CRAY CHANNELS

SPRING 1990 · A CRAY RESEARCH, INC., PUBLICATION

Structural analysis

CRAY CHANNELS

In this issue

Building prototypes, subjecting them to loads, and measuring the results. This is the traditional method by which engineers test the structural integrity of manufactured parts and products. If a part does not perform to standards, this costly and time-consuming process is repeated until acceptable results are achieved. Today, however, many manufacturers are using Cray Research supercomputers to analyze structures more efficiently. By modeling the structural properties of parts on Cray systems, designers and engineers can modify and retest the parts many times, quickly and inexpensively, without building prototypes.

In this issue of CRAY CHANNELS, we look at the latest developments in crashworthiness analysis, an application of structural analysis in the automotive industry for which supercomputers have proven invaluable. We also report on progress in modeling sheet metal formation at General Motors and on research at Lawrence Livermore National Laboratory into the stress behavior of composite materials, along with a look at some of the largest structural analyses ever performed. In addition, we profile the Perfect Benchmark Suite, a new benchmark of application codes that offers a more comprehensive look at computer performance than traditional benchmarks. Our regular departments describe the latest version of Cray Research's UNICOS operating system, engineering analysis programs for use on Cray Research systems, and a project to help minicomputers keep their cool.

Determining the structural properties of a product's components early in its design cycle will bring the product to market faster and at lower cost. By providing this capability to manufacturing industries, Cray Research systems are helping companies enhance their competitiveness and profitability.

Features

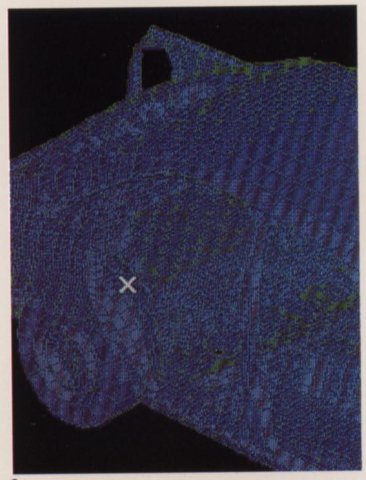

2

Departments

CRAY CHANNELS is a quarterly publication of Cray Research, Inc., intended for users of Cray Research computer systems and others interested in the company and its products. Please mail feature story ideas, news items, and Gallery submissions to CRAY CHANNELS at Cray Research, Inc., 1333 Northland Drive, Mendota Heights, Minnesota 55120.

Volume 12, Number 1

Editorial staff
Ken Jopp, editor
Elizabeth Knoll, associate editor

Design and production
Barbara Cahlander
Eric Hanson
James Morgan
Cynthia Rykken

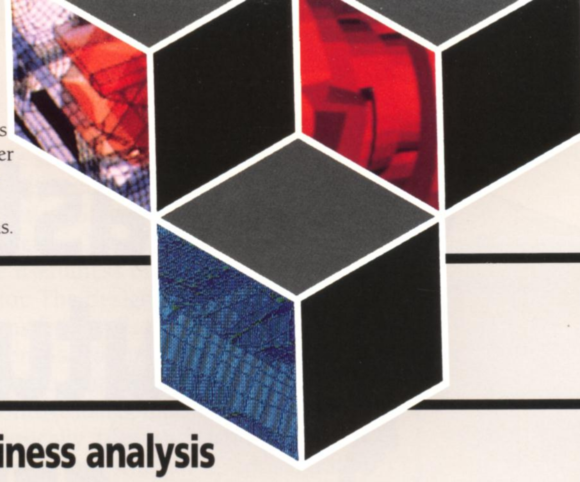

On the cover: Industrial designers and engineers are able to maximize the strength and efficiency of numerous products by performing structural analyses on Cray Research computer systems. Cray Research systems are used extensively in the automotive and aerospace industries to perform analyses that would be impractical to run on other computer systems.

Past, present, and future of industrial crashworthiness analysis

Paul Du Bois, consultant, Offenbach, West Germany

Porche 904 crash tested the old-fashioned way.

During the past decade, numerical crashworthiness analysis has evolved from a theoretical possibility to one of the primary applications of supercomputers and numerical methods in the automotive industry. The evolution of this type of analysis is marked by several noteworthy milestones, though the full potential of the method has yet to be realized. Without trying to give a full historical review of the evolution of crashworthiness analysis, I hope to convey an appreciation for the more significant developments that have occurred along the way and will try to identify some of the more general trends that characterize this kind of work.

The crashworthiness of newly designed vehicles traditionally has been verified by conducting a series of physical crash tests. Such tests, however, require prototypes that are time-consuming and expensive to manufacture. Also, they can be performed only relatively late in the design cycle. Consequently, design changes suggested by such testing are expensive to implement. Therefore, crashworthiness predictions obtained by numerical simulation are extremely valuable, even if they are obtained only a few months before the physical test.

Historical overview

Driven by the obvious time- and cost-savings potential of numerical simulation, automotive engineers performed simulations of vehicle crashes as early as the 1970s, primarily in the United States. The various approaches included beam models, mixed models, and finite element models. These approaches, however, all were limited in size to a few hundred elements due to the limitations of the available computing power. Because such models could accommodate only a small number of degrees of freedom, realistic crashworthiness predictions of full vehicles proved impossible.

From 1983 to 1985 the Arbeitskreis, or "working group," of the West German automotive industry sponsored two projects to investigate the feasibility and reliability of full-vehicle crash simulation. Simulations were run using explicit finite element codes on models that represented the front half of a vehicle with approximately 8000 shell elements. The projects modeled frontal crash of the front half of a car driven into a barrier by a sled. The vehicle models were furnished by Volkswagen and BMW, and the numerical simulations were performed using the PAMCRASH and CRASHMAS codes, respectively.

The usefulness of explicit finite element techniques still was questioned by many in the industry at this time. Although these methods proved to be computationally superior to the more traditional implicit approach for simulations of short-duration phenomena, at the time they were not practical for simulations of longer duration phenomena, such as automotive crashes. Short duration in this context meant a few milliseconds; an automotive crash, by contrast, typically lasts 100 milliseconds. Explicit methods had been applied most effectively to simulations of hypervelocity impact, such as from meteorites, and high velocity or ballistic impacts. Prior to the availability of supercomputers, explicit methods would have been impractical for crash simulations due to long run times. A crash simulation might require up to 100,000 explicit time steps.

During the mid-1980s, automotive engineers used supercomputers for the first time in Europe to conduct crashworthiness analyses. The engineers turned to supercomputers because the simulations required several thousand hours of CPU time on a Digital Equipment VAX 750 computer, a computer typically used for such applications at the time. The use of vectorized codes on a CRAY-1 supercomputer improved the performance of these applications by a factor of 200 with respect to a VAX 750 computer, and this performance doubled again when engineers exploited the hardware gather-scatter facilities of CRAY X-MP computer systems. This performance improvement reduced the typical run time for a crash simulation to a few hours. Thus, for the first time, a realistic crash simulation on a full vehicle model was feasible overnight.

Similar performance improvements were experienced in the United States, where the DYNA3D program from Lawrence Livermore National Laboratories was applied to the simulation of frontal car crashes.

During this time, the major drawback of the newly accepted explicit methodology was the accessibility of the results. A good understanding of the results of a typical crash simulation required the production and evaluation of many hundreds of plots, and interpreting this data required days, or even weeks. A solution to this problem was provided in 1985 at the first Cray Research-sponsored automotive conference, in Zürich, Switzerland, where Cray Research and Peugeot SA presented the first animated graphics of the Citroën BX frontal crash simulation. The use of animation allows engineers to evaluate their results quickly, and since 1985, much progress has been made in automating the display of results and in the fast production of animations. Engineers today typically hope to absorb the analysis results within one working day, which requires the ability to perform overnight computing.

Since 1985, dozens of frontal impact simulations have been performed in Europe, Japan, and the United States. Gradually, such analyses evolved from qualification tests that were performed on existing vehicles for comparison with test results to real-life production runs on prototype models, which enable engineers to check a structure's buckling modes before the structure is built.

In the increasingly stringent legal environment of recent years, a great diversification of crash simulation technology has taken place. For example, in 1987 Ford Motor Company in the United Kingdom was running side-impact simulations using a barrier proposed as a European standard. Many manufacturers followed this example and simulated various proposed side-impact standards such as the European CTP and EEVC standards and the U.S. NHTSA standard. Peugeot SA first performed skewed frontal impacts of under 30 degrees early in 1988, and in the same year Ford Motor Company in the United States was simulating light truck crashes, and Volvo was simulating the rear impact of passenger cars using a rigid barrier.

