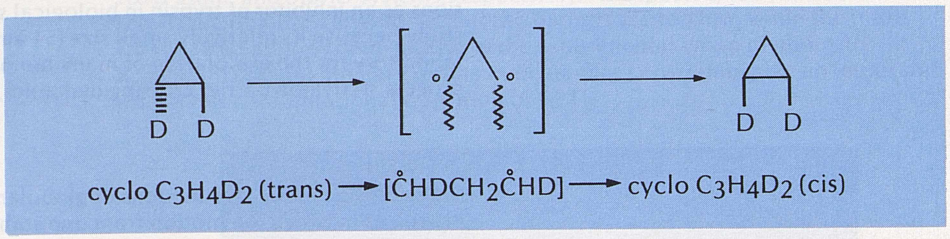

Two reactions involving biradicals whose rates were computed in this study. In the first, two deuterium atoms change their positions relative to the three carbons in the ring, passing through a trimethylene biradical structure. The two dots, indicating a broken carbon-carbon bond, is characteristic of biradicals.

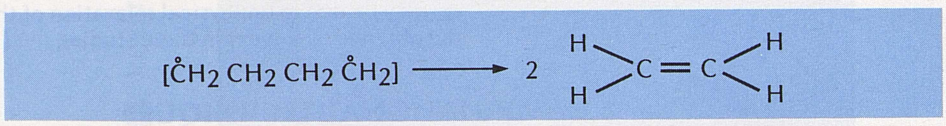

In the second reaction, a tetramethylene biradical splits apart in the middle to form two stable ethylene molecules.

4 BIOPHYSICS

Barrier Crossing In Biological Molecules

Douglas Tobias and Charles L. Brooks, III
Carnegie Mellon University

BACKGROUND

Globular proteins are linear polymers of amino acid which fold into structures consisting of one or more globular domains. They account for a large fraction of the material present in biological systems. These biological macromolecules exhibit a wide range of functional diversity in living organisms, contributing to the structure and executing most of the tasks necessary for function. The protein bovine pancreatic trypsin inhibitor (BPTI), which functions as an inhibitor of trypsin in biological systems, is used in the present study because its relatively small size (54 amino acids) makes it an ideal laboratory for the exploration of many biological processes. One such process, activated barrier crossing dynamics, is the focus of the present work.

THE QUESTION

Many of the biological functions of globular proteins involve activated processes in which the motion from one stable configuration of the protein to another is limited by the rate of overcoming an energy barrier. Special simulation techniques have been developed for the calculation of the free energy barriers associated with such processes. Using molecular dynamics methods, this study (underway at the Pittsburgh Supercomputing Center) assesses the roles of both implicit and explicit representations of solvent in empirical energy functions on an activated transition of the protein BPTI. Previous calculations, which ignored the aqueous solvent, have produced energy barriers that are too low to account for the experimental rate constant; this discrepancy may be due to unphysical relaxation of the surrounding protein matrix in the absence of solvent in these studies.

COMPUTATIONAL TECHNIQUES

In the current study, the potential of the mean force (PMF) (free energy change) is being calculated as a function of coordinates for a transition of BPTI involving the rotation of the protein's tryosine-35 sidechain (see figure below). To determine the trajectories necessary for the PMF calculations, the Cray X-MP/48 was used to implement the program known as CHARMM, a general-purpose software package which applies molecular mechanics (molecular dynamics, energy minimization, and normal mode analysis) to large systems such as aqueous solvated biopolymers. Several 22 picosecond trajectories were used to construct each PMF. The PMF for the tyrosine-35 ring flip in the full 580-atom protein was first computed in vacuo. To explore how different constraints, such as fixing atoms, might mimic some of the effects of solvent, a similar calculation was carried out in which all atoms beyond 10 A of the ring were fixed. In addition, the vacuum PMF was calculated using stochastic boundary molecular dynamics (SBMD). With this method only 343 atoms were simulated explicitly; the remainder were replaced by a stochastic heat bath and boundary forces.

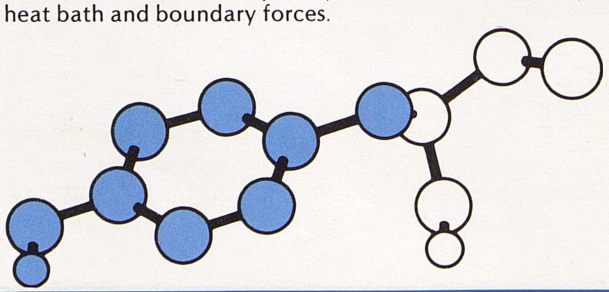

The tryosine-35 sidechain (dark area) of the globular protein BPTI. Rotation of the side chain about the single bonds connecting it to the rest of the molecule are possible if energy barriers can be overcome.

To calculate each 22 picosecond trajectory on a VAX 11/780 would require about 80 CPU hours for the full system and about 55 CPU hours for the reduced SBMD system. On the Cray X-MP/48, the same calculations take 45 minutes and 25 minutes, respectively. In the Cray implementation of CHARMM, the computationally intensive parts of the calculation were vectorized and the hardware gather-scatter was used. The tremendous time savings afforded by the X-MP allow calculations to be performed which are essentially intractable on other available machines.

RESULTS

Some preliminary results of these calculations are illustrated in the figure below. The energy barrier heights are 10.0, 17.3, and 8.2 kcal/mol for the full, fixed-atom, and SBMD systems, respectively. As expected, the full and SBMD results are qualitatively similar, but the barriers are too low, while the result for the fixed-atom system is close to the experimental value of about 16 kcal/mol. These preliminary results indicate that constraints play an important role in this conformational transition of BPTI.

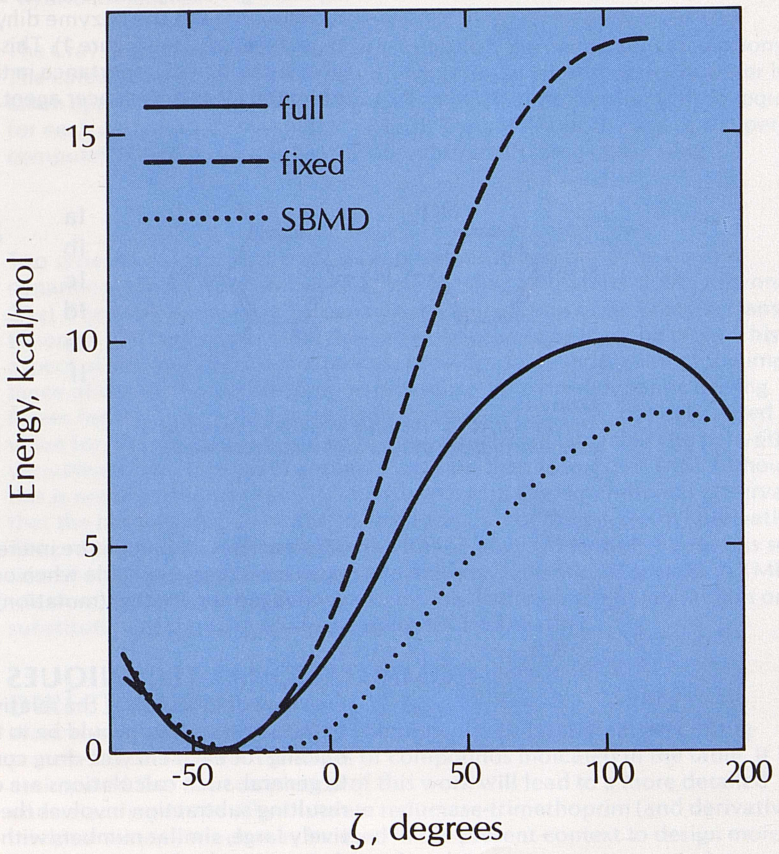

The potential of the mean force (PMF) as a function of position for the rotation of the tryosine-35 sidechain. Three different calculational assumptions were made; the resulting curves are shown. The higher the curve, the more energy is required for the transition.