Automotive engineers typically consider modeling of the entire vehicle to be important in crash simulation because modern vehicles, once assembled, are difficult to divide structurally into discrete components. The buckling behavior of a front rail, for example, largely will be a function of the rail's connection to the surrounding structure. This situation, however, does not undermine the usefulness of optimizing single component behavior. The design of energy-absorbing components, such as front rails and even frontal car assemblies, has given rise to many new applications of numerical simulation that run parallel to full vehicle analysis. Models of individual components are relatively small, containing typically a few thousand elements, but the computational effort still can be considerable due to the need for contact-detection in the so-called folding problems. At Ford in the United States, the numerical crash simulation of front rails was used as a fully functional design tool as early as 1986.

Trends

From the beginning, numerical crashworthiness analysis has been a tradeoff between feasibility and accuracy. Two trends of the past decade continue to affect this tradeoff. These are a reduction in the modeling effort due to automatic mesh generation and the refinement of meshes to increase the accuracy of the computed results. Both trends are contributing to a rapid increase in required computing capacity.

Despite many advances in the field, finite element models of vehicles still suffer from several major drawbacks:

- ☐ The mesh is nonuniform; zones that are expected to deform heavily are represented much better than zones where little deformation is expected to occur. As a result, the analyst has to anticipate part of the answer, which in many cases requires considerable experience and engineering intuition.
- ☐ Usually a very fine representation will be used for component analysis, which may not be affordable in the full model. As a result, the modeling effort cannot be passed on to full vehicle simulations, and thus the results are less accurate than those from component models.
- ☐ Due to the limitation set on the total number of elements that can be used practically for a full body analysis, a fine representation is possible for approximately half the car. Consequently, very different numerical models must be built to investigate multiple load cases such as frontal, side, and rear impact.
- ☐ The need to create an economical mesh and save elements wherever possible puts severe limitations on the use of automatic mesh generators. Therefore, considerable skill and patience still are required of the modeler.
- ☐ The complexity of the meshes makes them hard to adapt to design changes, hindering parametric studies.

Over the past five years the compromise between feasibility and accuracy has become less critical as typical vehicle models have grown in size from 10,000 to more than 20,000 finite elements. However, models with on the order of 60,000 elements would be necessary to overcome most of the problems noted above. The computational needs for executing

Prior to the availability of supercomputers, explicit methods would have been impractical for crash simulations due to long run times.

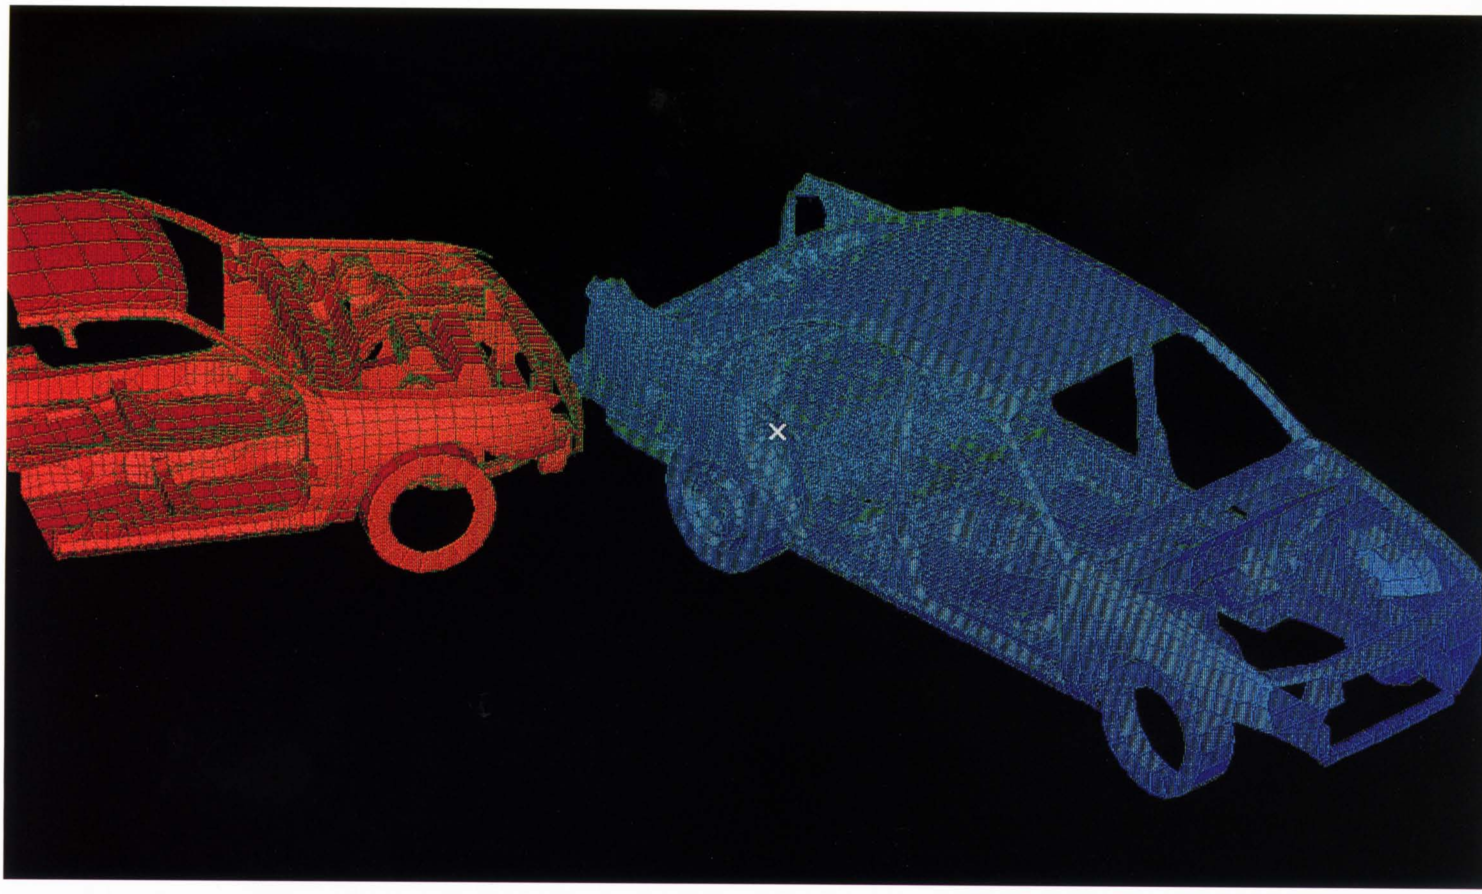

Figure 1. Simulation of car-to-car collision performed at Adam Opel AG on a CRAY X-MP computer system. The models are shown in their initial positions.

such models can be determined by basing the calculations on the operation count for a Belytschko/Tsay shell element with three integration points, as given by Hallquist: 560 + 3*33 = 659.

If we then assume 60,000 finite elements and a simulation with 100,000 time steps, we obtain 3.954^{12} operations to be performed. Such an undertaking would require a processing speed of 549 MFLOPS to complete within two hours or 183 MFLOPS to complete within six hours. Attaining the second of these numbers is feasible today, and even the first is within reach, on a multiprocessor CRAY Y-MP system.

Present applications

Thanks to refined numerical methods and more powerful supercomputers, the methodology of frontal crash simulation is finding its place in the design cycle. Once a car body is modeled, crash simulation can deliver the global collapse mode of the structure, the acceleration levels at the driver and passenger locations, and can even predict the response of air bag sensors. This information can be used to analyze important phenomena that essentially are uncoupled from the car body response itself. These phenomena include the motion of the steering column, the response of the air bag, and the response of the dummy. This is possible because much experience has been gained from tests over many years. These tests have revealed which interactions are important and which are not. Unfortunately, the same kind of experience is not available for all other load cases, such as side impact, to mention only the most critical example.

Even with the remaining limitations, numerical crashworthiness analysis can deliver extremely valuable results. The following examples describe some of the more advanced work being done in this field in the automotive industry.

Code developers are working to include the options necessary to simulate all the crash-related effects that do not directly concern the car body. Composite material models have been included in simulations to represent deformable barriers, special joint elements have been included to simulate steering columns and dummies, and arbitrary Lagrangian-Eulerian (ALE) techniques have been incorporated to allow air bag simulation. These developments have enabled engineers to extend their modeling capabilities far beyond the earlier goals of studying a structure's global buckling modes.

For example, engineers at Adam Opel AG, in Rüsselsheim, West Germany, are performing car-to-car collision simulations. The simulations (Figure 1) require detailed modeling of two full vehicles and an extremely accurate description of all possible contacts between them, including wheels and components. The model in Figure 1 comprised approximately 30,000 elements and required up to 70 hours of CPU time to solve on a CRAY X-MP system. This huge analysis, however, was performed to study the behavior of, for example, the fuel tank and other components in the target car, and examine its integrity during impact. Excellent agreement with experiments was obtained, demonstrating the accuracy of the simulation for this very local (as opposed to global) result. Car-to-car collisions regularly are simulated on the Cray system at Adam Opel using the RADIOSS finite element code.

New air bag model unfolds

As most major car companies introduce air bags as options, understanding the dynamics of air bag deployment is becoming critical to the auto industry. Cray Research engineers and the developers of the DYNA3D finite element program recently developed one of the most complex air bag deployment simulations of its kind. Using the DYNA3D code on a CRAY X-MP/464 computer system, John Hallquist and Doug Stillman of Lawrence Livermore Software Technology Corporation, and Mike Long of Cray Research have created a 13,000-element generic model of the deployment and inflation of an air bag. This is the first large-scale simulation that accurately models the second-order effects of wrinkling, creasing, and tearing due to overinflation. Because the finite-element model is symmetrical, it was necessary to model only one-fourth of the structure.

Upon automobile impact, an air bag is deployed from the steering column and inflated with gas within a split second. Air bags must be stout enough to stop the impact of a driver's chest against the steering column, while strong enough to withstand the pressure of inflation and the force of a human striking the surface.

Now that the researchers have developed an accurate model of the deployment and inflation of a generic automotive air bag, they are adding an additional variable to the model — the force of a dummy hitting the air bag. This significant variable will boost the model size to approximately 50,000 dynamic elements. The original model required three CPU hours on a CRAY X-MP system, and the new 50,000-element model will require approximately 10 CPU hours.

This research is part of Cray Research's effort to optimize the DYNA3D code for CRAY X-MP and CRAY Y-MP computer systems. The latest optimization works well on Cray Research's vector architecture — running an average of 25 percent faster than previously optimized and vectorized versions.

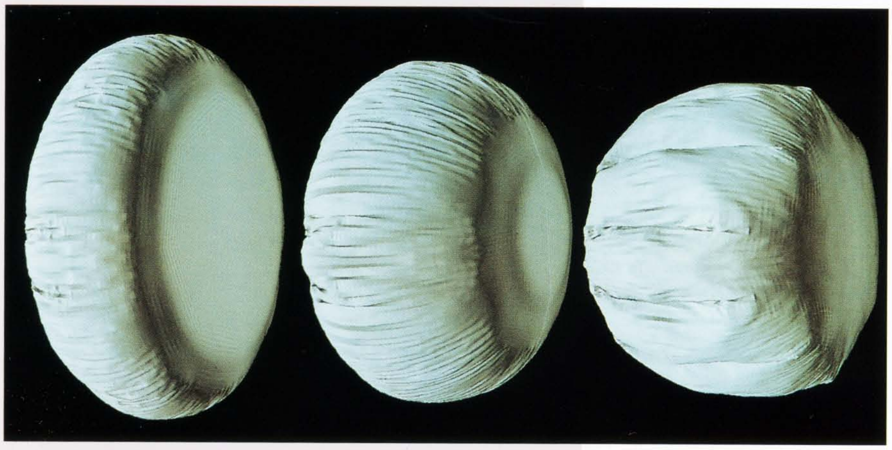

Three stages of air bag deployment and inflation. This is the first large-scale air bag simulation that accurately models the second-order effects of wrinkling, creasing, and tearing.

Another example of leading-edge work in this field involves side-impact simulation. Legislative proposals that set standards for side impact crashworthiness focus on acceleration levels at the dummy's pelvis, thorax, and head. These quantities cannot be obtained using dummy simulations that are uncoupled from the structural analysis, as in frontal impact studies, because considerable interaction occurs between car, door, and dummy during impact. To tackle this problem, engineers at the Ford Motor Company in Dunton, United Kingdom, included a finite element model of the Ford side-impact dummy in the Sierra side impact simulation (Figure 2). Again, very valuable information was obtained from these simulations concerning the dummy values, and this information was obtained within the context of a simulation comprising over 10,000 finite elements.

Applications such as these create a situation in which numerical simulation can be used as a design tool. These applications allow engineers to simulate real-life situations and obtain reasonable answers to very specific (legal) questions long before the actual hardware exists.

The value of numerical simulations in crashworthiness depends directly on the reliability of the results and the readiness of their availability. Both requirements reflect the need for a hardware environment capable of absorbing very large finite element models. Such an environment requires an integrated architecture of one or more supercomputers that allows overnight or faster computation and an environment of software and workstations that allows quick visualization of the analysis results and modification of the models.

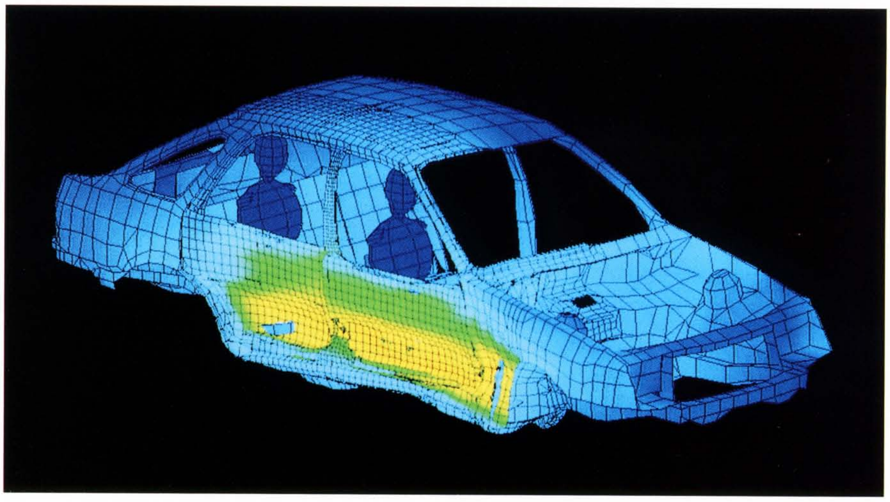

Figure 2. Simulation of a side impact with dummy performed at Ford Motor Company in the United Kingdom.

Acknowledgment

The author thanks Aruna Tilakasiri of Ford Motor Company and Ernst Nalepa of Adam Opel AG for their help in preparing this article.

About the author

Paul Du Bois is an independent consultant to the automotive industry. He received an undergraduate degree in civil engineering in 1979 and an MBA degree in 1980, both from the University of Gent, Belgium. He has worked for Engineering Systems International (ESI), where he gained experience with nonlinear finite element analysis. He was among the first engineers worldwide to promote the use of explicit integration techniques in automotive crashworthiness problems.

Sheet metal forming simulation using finite elements

Thomas B. Stoughton and Frank J. Arlinghaus
General Motors Research Laboratories
Warren, Michigan

Researchers at General Motors Research Laboratories have developed a finite element code that runs on Cray Research computer systems and enables them to simulate the three-dimensional deformation of sheet metal, a valuable capability in die design. These simulations permit designers to optimize and test designs before the dies are built, shortening lead times and lowering tooling costs, which are major bottlenecks in new car development programs. Computer optimizations enable GM to improve design performance during production by reducing scrap rates and press down times.

The formation of sheet metal automotive components is a multibillion dollar industry in the United States. Half of the costs associated with this industry can be attributed to producing and testing tool dies. This is because virtually every sheet metal forming die must be tested until a part is formed to the specified shape without splitting or buckling the sheet metal. Ideally, this trial-and-error process involves modifying the boundary forces to obtain the appropriate amount of resistance to metal flow. On average, this process takes several months for each major die.

The current method of die design, which emphasizes physical testing, accounts for the high overhead costs of tool dies and for the long lead times associated with the release of new products. Consequently, short-term solutions often are used during tryout, such as the application of special lubricants, which during production often are found not to work due to slight changes in the properties of the incoming material or other process variables. Such a situation increases the rates of scrap and eventually shuts down the production line. Thus, die designers would benefit from having analytical tools to test and prove their designs before the dies are cast and to optimize these designs so that the resulting panels would be insensitive to process variations.