FUTURE DIRECTIONS

This work is currently being extended to include computation of the PMF using SBMD with the explicit inclusion of several hundred water molecules. The results of these studies will provide a valuable base of information regarding the roles of external constraints, particularly those due to solvent, on the rates of activated processes in biological molecules.

Drug-Protein Binding

Stephen H. Fleischman and Charles L. Brooks, III
Carnegie Mellon University

BACKGROUND

Molecular dynamics, the calculation of the equations of motion of atoms in a system of molecules, can be used to simulate various processes of biological interest involving proteins. Of specific interest is the binding of drug molecules to the active sites of enzymes. The goal is to increase our level of understanding of the complicated interactions that determine drug-protein binding to the point of rational prediction of drug design changes. Such a process is quite complicated, however, involving desolvation (the stripping away of solvating water) of the drug molecule and parts of the protein. In many cases, free energy of binding determines the rate at which the drug binds to a protein. Unfortunately, the computational time involved generally makes the calculation of free energies for large systems intractable, even with current supercomputers.

THE QUESTION

However, it is often sufficient to know the *relative* free energies of binding for closely related systems. For example, one may wish to compare the free energies of binding of different derivatives of a given drug molecule to an enzyme or, more generally, to a receptor site. The binding of trimethoprim (TMP) and its derivatives to the enzyme dihydrofolate reductase from chicken is currently under study (see Figure 1). This system was chosen for study because of its pharmacological importance in the treatment of bacterial maladies and its potential as an anti-cancer agent.

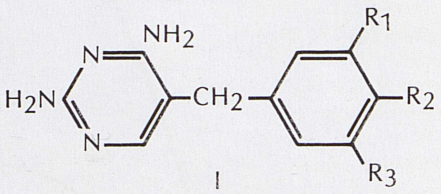

I

	R_1	R_2	R_3
Ia	OCH_3	OCH_3	OCH_3
Ib	CH_2CH_3	OCH_3	OCH_3
Ic	OCH_3	CH_2CH_3	OCH_3
Id	OCH_3	CH_2CH_3	CH_2CH_3
Ie	CH_2CH_3	OCH_3	CH_2CH_3
If	CH_2CH_3	CH_2CH_3	CH_2CH_3

Figure 1:
The derivatives of the drug trimethoprim (Ia) each have one or more oxygen atoms replaced with CH_2 groups.

Alternatively, one might be interested in the relative binding energies of a particular drug molecule when one amino acid residue of the enzyme is exchanged for another (mutation).

COMPUTATIONAL TECHNIQUES

One way of calculating the relative free energy of binding two drug molecules to the same enzyme would be to first calculate the absolute free energies of binding for each enzyme-drug complex separately, and then subtract the two. In general, such calculations are computationally infeasible. Moreover, the resulting subtraction involves the small and unreliable difference of two relatively large, similar numbers with high statistical uncertainties. The thermodynamic cycle shown in Figure 2 suggests an alternative approach.

ΔA_1 and ΔA_2 are the absolute free energies of binding for each individual drug molecule (D1 or D2) to an enzyme (E). ΔA_3 and ΔA_4 represent free energies for the non-physical processes in which one drug molecule is transformed into the other. However, the free energy differences $\Delta A_2 - \Delta A_1$ and $\Delta A_4 - \Delta A_3$ are thermodynamically equivalent and correspond to the relative free energy of binding for the two drugs. Only the latter difference is computationally tractable.

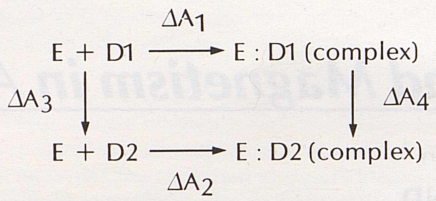

Figure 2:
Theoretical Thermodynamic Cycle for Drug-Protein Binding.

To effect the calculation, the complexed and uncomplexed drug molecules are transformed into the desired "perturbed" systems in several steps. Each step involves a complete molecular dynamics calculation. Thermodynamic-cycle perturbation theory is used to calculate the free energy required to go from one step to the other. The intermediate results are summed to give the total free energy difference. The intermediate steps represent imaginary "hybrid molecules."

The Cray X-MP makes it possible to accomplish the necessary calculations. The calculation of 75 picoseconds of dynamics of trimethoprim in water has taken approximately 12 hours of Cray X-MP CPU time. Similar time is required for each derivative and for each "hybrid" step. Without the use of a supercomputer, the calculations would be prohibitively time-consuming.

RESULTS

Two series of calculations must be carried out to complete the thermodynamic-cycle depicted in Figure 2 (one for the calculation of ΔA_3 and one for ΔA_4). The work to date has focused on the calculation of ΔA_3 for the transmutation of TMP to its para-ethyl derivative, species Ia to Ic in the table. This aspect of the thermodynamic process provides some indication of the importance of desolvation of the drug in providing the thermodynamic driving forces for binding to the enzyme dihydrofolate reductase. The calculated value for ΔA_3 for this step is 2.3 kcal/mol, which indicates that the para-ethyl derivative is less favored in aqueous solution than the parent TMP. Although this is not the whole picture, it is consistent with the experimental observation that the binding affinity of TMP is less than that of the para-ethyl derivative. The present result is also consistent with the physical model of aqueous solvation of an ethyl substituent versus the methoxy substituent present in TMP, which would indicate the loss of a favorable hydrogen bond interaction on the substitution of the para-methoxy group for a para-ethyl.

FUTURE DIRECTIONS

Work is continuing at the PSC to complete the characterization of drug-protein interactions for the series of compounds indicated in the table. It is anticipated that the completion of this work will lead to a more detailed understanding of the dihydrofolate reductase-trimethoprim (and derivatives) inhibitor system, which may be used in the present context to design more effective trimethoprim-based antimalarial drugs. In a more general sense, the thorough treatment of one explicit system should provide guidelines for rational drug design.

5 PHYSICS

Chemical Disorder and Magnetism in Alloys

F.J. Pinski
University of Cincinnati and Oak Ridge National Laboratory

BACKGROUND

Many alloys exist as solid solutions rather than as compounds. For transition metal alloys, these solid solutions can exist at a variety of temperatures and pressures. Many of these alloys have important technological applications, such as in the steel and magnetic recording media industries. One can think of these solid solutions, A_cB_{1-c}, as consisting of two types of atoms, A and B, randomly located on an underlying lattice. In this notation, c denotes the concentration of type A atoms, that is, the fraction of all atoms that are type A. Standard band theory techniques, which are often applied to crystal materials, cannot be applied here because of the broken symmetry introduced by the random location of the atoms in the alloy. Band theory requires long-range order, that is, a unit cell that repeats throughout space, and is not applicable to these solid solutions.