Computer optimizations enable General Motors to improve design performance during production by reducing scrap rates and press down times.

The process of sheet metal forming

Developing analytical tools for use in die design is not a trivial task. The nonlinearities of tool geometry, the plastic deformation of sheet metal, and the development of boundary forces demand an incremental solution to track the process from the undeformed sheet to the final part shape. Any other type of analysis or process of simulation is incapable of predicting the consequences of the nonlinear path-dependent nature of the problem.

The finite element method (FEM) is ideally suited to this problem. To be successful, the application must include a realistic model of the material behavior and be able to simulate the forming conditions found in dies. These attributes will enable the finite element method to predict accurately what will happen to sheet metal in a given die. To be useful in die design, however, the analysis also must be fast.

Wang and Budiansky developed a finite element computer code for sheet metal forming applications in simple dies.[1] Although the accuracy of this code was impressive for axisymmetric applications, the cost of computation was a strong deterrent against applying the code to production problems. Software optimization and hardware upgrades removed cost as an issue for simple production problems.[2] Development also proceeded in the area of improved material models, including strain-rate sensitivity, bending effects,[3] and boundary conditions.

Model of the material behavior

To simulate sheet metal forming with finite elements, one first creates a model by dividing the sheet metal into triangular elements, each composed of three nodes. The equations for this model are given in rate form, so that they are integrated from one time step to the next. The step size is chosen to minimize errors due to the linearization of the rate equations after they have been solved.

The basic equations are derived from the theory of membrane shells.[4] A constitutive relation between the components of the stress and strain tensors for finite deformation of three-dimensional elastic-plastic solids under plane stress and modified to account for normal anisotropy is taken from Hutchinson.[5] These relations are simplified further by allowing the material to have a specific elastic normal anisotropy related to the plastic behavior. Such simplification permits easy inversions of the constitutive relations and introduces negligible error. In the original model, the material was assumed to satisfy the Remberg-Osgood relation:

$$\varepsilon = \frac{\sigma}{E} + \left(\frac{\sigma}{K}\right)^{1/n},$$

where E is the Young's modulus, n is the work-hardening exponent, and K is the tensile strength from a power law fit to the tensile test data. Because strain rate effects are important in the forming of steel products, some work has been done to incorporate rate-sensitive relations between the stress and plastic strain.

Further work is required to generalize the constitutive relation used in this model to account for in-plane anisotropy and nonisotropic or kinematic work-hardening. Although low-carbon steel exhibits in-plane anisotropy, the effect is small and usually plays a secondary role in its formability.

The original formulation was based on membrane theory because it was thought that this was acceptable for most sheet metal forming applications. However, tests of the model against laboratory experiments showed that bending effects begin to play an important role in the flow of metal when the sheet thickness exceeds about 10 percent of the tool radius.[3] Because the radii of die designs range from 2-10 mm, bending effects often are an important factor in the formability of sheet metal panels. The applicability of the membrane model can be extended to cases in which the sheet thickness exceeds 30 percent of the tool radii with a negligible increase in the computation costs.[3] This improvement is made by separately calculating the force due to bending that resists the flow of metal and by modifying the force rate equations to account for its absence in the membrane formulation.

Model of the tool surfaces

Describing the tool surface is one of the most important and difficult problems in sheet metal forming simulation. Not only must the tool surfaces be capable of providing a z coordinate for a given x and y, but a rate formulation also requires the first and second derivatives of the surface. These parameters are required to impose constraints on the flow of metal and translate the developing nodal forces. A single metal forming analysis on a production die may require on the order of a hundred million surface evaluations. Finally, to be useful in die design optimization, the surface must be easy to modify.

These requirements are difficult to satisfy for three-dimensional surfaces generated from a clay model of an automobile. In an attempt to extend the surface description capability to arbitrary shapes, the surface was described by interpolating between scan line data using B-spline functions. Although this technique was fast and reproduced the desired surfaces with reasonable accuracy, it exhibited oscillations on the surface when the data were too sparse or when the surface took a sudden jump in value. These problems could be solved by increasing the amount of data in the scan lines, and although this method was capable of exceeding the computer's available memory quickly, it nonetheless was applied to production dies.[2]

A more severe problem with the B-spline interpolation method is that the surfaces are extremely difficult to modify. Changing a radius requires a change in all data points in the vicinity of the radius and in each scan line that crosses the radius. For this reason new methods of describing tool surfaces currently are being studied. The most promising method appears to be rational bicubic polynomial B-spline functions used as interpolating functions, combined with a contour-delineated surface blending function. This new approach permits the user to modify the surface rapidly during a pre-processing stage of analysis, and later allows the code to calculate the surface with sufficient speed and accuracy.

Model of friction

The model developed by Wang and Budiansky calculated the normal force at each node in contact with the tool surface as well as the tangential component of the force.[1] As long as the node continued to slip across the tool surface, the magnitude of the tangential force was proportional to the normal force by a coefficient of friction. Unfortunately, this coefficient is not well-known for metal forming applications and, in fact, may not be a constant at all.

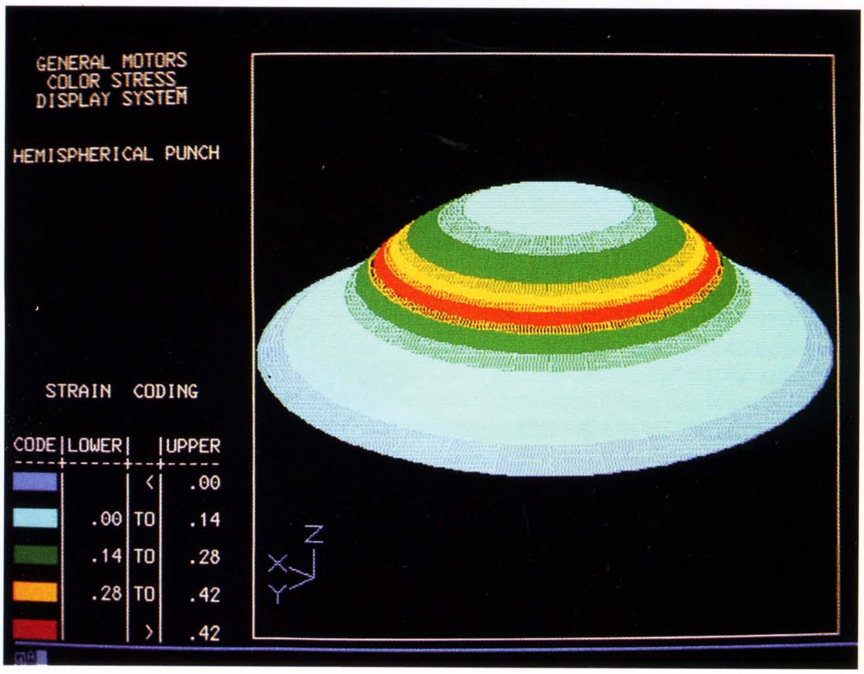

Analysis of hemispherical geometries show that the punch depth at failure and the location of the failure (shown in red) are strongly dependent on frictional conditions. As the coefficient of friction is lowered from 0.25 to 0.00, the depth at failure increases, and the location of failure moves toward the pole of the geometry — observations that are confirmed by the physical experiment.

Thus, one would expect that a model of the sheet metal forming process would include a model for the coefficient of friction that describes how it varies as a function of metal deformation and lubricant thickness. Although this eventually may be discovered to be necessary, usually a single coefficient of friction is adequate to reproduce the experimental results in all of the tests conducted on the FEM model to date. The only sophisticated feature of the model with regard to the coefficient of friction is that the user can specify different coefficients for each tool surface. This allows the simulation of effects due to differences in the lubricant applied to the two sides of the sheet metal, differences in its surface characteristics, as well as differences in the surfaces of the tool itself.

Implementing a model of frictional forces requires more than equations relating the normal and tangential forces using a coefficient of friction. Often in metal forming applications the assumptions of Coulomb friction do not hold. For example, sometimes the normal force is so large that the expected frictional force exceeds that required to stop the flow of metal. Obviously, this condition can occur in local areas of the panel while in other areas the metal continues to flow over the tool surfaces. To detect this condition the model keeps track of the velocity of the metal at each point from one time increment to the next. When the direction of this velocity reverses, that point on the metal is assumed to stick to the tool surface for that and succeeding time step increments.