THE QUESTION

In this study the solid solutions are not modeled as a large collection of two different types of atoms. Instead, a technique has been developed to replace the random arrangement of atoms with an "average atom." This procedure provides a *parameter-free*, first principles method for studying alloys, calculating total energies, and determining phase diagrams. The goal is to understand not only how magnetism is modified by disorder but also how magnetism affects other properties of these alloys.

To understand the stability of transition metal alloys, one must have a basic understanding of the intrinsic role played by magnetism. In the iron-nickel alloy Fe_cNi_{1-c}, magnetism seems to enhance the ordering tendencies, while in the iron-vanadium alloy Fe_cV_{1-c}, magnetism inhibits ordering. The magnetic properties are, in turn, affected by the existence or absence of order. The most striking example is in the nickel-manganese Ni_cMn_{1-c} alloy. At a nickel concentration of $c = 0.75$, an ordering transition (into a crystalline structure like the copper-gold compound Cu_3Au) exists at 510°C. Samples prepared above this temperature can be quenched very rapidly to much lower temperatures without allowing short- or long-range order to develop. At these low temperatures both Ni_3Mn (ordered, crystalline) and $Ni_{0.75}Mn_{0.25}$ (disordered, solid solution) become ferromagnetic. The net magnetization in the ordered state corresponds to $\approx 1 \mu_B$/atom, but despite the fact that the concentrations of the nickel and manganese are the same, the magnetism is only $\approx 0.25 \mu_B$/atom in the disordered state.

Over the years, many have speculated on the origins of this large difference. Most theories have suggested that the net magnetization is decreased because the magnetic interactions are dependent on the local environment; that is, the magnetic moments of manganese atoms sitting next to one another are coupled anti-ferromagnetically and produce no net magnetization. In a disordered alloy, many manganese atoms could sit next to one another and could reduce the net magnetization of the sample due to their anti-parallel alignment.

RESULTS

The results of the PSC study do not support this view. Instead, the calculations show that the small magnetization in the solid solutions results directly from the effect of the disorder on the electronic states. Much of the bonding seen in the ordered (band theory) calculations cannot be supported in the presence of disorder. This is the first time that equilibrium properties, such as the magnetic moment and bulk modulus, have been determined for any solid solution in a parameter-free method.

COMPUTATIONAL TECHNIQUES

The "average atom" approach used in this study to handle disorder produces energy-dependent scattering, which complicates the computations and eliminates the use of linearization which has been important in speeding up band theory codes. In addition, both spin and charge self-consistency are necessary to determine the ground state properties in magnetic systems. These calculations are only possible with the large increase in computational speed that supercomputers can provide. The calculations described above have taken a total of approximately 100 hours of CPU time using the Cray X-MP at the Pittsburgh Supercomputing Center. Such calculations would have taken 70 to 80 times longer on an IBM 3033 or similar machine.

FUTURE DIRECTIONS

In the future, many other materials will be investigated. Recently, a new class of superconductors with transition temperatures above 90K has been found. Many properties of these materials remain unknown, including the effect of alloying. The programs developed and tested in this investigation are ideal to study these new superconductors, and plans are underway to begin that work.

Electronic Structure of High Temperature Superconductors

Warren Pickett
Naval Research Laboratory

BACKGROUND

The discovery by Bednorz and Mueller of superconductivity at 30K in 1986 was a tremendous advance, especially since two decades of work by a large number of researchers were necessary to raise the superconducting critical temperature from 17K in 1956 to 23K in 1973. At that time, the best superconductors were based on transition metals near the middle of the transition series, with the best material being a niobium germanium compound. The new superconducting materials are based on lanthanum copper oxide crystals, which are poor metals, and thus represent an entirely new class of superconductors. Since this first startling discovery, a large number of laboratories have reproduced and improved upon this result, so that now "anyone" can make materials that superconduct around 90K. A wide variety of specific materials based on lanthanum copper oxides which superconduct near this temperature now exist.

THE QUESTION

Several important and fundamental questions arise: what is the electronic structure of these materials, what is the mechanism that binds the pairs of electrons, is the atomic vibrational character important, how stable are the materials, and how high can the critical temperature be raised? This study applies first-principles local density methods to determine the electronic structure of these materials and to find changes due to both alloying and to motion of the atoms.

COMPUTATIONAL TECHNIQUES

These first-principles calculations involve heavy "number crunching." The solution to the pertinent local spin density equations requires the repeated set-up and diagonalization of large (dimension of 650 or more) matrices. The Eispack matrix manipulation software, available at the Pittsburgh Supercomputing Center, is used extensively and care was taken to optimize the most time-consuming parts of the code specifically to the Cray hardware. Large temporary disk storage of huge files is required, and the necessary adjustments to utilize the SSD at the PSC improved the computational speed of the code. These calculations require many tens of hours even on the fastest supercomputers; the required CPU time on a non-vector machine would be many times greater.

RESULTS

To understand the superconducting properties of these new materials, it is crucial to know the characteristics of the electrons in the most highly occupied states since these are the ones that become superconducting. The results of these calculations indicate that the electrons have a highly two-dimensional character, such that they stay entirely on the copper-oxygen layers and stay off the lanthanum-oxygen layers. When a small fraction x of the lanthanum (La) is replaced ("alloyed") with barium (Ba), electrons are removed from the highest occupied energy levels. This removal of some of the electrons results in an increase in the number of electrons that can participate in superconduc-

tivity, and for an alloying fraction of x=0.15, the character of the electronic system is altered. It is exactly at this concentration that the measured critical temperature is found to peak, at 35-40K, depending on sample quality. This agreement of theory with experiment for the critical value 0.15 clearly indicates that the change of character that occurs there is important for the superconductivity.

Since nearly all known superconductors rely on the interaction between the electrons and atomic vibrations to become superconducting, this study made the first estimate of the strength of this interaction. Assuming that the vibrational energies are relatively low, as seems to be the case from experiments to date, it has been found that this interaction may be strong enough to account for critical temperatures in the 30-40K range. Since the discovery of these superconductors, another class of materials with even higher critical temperatures of 90-100K has been studied by a number of experimental groups. These materials are also based on copper oxides, but include one of the group of atoms known as "rare earths." The most well-studied example is yttrium barium copper oxide. It is not yet known whether this new class is closely related to the 30-40K class or whether the superconductivity results from the same mechanism. Studies of these newest materials are well underway.

Computational studies such as these can provide a wealth of information not available by other means. For example, one can perform "computer experiments" in which a single atom is moved, watch the response of the electrons, and calculate the forces on the other atoms. In most metals that become superconducting, the electronic response to the motion of an atom is local, due to good screening by the electronic charge density. In the new ceramic superconductors, however, a new combination of ionic character and metallic character makes these materials behave very differently from ordinary metals. Such information is available from measurements only in a very indirect way. The current computational theory of electronic structure is accurate enough to describe the electronic character as well as to make specific predictions of material properties.

FUTURE DIRECTIONS

The current studies on the 40K superconductors will presently be extended to the 90K materials. The studies will involve calculation of the coupling strength due to atomic vibrations to test the prevailing opinion among most researchers that some new mechanism must be responsible for the very high critical temperatures. Also under investigation are the possible magnetic interactions that have been suggested by neutron-scattering experiments, as well as the important role of the oxygen vacancies that always seem to occur in these materials. Finally, the constants related to the electrical conductivity, with and without an applied magnetic field, will be evaluated from the work underway. The comparison of these quantities with actual measurements will be crucial in determining which technological applications of high temperature superconductivity are likely to be feasible with improved materials preparation.

Numerical Studies of the QCD Confinement Problem

R.E. Cutkosky
Carnegie Mellon University

BACKGROUND

Quantum chromodynamics (QCD) is the gauge theory of the interactions of quarks and gluons. Although they are the smallest known constituents of matter, free quarks and gluons are not observed in nature. It is presumed to be a consequence of QCD that they are permanently confined inside the observed physical particles. A crucial central problem of elementary particle physics is to understand how this confinement mechanism operates and, in turn, to understand the physical quark-gluon structures that can exist. To study the mechanism by which gluons become confined to the interior of physical particles, a series of numerical calculations has been carried out at the Pittsburgh Supercomputing Center.

One popular approach to QCD calculations is provided by lattice gauge theories, in which continuous space-time is approximated by a finite set of points. Each year, many thousands of supercomputer hours are devoted to calculations based on these theories. However, the lattice calculations do have some technical problems; for example, it is difficult to include quarks properly. To try to bypass these problems, a new method for numerical study of the QCD confinement problem is being developed and investigated at Carnegie Mellon University. The continuum theory, rather than the discrete lattice theory, is employed in a spatial region which is just big enough to contain an observable collection of quarks called a hadron. The aim of these calculations is to determine directly the energies of low-lying excited states (physical particles) and to help develop insight into their structure.

THE QUESTION

To check the feasibility of this method, a series of calculations using a simplified version of QCD has recently been completed. In these initial studies, only three of the eight kinds of gluons were included. These calculations explore the Gribov confinement scenario, which is based on properties of the color-electric energy. Ordinary electromagnetism involves massless photons and a Coulomb interaction between electric charges; in a similar way, QCD involves gluons and a Coulomb-like interaction between "color" charges. In QCD, however, the gluons themselves have charges and magnetic moments. In addition, Coulomb's law is replaced by a more complicated formula that not only depends on the distance between the charges but gives an increased force when other gluons are present. As a result of this property of the QCD Coulomb energy, the amplitudes of gluon waves cannot become too large. Long wavelengths should be especially affected, preventing gluon waves from extending beyond small regions and effectively confining the gluons.

COMPUTATIONAL TECHNIQUES

The convenient, accessible facilities provided by the Pittsburgh Supercomputing Center made it possible to complete this work quickly and easily. The present code has been vectorized well enough to run about 120 to 170 times faster on the Cray X-MP/48 than on a VAX 11/780. The speed is somewhat parameter-dependent, however, and has only been checked for those parameter values for which a VAX calculation is even feasible.

RESULTS

The numerical calculations support the Gribov confinement picture of gluon propagation. Electrical energy effects were found to have a crucial influence on the behavior of gluons and the properties of physical states. A contrasting view, that confinement is an indirect consequence of magnetic effects, is not supported by these results. Nevertheless, the average electrical energy in the vacuum state is actually less than the magnetic energy. A "glueball" emerged from the calculations as a colorless physical particle composed of gluons. The gluons are arranged so that color charges tend to cancel throughout the glueball, and the internal lines of color-electric flux are endless. The precise configuration of these flux lines has not yet been determined, but they may have a somewhat doughnut-like arrangement.

FUTURE DIRECTIONS

Besides providing these specific results, the calculations show that this new method provides a practical and promising way to investigate the properties of QCD. The work is now being extended to the realistic three-color (eight-gluon) version of QCD. In the future, quarks will be added as well, and additional hadronic states will be examined. These more complicated calculations seem entirely feasible with the supercomputing resources available at the PSC. In the meantime, even the simplified calculations provide insight into hadronic physics.

Quasi-Crystalline Order in a Simple Atomic System

Michael Widom, Katherine J. Strandburg, and Robert H. Swendsen
Carnegie Mellon University

BACKGROUND

It is a well-known result of solid state physics that one cannot build a crystal with five-fold symmetry, just as one cannot tile a floor with regular pentagons, as shown by the "frustration" in the figure below. (Triangles, squares, and hexagons *will* tile a floor.) It came as a great surprise when rapidly-cooled liquid manganese-aluminum alloys were discovered to exhibit five-fold symmetry. The structures were given the name "quasi-crystals." In this study, the physics of a simplified two-dimensional world was simulated to show how quasi-crystals could be formed. Studies like this will lead to an understanding of the structure of real quasi-crystals.

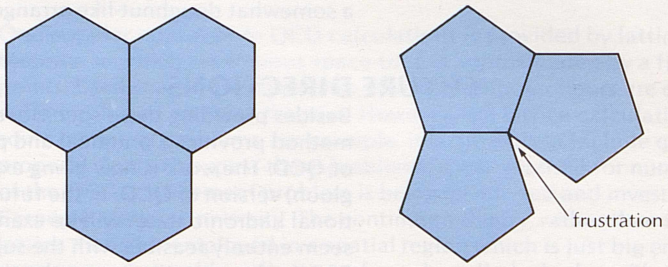

Tiling a floor is possible with hexagonal tiles (left) but not with pentagonal tiles (right).

Roger Penrose of Oxford University has shown that a plane can be tiled to exhibit five-fold symmetry by using two fundamental "tiles" whose shapes are shown below. The rules for arranging these "tiles" require that adjacent sides have the same number of arrows pointing in the same direction. One can see that this tiling has some aspects of five- or ten-fold symmetry because the angles are 1/5 or 1/10 of a complete 360° rotation. Unfortunately, these tiling rules seem to have no simple connection to the forces between atoms and, therefore, no connection to the laws of physics.

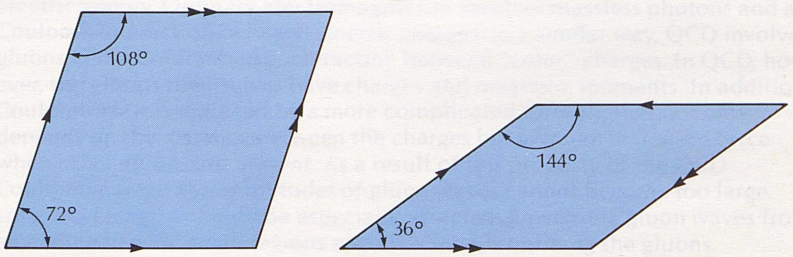

The two fundamental tile shapes of a Penrose tiling have angles that are multiples of 36°.

THE QUESTION

The work done at the Pittsburgh Supercomputing Center models a simple two-dimensional physical system, which has a hope of exhibiting five-fold symmetry. The model has two kinds of particles, large ones and small ones. There are three kinds of forces: those between large particles, those between small particles, and those between large and small particles. The forces are chosen to be repulsive at short distances and attractive at long distances. The equilibrium positions of the particles are chosen so that five large particles fit around a small one, and ten small ones fit around a large one. Having chosen the forces, standard methods from statistical mechanics are used to introduce temperature into the system. In the hot state, these particles are randomly distributed. The interesting question is whether the system will condense as it is cooled into a conventional crystal or into a quasi-crystal. If the latter, will it be like the Penrose tiling?