Model of the boundary conditions

Original applications of the model were limited to stretch-form simulations.[1,2,3] Extending the capability of the model to processes in which the boundary is not held fixed during the forming of the panel introduces three major difficulties for a FEM analysis. Because pure stretch-forming processes rarely are used, this extension is essential to the use of FEM analysis as a die design and tryout tool.

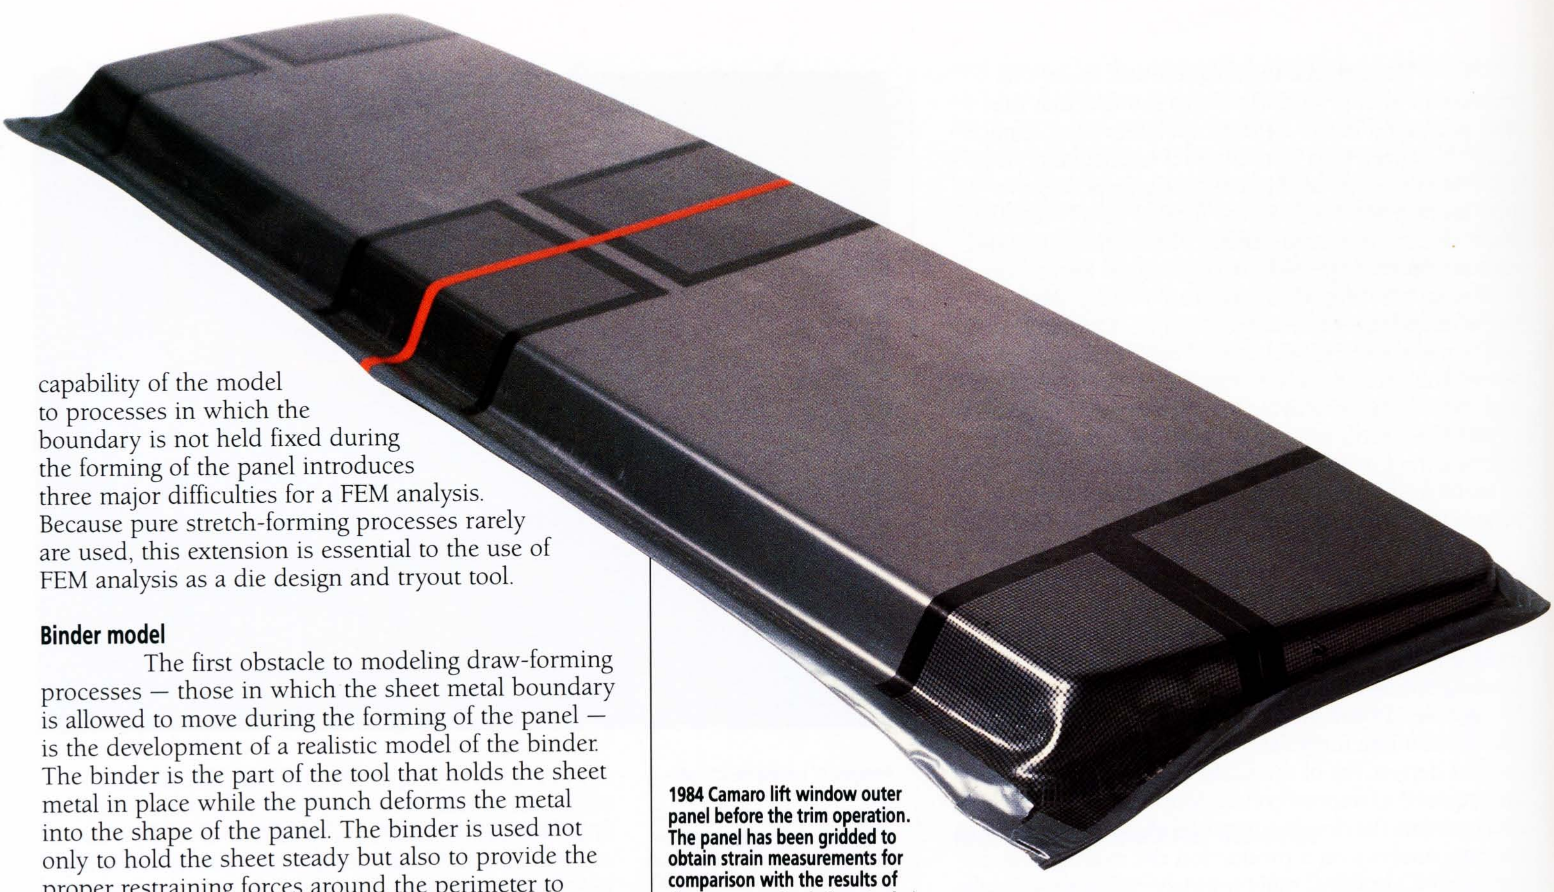

1984 Camaro lift window outer panel before the trim operation. The panel has been gridded to obtain strain measurements for comparison with the results of the computer simulation. Analysis was needed only for one half of the panel because of the symmetry of the subject.

Binder model

The first obstacle to modeling draw-forming processes — those in which the sheet metal boundary is allowed to move during the forming of the panel — is the development of a realistic model of the binder. The binder is the part of the tool that holds the sheet metal in place while the punch deforms the metal into the shape of the panel. The binder is used not only to hold the sheet steady but also to provide the proper restraining forces around the perimeter to control the flow of metal into the die cavity. If these restraining forces are too high, the sheet metal may split. If the forces are too low, the sheet metal may buckle. The major task of die tryout is to determine the optimum restraining forces at each point around the perimeter of the binder that do not cause the sheet metal to split or buckle.

Determining the optimum restraining forces from the binder requires a model of the binder forces. One mechanism for generating these forces involves applying a load to the binder. This load generates restraining forces on the metal in the binder through friction. The model for the binder load forces is complicated by the fact that as metal flows from the binder into the die cavity, the area of metal in the binder decreases. Depending on the initial blank shape and the variation in the amount of flow around the perimeter of the binder during the forming of a panel, these forces can vary in time and position in extremely complex ways around the perimeter. The model of the binder load forces must also account for the effects of these thickness variations, which may, in cases of extreme local thinning or thickening, transfer all the binder load to the thickest metal in the binder.

In addition to the binder load, drawbeads commonly are used to control metal flow from the binder. These drawbeads, which look like ridges on the binder surface, force the metal to deform as it flows across the bead. Models of the restraining forces generated by the drawbead have been developed to present these forces as a function of the drawbead geometry and material properties. Thus, instead of modeling the deformation of the sheet metal over the drawbead, the binder shape is smoothed in the drawbead area. The drawbead forces are calculated and applied as external forces in these locations.

Mesh generation

The second obstacle to modeling draw-forming processes is a set of additional requirements on the mesh. In stretch-forming applications, nodes that represent points on the sheet metal do not move very far during the forming of a panel — typically less than a few millimeters in the projection perpendicular to the direction of punch travel. This means that it is fairly easy to generate a mesh of nodes and to concentrate them in regions where high density is required, such as on sharp tool radii. However, in draw-forming the metal can flow 10 cm or more in this projection. Unfortunately, unknowns of the forming process include the questions of where the metal will flow and how much metal will flow into specific areas. Therefore, one cannot predetermine where nodes will be required on the original undeformed blank.

The best solution to this problem is adaptive mesh regeneration. This means that as the FEM calculation proceeds and metal begins to flow, the code automatically determines if the nodal density is too sparse in certain areas. When this condition occurs, a new mesh is generated and the calculation continues.[1] The difficulty with this approach is the danger of introducing artificial instabilities due to the algorithm that maps the elemental properties from one mesh to the next, and that may result in an erroneous result. A simpler solution is adequate for preliminary applications to draw-forming processes. This solution involves using the original tool surfaces to determine nodal density and generating a mesh using Delaunay triangulation.[2] During analysis the model determines an improved nodal density function, which is used to generate a new mesh to be used in a succeeding analysis. Although this method works, it is expensive. Nevertheless, this method works as fast as an adaptive mesh generator