COMPUTATIONAL TECHNIQUES

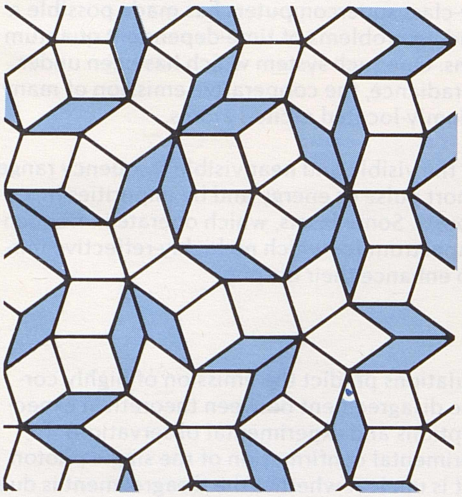

Penrose tiling.

The simulations began on a PC, with seven particles (six large and one small) to see whether, at low temperatures, the small particle would be surrounded by five large ones, with the sixth far away, or whether the six particles would prefer to share in surrounding the small one. Successive simulations modeled more and more particles; 30 particles were included on a microVAX, and finally 272 particles were studied using the Cray X-MP at the PSC. Since the computational complexity grows as the square of the number of particles, the Cray problem is about 80 times more difficult than the microVAX problem. The same code was used on the Cray as on the microVAX but ran about 500 times faster on the Cray. (The code was written with vectorization in mind.) The portability of the code enabled "friendly" workstation tools to be used to debug the program. The workstations were also crucial in representing the results of the simulations graphically.

One can exploit the unnatural dynamics of Monte Carlo algorithms to speed up certain atomic motions that occur on long timescales in nature; conventional simulation techniques rarely produce such atomic rearrangements. Dramatic reductions in the computation time required for equilibration result from inclusion of this mechanism.

The simulations on the Cray were necessary because phase transitions are not sharp in systems with few particles and because it is necessary to minimize surface-to-volume effects. Moreover, the system must be observed for long periods of time in order to be confident that the observed cooled states are truly equilibrium states and not simply metastable states that would revert to more traditional crystalline form.

RESULTS

The results of this study showed that this system exhibits an abrupt transition to a quasi-crystalline state as the temperature is lowered but that the quasi-crystal pattern is not of the Penrose form. The quasi-crystalline order is evident in the existence of rows of atoms along pentagonal directions, which is best seen by sighting along the figure at left at grazing incidence. There is great freedom in arranging the atoms in quasi-crystalline states. In fact, these results suggest that quasi-crystals may appear to violate the Third Law of Thermodynamics. The Third Law states that as the temperature of a system is lowered, it should become more and more ordered, restricting the choices available for atomic positioning.

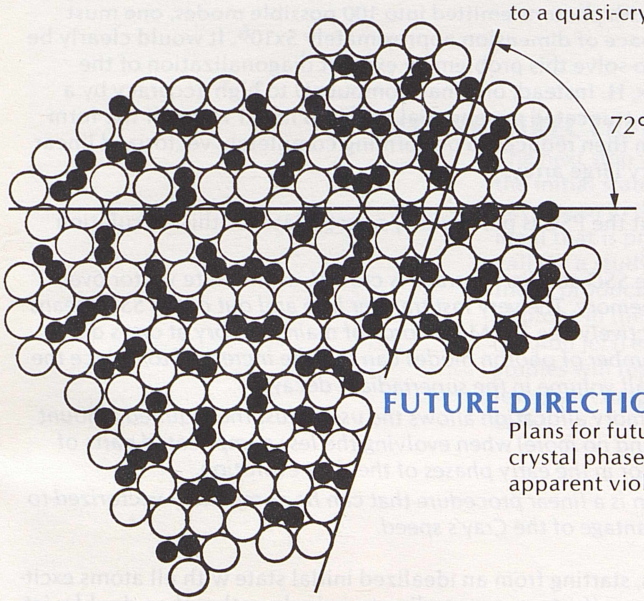

The long-range order of quasi-crystals is evident along lines which intersect at angles of 72°. These quasi-crystals do not, however, exhibit a conventional Penrose pattern.

FUTURE DIRECTIONS

Plans for future calculations include detailed studies of the liquid to quasi-crystal phase transition and low temperature simulations to elucidate the apparent violation of the Third Law of Thermodynamics.

Supercomputing and Quantum Optics

Anthony H. Duncan and Philip Stehle
University of Pittsburgh

BACKGROUND

The increasing availability of Cray-class supercomputers has made possible a renewed attack on many long-standing problems of time-dependent quantum theory in complex dynamic systems. One such system which has been under study is the phenomenon of superradiance, the cooperative emission of many photons from an assembly of randomly-located excited atoms.

Superradiance can be observed in the visible and near-visible frequency range in gases which are pumped by a short pulse of energy and by impurities in a crystal that are excited simultaneously. Some lasers, which operate in frequency ranges of the electromagnetic spectrum for which no highly-reflective mirrors exist, rely on superradiance to enhance their output.

THE QUESTION

Quantum and semi-classical calculations predict the emission of highly correlated radiation, but there is some disagreement between theoretical expectations based on particular assumptions and experimental observations. In particular, there has been no experimental confirmation of the single photon rays, or jets, predicted by theory. It is unclear whether the disagreement is due to the theoretical treatment or to the difficulty in achieving the assumed initial state (with essentially all atoms excited) in the laboratory.

COMPUTATIONAL TECHNIQUES

A supercomputer is an ideal tool for numerical studies of the time-evolution of quantum systems with many degrees of freedom. The technical problem is to apply the quantum mechanical time evolution operator, the matrix $U=e^{-iHt}$, to some initial state vector of very large dimension which reflects the large number of ways, or "modes," in which the photon can be emitted. This number increases with the size of the simulational box containing the atoms. For a four-atom system with photons emitted into 100 possible modes, one must deal with a state space of dimension approximately $5x10^6$. It would clearly be impractical to try to solve this problem by explicit diagonalization of the Hamiltonian matrix, H. Instead, one may compute U to high accuracy by a factorization of the truncated exponential e^{-iHt} into linear terms of the form (1-aH). The problem then reduces to performing completely vectorized linear operations on a very large array.

The Cray X-MP/48 at the PSC is particularly appropriate for this calculation for several reasons:

1. *The use of the SSD auxiliary storage is crucial, as the state vector overflows main memory. The very fast transfer into and out of the SSD means that one effectively has 128 Megawords of main memory at one's disposal. The number of photon modes can thus be increased to reduce the effects of small volume in the superradiant decay.*
2. *Dynamic memory allocation allows the use of just the required amount of memory (and no more) when evolving the less complicated parts of the state vector in the early phases of the time evolution.*
3. *The evolution is a linear procedure that can be completely vectorized to take full advantage of the Cray's speed.*

Preliminary results, starting from an idealized initial state with all atoms excited, show collimation effects corresponding to a single, rather than double, jet of photons and agreeing with earlier theoretical analyses. As an example of the sort of explicit information provided by a Cray simulation in this problem, refer to the graph of the multiphoton occupation probabilities (for the number of photons, n_γ, equal to 0, 1, 2, 3, or 4) as a function of time in a four-atom superradiant decay. At large times, the photons are reflected by the

walls of the finite container and induce oscillations, evidence for which can be seen in the diagram. At earlier times, before oscillations are induced, the occupation probabilities show the (roughly) exponential decay expected in an infinite volume problem.

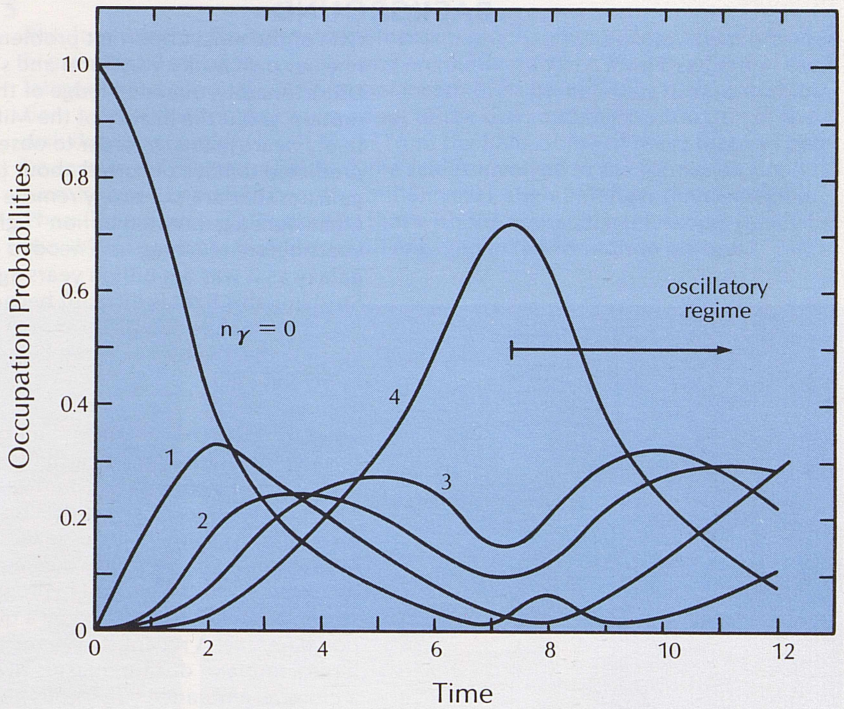

Probabilities as a function of time that multi-photon states will result from the superradiant decay of a four-atom system. The numer of photons in each state is denoted by n_γ. As the photons are reflected by the walls of the container, the probabilities exhibit undulating patterns (oscillatory regime).

FUTURE DIRECTIONS

The next step will be to expore how the final state of the photon depends on the initial state of the atoms. In the numerical studies to date, all the atoms have been in their excited state at the beginning of the simulation (t=0), an ideal that is probably impossible to achieve in the laboratory. Future plans call for a study of five atoms with four units of excitation distributed over them, producing a four-photon final state. The presence of an even number of photons in the final state makes the identification of directional correlations (photon jets) easier since there is no unpairable photon present. These new studies will represent the first time that systems of more than two atoms with less than complete population inversion have been studied in detail theoretically.

6 ASTRONOMY

Hydrodynamic Collapse of Protogalaxies

Allen V.R. Schiano and Arthur M. Wolfe, University of Pittsburgh
and C.A. Chang, Canadian Institute for Theoretical Astrophysics

BACKGROUND

One of the most important problems in modern astrophysics is how galaxies like our own Milky Way form and subsequently evolve to their present states. Unfortunately, our knowledge of this process is based almost entirely on information about the history of the Milky Way which has been inferred from relatively nearby stars. In order to observe the formation events directly, events predicted to have occurred about ten billion years ago, it is necessary to find galaxies that are extremely remote in space, and, therefore, in time as well. Light from a galaxy ten billion "lightyears" away has traveled for ten billion years before reaching us. Encoded in this light is the image of the distant galaxy as it was ten billion years ago, perhaps near the time of its formation. Studying this light is difficult because the starlight emitted by galaxies is not bright enough to be observed at distances of ten billion lightyears. Fortunately, the light emitted by equally remote quasars is sufficiently bright to be detected.

Quasars are extremely distant and powerful sources of many kinds of electromagnetic radiation, including visible light. The light from quasars will reveal the presence of a distant galaxy if the quasar lies behind the galaxy. The diffuse interstellar gas that resides in the disk of the Milky Way and in the disks of most spiral galaxies leaves a recognizable spectral imprint on the quasar light. The imprint is a discrete resonance absorption line, Lyman alpha (Lα), created by the neutral hydrogen gas in the disk. The Lα line will be very strong since it is broadened by a quantum mechanical effect known as *radiation damping*. A recent survey for these damped Lα lines in the spectra of a large sample of distant quasars revealed a population of remote objects with a close resemblance to gaseous galactic disks. Moreover, the cosmological mass density of this population is comparable to the mass density of stars in present-day disk galaxies. This coincidence raises the exciting possibility that for the past ten billion years, gas in the remote population has been converted into stars like those we see today; these remote objects may be the progenitors of current disk galaxies.

THE QUESTION

The goal of this project is to understand the formation of these gaseous disks using two-dimensional hydrodynamic simulations. The idea is to follow the collapse of a spinning, gaseous protogalaxy in the background gravitational potential of a massive halo of non-visible "dark matter." The gas dissipates energy radiatively, while the non-dissipative halo only interacts gravitationally. Fluid flow equations are solved starting with initial conditions, such as the angular momentum distribution and the ratio of disk matter to gas mass, which are set by cosmological constraints. The calculation ends when the gas collapses to a rapidly spinning disk that is supported by its centripetal acceleration.

COMPUTATIONAL TECHNIQUES

The PSC Cray X-MP was invaluable in these simulations not only because of the large spatial and temporal extent of the calculations but also because of the large range of initial conditions possible in the model. The calculational region was divided into a 130 x 130 mesh and evolved through approximately 10,000 variably-sized time steps. A typical simulation took one hour on the X-MP, compared to approximately 50 hours on a VAX 8600. More than 40 different simulations were created to study a wide range of initial conditions.

RESULTS

Results show that the formation process is much richer in detail than previously expected. Among the phenomena encountered were: long-lived radial oscillations of the gas disk, gravitational instabilities modified by differential rotation, propagation of shock waves out of the disk, and the contraction of the disk due to secondary infall of gas from the halo. Some of these features can be seen in the diagrams below. Time scales involved in the formation process have been estimated to be a few billion years, slightly longer than expected. These longer time scales suggest that with the new space telescope, it may be possible to see distant galaxies still engaged in the formation process.