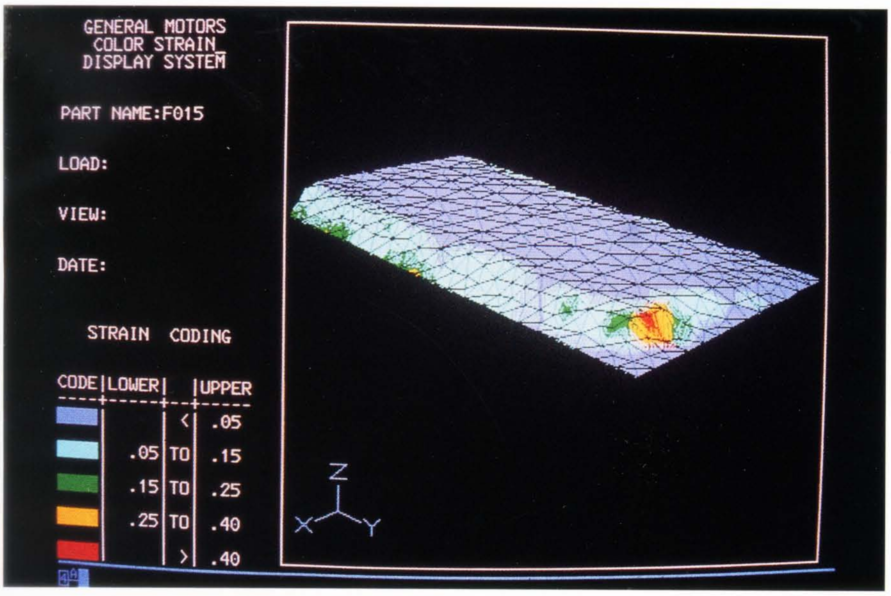

Analysis of the 1984 Camaro lift window outer panel predicted a tear in the corner during the stretch forming process (fixed boundary). This problem was discovered during an actual tryout using prototype tooling — before computer analysis was available. Computer simulation of this event would have saved great time and expense.

for the applications in which the amount of metal flow is small.

Computational efficiency

The last obstacle to modeling draw-forming processes is a set of additional computational requirements that can be attributed to two factors. The first is the observation that as boundary forces become smaller, metal is allowed to flow more. This results in an increased tendency to buckle, which creates instabilities in the equations for the nodal velocities. To limit the introduction of errors in the linearization of the equations, very small time step increments are required. This can increase the analysis time by more than an order of magnitude for cases of substantial metal flow from the boundary compared to a case of no metal flow. Fortunately, for a significant range of metal flow, the computational requirements are the same as the stretch-formed boundary conditions for some applications. Nevertheless, because the occurrence of buckling instabilities is application dependent, the cost of analysis can be expected to increase by an order of magnitude as the boundary forces are lowered during a search for the optimum forming conditions.

The second factor contributing to additional computational requirements of the draw-forming process is that the die designer must determine the optimum boundary forces around the entire binder that will form the panel without splitting or wrinkling the sheet metal. The designer essentially has an infinite number of degrees of freedom. The shape of the blank determines the distribution of the binder load, which factors into the distribution of the boundary forces. In addition, the number, location, and geometric characteristics of the drawbeads allow die designers to add and vary forces anywhere in the binder. These variables play a crucial role in determining the flow of metal when the boundary is allowed to flow freely during draw-forming processes; but none of the variables plays a significant role when the boundary is held fixed during stretch forming. With these infinite degrees of freedom, a method for automatically or systematically determining the optimum forming conditions is not obvious. Although such an improvement may be possible, it is not essential because die engineers have sufficient knowledge and experience to sift through these options fairly efficiently. Although die design optimization always will take longer for draw-forming applications than for stretch forming, proper integration of this new technology with the people and methods involved will minimize this cost.

Computational requirements

Although use of the finite element method in sheet metal forming requires large computational resources, this no longer is considered a serious obstacle to the utility of this method. However, applications of the model described here were limited to axisymmetric parts when the model first was developed on an IBM mainframe computer. At that time, an analysis of even a simple production die design required hundreds of CPU hours on an IBM-3081 computer. This was not a serious consideration because the nature of the FEM code made it ideally suited for vector and parallel computer architectures. By 1983, a restructured and optimized version of the original code was developed to run on a CRAY-1S computer system at General Motors Research Laboratories. This version was running applications 38 times faster than the original version of the code running on an IBM-3081 computer. Further optimization and a computer upgrade to a CRAY X-MP/24 system have resulted in substantial improvements. Depending on the panel complexity and on the type of process, real applications now are analyzed in times ranging from 20 minutes to 16 hours.[2] A performance increase of 30 percent is expected with GM's new CRAY Y-MP system.

Although these analyses are very useful, the efficiency of the computer code needs improvement. The largest panels and some complex inner panels that exhibit many details require meshes with 10 times more nodes than can be analyzed realistically today. Such an analysis would increase the computation time by several orders of magnitude. In addition, the times listed above are for a single analysis. To develop and optimize a design on a computer obviously requires

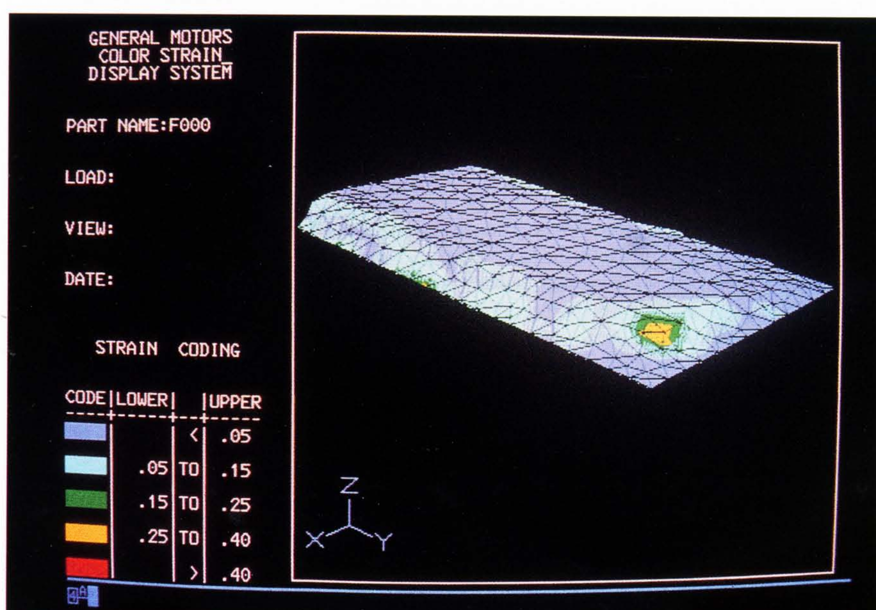

Analysis of the panel using ideal lubrication conditions shows that even under the best forming conditions, this panel could not be stretch formed. Panel formation would require the conventional draw process, in which the boundary is allowed to flow. Because this process requires a different type of press, the availability of such knowledge early in the design process greatly reduces the lead times and costs associated with new car programs.

several analyses. Unfortunately, the most expensive analyses — those using draw forming that results in buckling instabilities — also have the highest number of degrees of freedom. Thus the most expensive analyses are the ones that must be repeated often during optimization.

Fortunately, the FEM model is expected to work well on a parallel architecture system that has individual processors with vector capability. Only on these machines will we find the order of magnitude improvements required to see the full benefit of computer simulation of the die design process. The efficiency of the FEM model reduces significantly the costs of the conventional method of die design and tryout.

Experimental tests of the FEM model

Originally, the model was applied to the analysis of a hemispherical dome.[1] Although the analysis and the experiments were done by different departments at the General Motors Research Laboratories, the test was a collaborative effort. Determining the predictive accuracy of the model was difficult because experimentally predetermining the coefficient of friction is impossible. Thus, the predictions of the model were made for many values of friction and the value that resulted in the best agreement in the strain distribution was selected for each material tested. Although this may appear to be simple curve-fitting, the behavior of the strain distribution as the coefficient was lowered to zero reproduced the experimental observation that the peak strain shifts to the pole of the dome as the lubricity is improved. In addition, reasonable agreement was found in both the major and minor strain distributions for a given material using the same coefficient of friction. Finally, the value of the coefficient of friction was physically reasonable for the lubrication condition used. Although these observations are compatible with the predictive accuracy of the FEM model, they are far from a definitive test.

The next experimental test of the model was performed using the new version of the code running on a CRAY-1S system. In this test the model was applied to forming three rectangular pockets, each with a different profile-radius.[3] Again, because no means existed to determine the coefficient of friction experimentally, a value was selected that gave the best agreement in the strain distributions. These results are shown in Figures 1 and 2. Although this test also was incapable of establishing the predictive ability of the model due to the selection of the friction coefficient, the same compatibility with the observations was found as with the hemispherical dome test. The fact that the same coefficient was used for all three sets of tools also indicates that the FEM model was truly predictive.