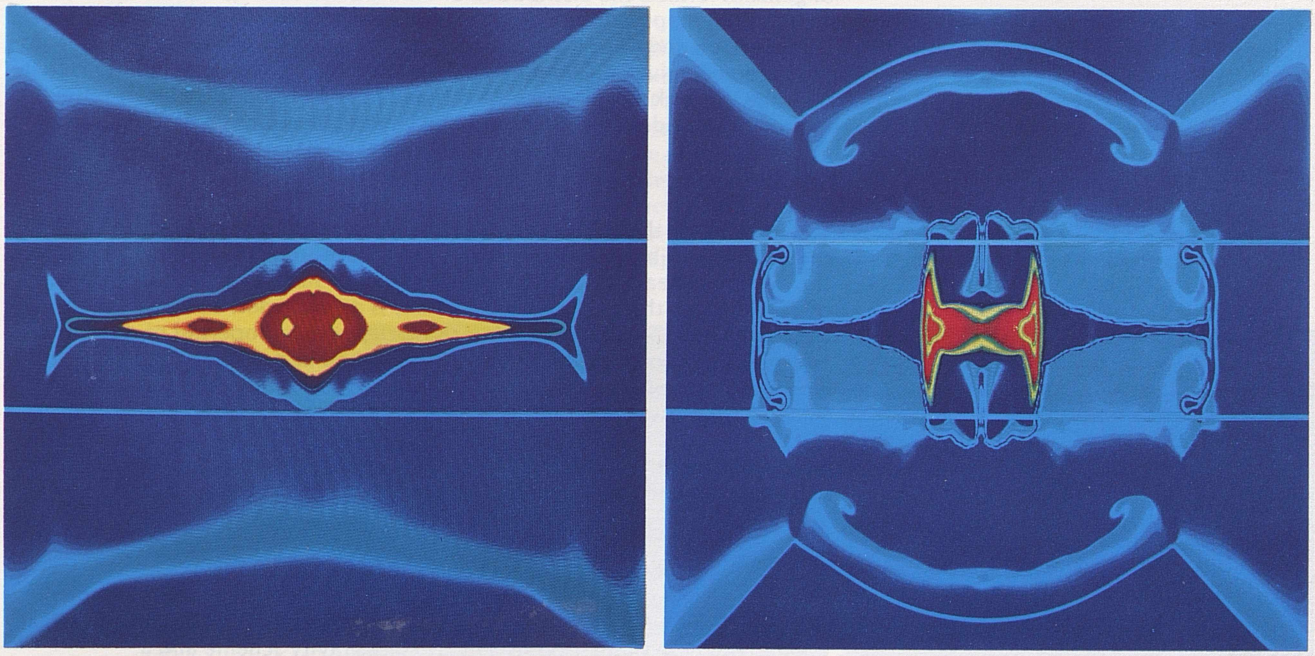

"Snapshots" representing early (left) and late (right) phases in a computational simulation of galaxy formation. The (reddish) disk material contracts into a rotating disk which then produces shock waves (light blue) that propagate outward into the "halo" of non-visible dark matter (blue).

The vertical distance scale of each figure spans 120 kiloparsec (kpc); the central 16 kpc have been expanded to show detail near the plane of the galactic disk.

The results of this study also indicate that in order for protogalaxies to form the type of large disk galaxies which are currently observed, they must either be initially larger than previously thought (destroying theories proposing dynamic independence between them) or be surrounded by large, hot halos of dark matter. Either possibility is tantalizing.

FUTURE DIRECTIONS

Future plans include allowing the halo to adjust to the formation of the disk and the formation of the first generation of galactic stars. Comparisons between the nearby disks that will be found with the space telescope and their remote counterparts may constrain the possible evolutionary tracks followed by the model disks. These comparisons may provide the first solid clues as to how disk galaxies form and evolve.

Gas and Dark Matter in Cosmological "Experiments"

Raymond Carlberg
University of Michigan

BACKGROUND

Almost all of the visible mass of the universe is in galaxies, yet we know very little about galaxy origin and evolution. The growing power of computers, however, is allowing astronomers to perform realistic, and otherwise impossible, "experiments" in galaxy formation. These simulations provide precise predictions that will soon be tested as large ground-based and space telescopes begin to allow us to see the universe as it was at the time of galaxy formation. The "experiments" are no replacement for observation but do offer an important advantage: millions of years in the evolution of a galaxy can be simulated in just a few hours on a supercomputer.

Most of the mass of the universe is not visible; astronomers call this nonluminous mass "dark matter." The properties of the unseen dark matter are likely to be the main determinant for the properties of galaxies and their distribution in space. Most methods for measuring the amount and distribution of dark matter depend critically on the assumption that galaxies are simple, unbiased tracers of the underlying mass distribution; that is, the distribution of galaxies reflects the distribution of dark matter. Ideally, one would like to take a set of proposed cosmological initial conditions, work out where the galaxies form and what they look like, and then compare this "experiment" to the observed universe. Completely dissipationless numerical simulations that describe the dark matter alone do this and have become sufficiently refined to provide considerable insight into the large-scale structure of the universe.

THE QUESTION

So far, however, there has been little numerical work that includes the gas that leads to visible galaxies in cosmological evolution. This is due, in part, to the extremely heavy computational demands of such a study. The conversion of gas into stars that make a visible galaxy is an extremely complex process. However, a number of studies have shown that gross properties of dissipation and star formation may be adequately described using equations that give averages over large regions, thereby reducing the computational complexity. In dissipationless simulations, a severe problem called "overmerging" occurs; as clustering develops, all of the substructure of the group is wiped out because dissipation is not included. Part of the solution is to include gas that can dissipate and shrink to a smaller size as it forms into galaxies.

COMPUTATIONAL TECHNIQUES

A test case for a cosmology, including gas and dark matter, is being studied using the Cray X-MP at the Pittsburgh Supercomputing Center. The numerical code to do these simulations is a hybrid of a gravity scheme and a "superatom" gas code. First, a large scale gravity field is calculated using a Particle-Mesh (PM) technique on a 64^3 grid. Since the mesh underestimates gravity for nearby particles, a second pass is made, the Particle-Particle (PP) step, in which neighboring particles are individually given additional forces to bring their attractions back up to the full Newtonian value. A small softening is added to avoid problems for very close particles. The PM step is easily vectorized to take best advantage of Cray speed, but full vectorization of the PP step is impossible. Extensive use of special Cray subroutines, however, increases the speed of the PP portion of the program by a factor of five or so.

The gas is then treated as if it were atoms. That is, particles are examined first to find neighbors, then tested to determine whether they are approaching one another, and then to see if they are close enough to collide in the next time step. A collision can be either "adiabatic," a perfect bounce along the line of centers, or "isothermal," where the outgoing velocity is randomly drawn from a distribution corresponding to a predetermined temperature. The isothermal gas is the simplest scheme that allows an effective dissipation to be included in the simulation. As the visible gas and dark matter are traced through their

evolution, a "snapshot" is taken of the distribution of matter every few time steps. The computation time on the Cray is approximately 20 CPU seconds near the beginning of the simulations; but as clustering develops, each step may take as long as 200 seconds.

RESULTS

The results of such a numerical "experiment" for the formation of a stable cluster of galaxies are shown in the accompanying diagrams. The success in producing this cluster is partly a consequence of the particular choice of initial conditions. The initial cosmological density perturbations were chosen according to the popular "Cold Dark Matter Spectrum," which enhances power on long wavelengths relative to white noise. The Cold Dark Matter Spectrum is the form of the initial distribution of matter (that is, its "lumpiness") predicted by theory during the first one-millionth of one second in the Big Bang.