Although these tests offer significant insight into the probable accuracy of the FEM code, a more definitive test was needed. In addition, the above tests were limited to stretch-forming applications and to small laboratory-sized specimens. For this reason the advanced engineering staff at General Motors conducted three experiments using a production press. These three experiments formed large panels that separately used the stretch- forming process, the draw-forming process with binder load, and the draw-forming process using drawbeads. The geometries were selected to deform the sheet metal over the whole gamut of final strain states near the forming limit of the material. This was a blind test; the modeling group was provided only with information about the material properties and the die geometry. From that information the predictions were made and returned to the advanced engineering staff for comparison with the experimental results, which were conducted independently. The result of this blind test was the strongest confirmation of the accuracy of the model. Not only were the strain distributions in excellent agreement for all three geometries, but the model correctly predicted the amount of metal flow from the binder for both draw-forming processes.

Figure 1. Comparison of numerical results using a friction coefficient of 0.25 with experimental major strain distributions for AK steel stretched over a rectangular punch with a 12.7 mm punch profile radius.

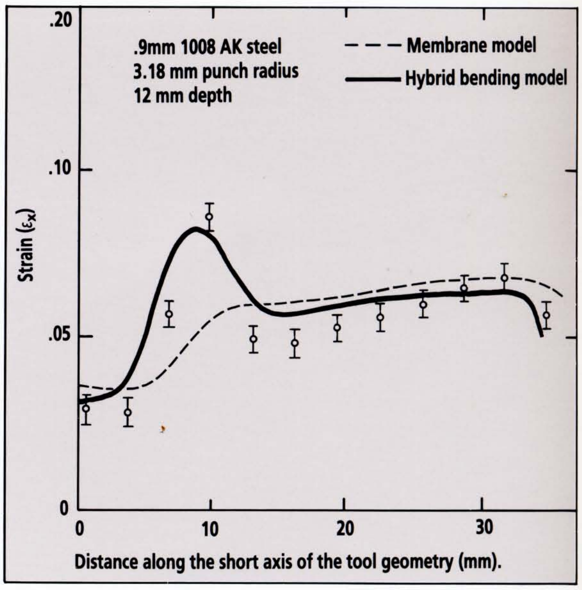

Today, in less than an hour we can determine answers that previously took several months of labor and a large capital investment.

Applications of the FEM model

Wide application of this FEM model to die design has been slowed by the lack of many of the capabilities described in preceding sections. Now that most of these capabilities are in hand, the model is being applied on a wider scale and to more difficult designs. The code now is being applied not only to determine the formability of new designs, but also to optimize design parameters of existing dies.

The first production applications of this model were not used to aid in the die design process, but rather to demonstrate the utility of computer simulation. The best example of this application is the 1984 Camaro lift window outer panel. This panel originally was designed to be stretch formed, which required special expensive tooling. During die tryout it was found after several months of trial and error that the part could not be stretch formed. The binder features that prevented the metal flow had to be ground down to allow sufficient metal to flow. In the end, the part could have been made by the conventional draw-forming process, but researchers had no way of knowing this at the time.

By 1985, computer technology had advanced enough to enable researchers to determine whether that panel could have been stretch formed. An analysis found that for a typical coefficient of friction of 0.15,

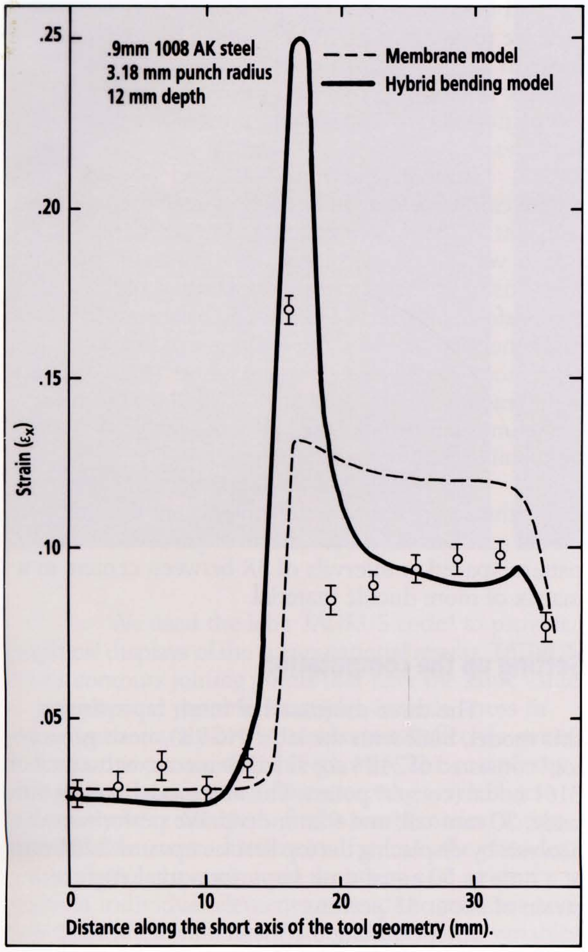

Figure 2. Comparison of numerical results using a friction coefficient of 0.25 with experimental major strain distributions for AK steel stretched over a rectangular punch with 3.18 mm punch profile radius.

the strains in the corner of the panel would be 40 percent over the forming limit of the material.[2] In other words, the part would split. A study was done of the strain distributions as a function of the coefficient of friction to determine if the panel could be formed using a better lubricant. The conclusion was still negative. Those analyses took approximately 20 CPU minutes each. Today, in less than an hour we can determine answers that previously took several months of labor and a large capital investment.

Conclusions

The FEM code described here is being applied to production to influence the design process. Although most of the analysis still is performed by research personnel, engineers in the production area are gaining proficiency with the code and soon will be using this analysis tool independently.

This analysis tool reliably predicts the behaviors of given types of sheet metal in given tool designs. The value of this information is obvious when viewed over the entire design and tryout industry. A die design that has passed the test of an accurate computer analysis virtually never will require a long die tryout. The costs and labor associated with the tryout process, which accounts for up to half of the cost of tools, will be reduced significantly. The cost associated with rebuilding a tool die after a tryout ends in failure should be eliminated. As the cost associated with tryout is reduced, so is the lead time.

One might wonder why tryout would be necessary at all. If a design passed a computer analysis, why wouldn't it produce a perfect part on the first try? Several reasons suggest this will not happen, at least not in the near future. Shrinkage during the die-casting process requires some trial and error compensation to produce a part within specifications. The machining of the die also leaves burrs and other irregularities that are discovered and polished off most effectively through tryout. Application of lubricant, fine tuning of the blank size, binder loads and other boundary forces, and compensating for the variability of the material properties are issues best addressed at the end of the design process.

This will not always be true. Today, computer analysis is a cost effective way to weed out poor designs that will not pass a conventional die tryout or will pass only after a great deal of effort. This is because the costs of computer analysis have dropped several orders of magnitude in the past 10 years, while the costs of labor, plant space, and other factors have increased. The costs of computer analysis will continue to drop an order of magnitude every five years, while the costs of conventional tryout will continue to rise. Someday it will be cost-effective to optimize die designs so that they are less sensitive to the process variables over which we have limited control. When this happens, a good panel may be formed on the first or second try. Finally, when that die goes into production, it will form one good panel after another, ideally never requiring adjustments to compensate for changes in the humidity or other variables over which we have no control. This virtually will eliminate die tryout, production scrap, and production down times due to forming problems. The potential annual cost savings will be millions of dollars.

About the authors

Thomas Stoughton and Frank Arlinghaus are staff research scientists in the Computational Physics Group of General Motors Research. Stoughton received a B.S. degree in physics from the California Institute of Technology in 1976 and a Ph.D. degree in high-energy physics from the Massachusetts Institute of Technology (MIT) in 1981. Arlinghaus received a B.S. degree in chemistry from the University of Detroit in 1962 and a Ph.D. degree in physical chemistry from MIT in 1966.

References

1. Wang, N. M. and B. Budiansky, "Analysis of Sheet Metal Stamping by a Finite-Element Method," *Journal of Applied Mechanics,* Vol. 45, pp. 73-82, March 1978.
2. Arlinghaus, F. J., et al., "Finite Element Modeling of a Stretch-Formed Part," *Computer Modeling of Sheet Metal Forming Process,* 12th Automotive Materials Symposium, The Metallurgical Society, April 29-30, pp. 51-64, 1985.
3. Stoughton, T. B., "Finite Element Modeling of 1008 AK Steel Stretched Over a Rectangular Punch with Bending Effects," *Computer Modeling of Sheet Metal Forming Process,* 12th Automotive Materials Symposium, The Metallurgical Society, April 29-30, 1985.
4. Budiansky, B., "Notes on Nonlinear Shell Theory," *ASME Journal of Applied Mechanics,* Vol. 35, p. 393, 1968.
5. Hutchinson, J. W., "Finite Strain Analysis of Elastic-Plastic Solids and Structures," *Numerical Solution of Nonlinear Structural Problems,* ed. R. F. Hartung, 17, ASME, New York, 1973.