When these simulations are "observed" in the same way that the universe is observed, the galaxies tend to be much more clustered than the dark matter, giving a severe underestimate of the total amount of mass present. Thus, the assumption that the distribution of galaxies accurately reflects the distribution of total mass in the universe may not be valid. Furthermore, these studies indicate that the relatively low mass density that is inferred from visible matter may underestimate the total density by such a large degree that the true density is actually in agreement with predictions from the Inflationary Universe Theory. (This theory predicts that the density is just the value needed to slow the expansion of the universe to a complete and permanent stop.)

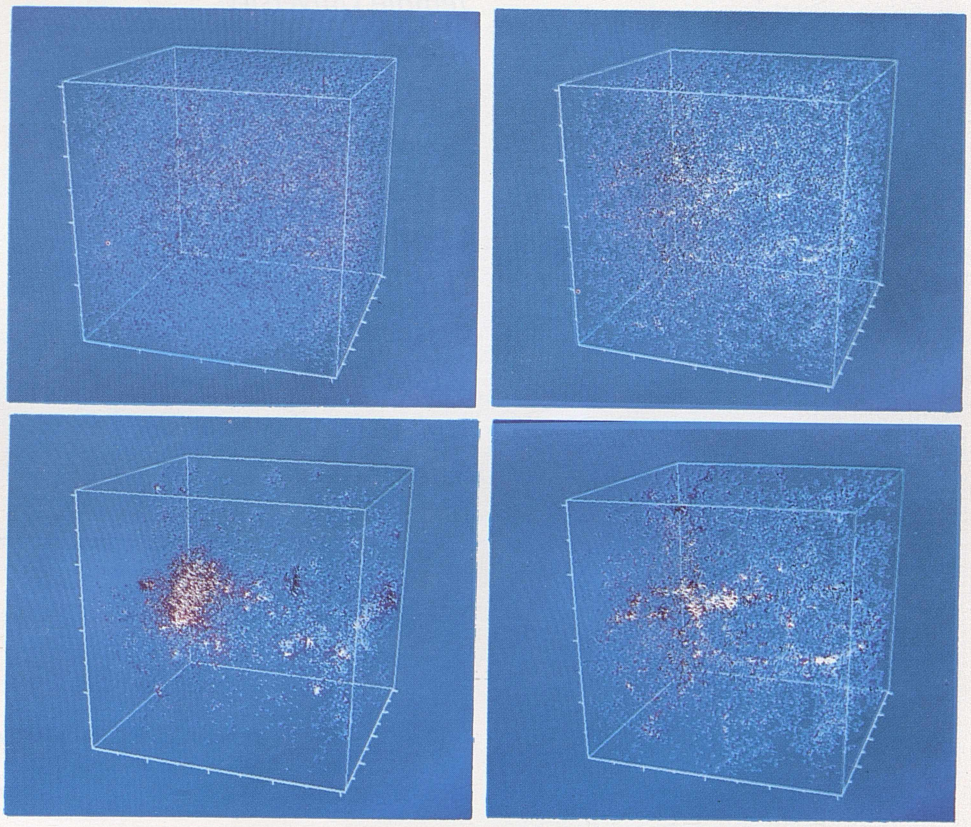

Four stages of large-scale galaxy formation are shown beginning with a time early in the evolution of the universe (top left) and proceeding clockwise to a time late in the evolution (bottom left). The "dark matter" is shown in red; the visible gas in white. This simulational study clearly shows that as time proceeds, the visible gas forming galaxies exhibits much stronger clustering than does the dark matter.

The Academic Affiliates of the Pittsburgh Supercomputing Center

Carnegie Mellon University
Case Western University
University of Florida
Georgia Institute of Technology
Johns Hopkins University
Lehigh University
University of Maryland
University of Michigan
Michigan State University
Northwestern University
Ohio State University
University of Pennsylvania
Pennsylvania State University
University of Pittsburgh
University of Rochester
Temple University
Triangle Universities:
- *Duke University*
- *North Carolina State*
- *University of North Carolina*

University of Tennessee
University of Virginia
Washington University
West Virginia University
University of Wisconsin
Yale University

Vivian Benton: Production Editor
Joel Welling: Computer-Generated Graphics Specialist
Penny Sackett: Scientific Editor
Jeff Lunderstadt: Graphic Designer

The Pittsburgh Supercomputing Center would like to thank all those who aided in the production of this booklet, with special thanks to the contributing researchers.

All of the full-color figures in this booklet were prepared using graphics devices at the Pittsburgh Supercomputing Center.

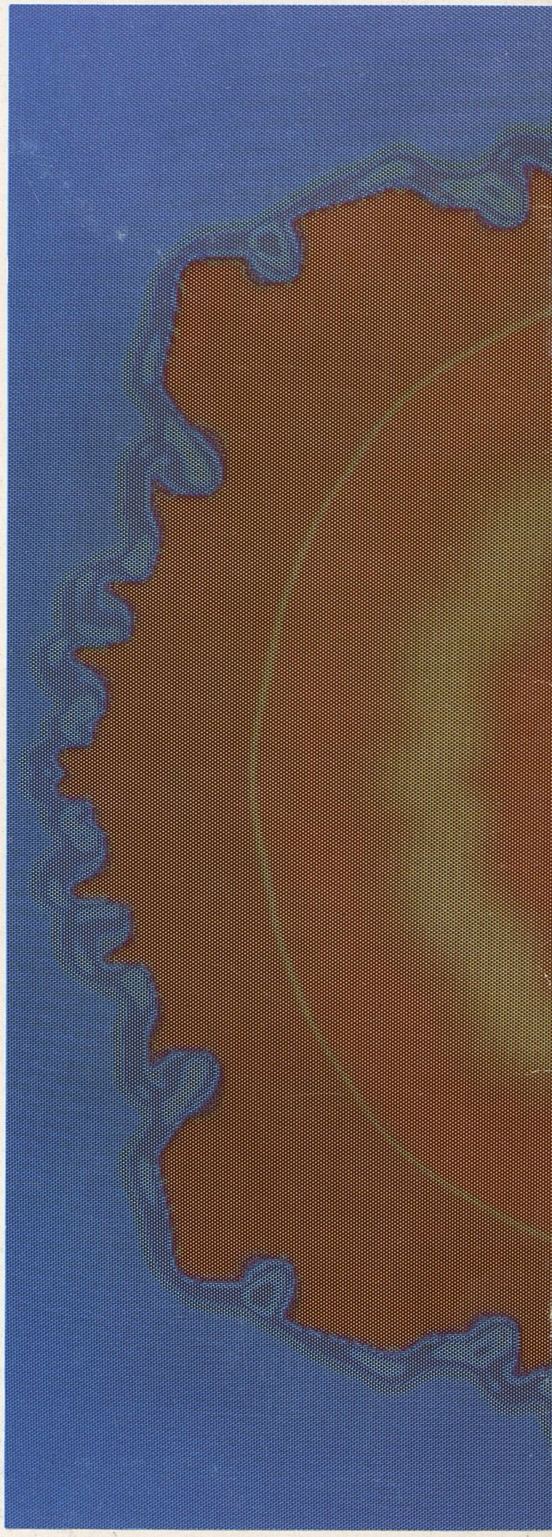

Pittsburgh Supercomputing Center
4400 Fifth Avenue, Pittsburgh, PA 15213-2683
412/268-4960