Modeling plasticity in a two-phase ductile material

Elane C. Flower, Lawrence Livermore National Laboratory, Livermore, California
Dan Watt, University of Windsor, Ontario, Canada

The laboratory's CRAY X-MP/416 supercomputer makes full three-dimensional modeling not only possible but also economically feasible.

A number of commercially important alloys, including normalized steels (ferrite-pearlite), dual-phase steels (ferrite-martensite), and the titanium alloys, consist of large, strong particles thinly distributed in a softer ductile material. Although well-developed theories describe the deformation and strengthening behavior of alloys with precipitated particles that are small and close together, as in ordinary steels hardened by quenching and tempering, the theories do not apply to alloys with particles that are large and far apart. In other materials, such as glass fiber/epoxy composites, the softer matrix binds together long fibers of a stronger material. Well-developed equations based on the "law of mixtures" can be used to calculate the strength of such composites. These equations apply fairly well to the large-particle alloys, which contain no fibers.

Over the past 30 years, various investigators have studied the application of these equations to the large-particle alloys and have achieved many important insights into the behavior of the large-particle materials, especially behavior close to the particles. However, the complexity of the stresses and strains around and between the particles could be described adequately only with full three-dimensional modeling using, for example, the Lawrence Livermore Laboratory's NIKE3D finite-element Lagrangian computer program.[1] This approach, however, was very time-consuming and expensive on the computers available at the time. The computing power of the laboratory's CRAY X-MP/416 supercomputer at last makes full three-dimensional modeling not only possible but also economically feasible.

The model

Researchers at the Livermore Laboratory have been analyzing a model that consists of three spherical particles whose centers lie in a vertical plane. We chose periodic boundary conditions that enabled us to reflect this pattern of particles horizontally and vertically to infinity. The process can be repeated indefinitely on all sides.

To make this model behave as it would if it were one unit in a repeated pattern stretching out in all directions to represent an infinite medium, we applied a set of constraints on the motions of those nodes that are on the model's surfaces. Nodes on the top and bottom planes, for example, are free to move in these planes, that is, in the x and z directions, but those in the bottom plane are not allowed to move in the y direction (either up or down). Those in the top plane can move up only in unison as the model stretches.

Similarly, the front, back, and sides of the model can move inward enough to keep the volume constant as the model stretches. To ensure that they remain vertical, the calculation applies a set of horizontal, tractions to each node in these planes at the end of each deformation step. This exactly balances any forces that otherwise would move them out of line with the other nodes in the same vertical plane. These are the same tractions that would arise if a full set of mirror image units surrounded the one being analyzed, extending in all directions to infinity.

The radius R of the larger particles is two-thirds the depth of the model unit. Hence, the complete model consists of vertical sheets of particles in this pattern spaced at intervals of $3R$ between centers in a matrix of more ductile material.

Setting up the computation

The three-dimensional mesh representing this model, built with the lab's INGRID mesh generator,[2] consisted of 2484 six-sided elements with a total of 3161 nodal (corner) points. The unit model was 24 mm wide, 30 mm tall, and 9 mm deep. We performed the analysis by displacing the top surface upward 0.205 mm at a time in 50 equal time steps, for a total engineering strain of about 33 percent.

After building the mesh that defines the model's geometry we needed to enter the engineering properties of the model's two components in the bilinear, elastic-plastic material model used in the calculation. For the soft matrix, we chose a Young's modulus of 130 GigaPascals (GPa), a yield stress of 400 MegaPascals (MPa), and a hardening modulus of 400 MPa; for the hard particles, the corresponding numbers were 207, 700, and 200 MPa. Poisson's ratio (ratio of principal strain) for both materials was 0.30.

For large-scale, three-dimensional, finite-element calculations such as these, the cost and memory requirements increase as the fourth power of the resolution (number of elements); resolution is usually given in terms of the mean bandwidth. The mean bandwidth of the solution matrix more than doubled when we included the symmetry boundary conditions in the problem, which increased the cost more than 30-fold. Even on our CRAY X-MP/416 supercomputer, using 8 million 64-bit words of memory, we estimated on the basis of a two-time-step trial using conventional techniques that the full calculation would have taken 156 minutes, 136 of which would have been devoted to solving systems of linear equations directly by Gaussian elimination with the Fissile package in NIKE3D.

To shorten this calculation time, and reduce the computation expense, we turned to the Crout element-by-element, preconditioned conjugate-gradients method.[3] This recently implemented iterative procedure greatly accelerates the solution process. In our case, it reduced the time spent on equation solving to 16 minutes and the total computation time to 36 minutes, a four-fold improvement.

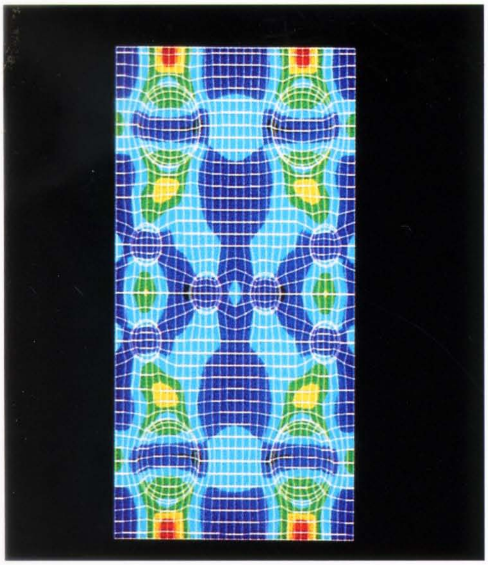

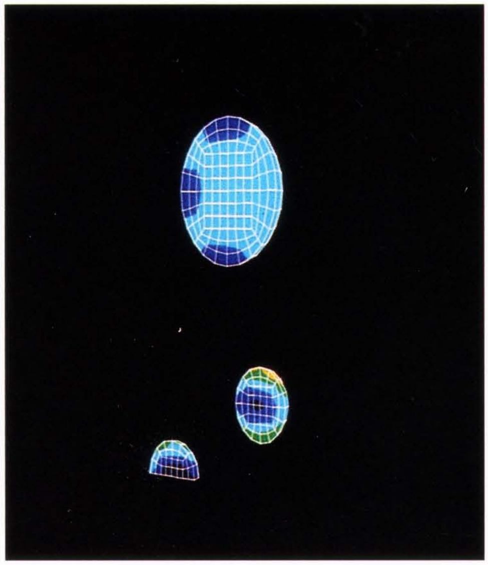

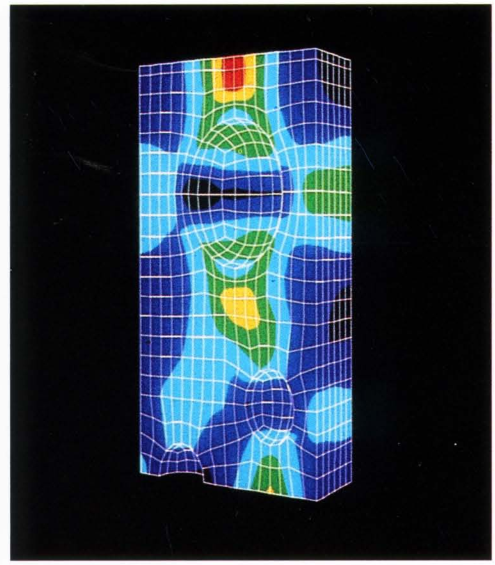

Figure 1 (left). Effective plastic strain in the soft matrix after being stretched vertically by about 33 percent. The plastic strain in any one element ranges from 16.35 to 40.6 percent with fringe levels set at 20.1, 24.3, 28.5, 32.7, and 36.9 percent.

Figure 2 (center). Effective Von Mises stress at maximum displacement in the spherical particles. Stresses in the particles range from 728.9 to 745.4 MPa with fringe levels set at 731, 734, 737, 740, and 743 MPa.

Figure 3 (right). Effective Von Mises stress at maximum displacement in the matrix. Stresses in the matrix range from 480.3 to 599.2 MPa with fringe levels set at 498, 519, 540, 560, and 581 MPa.

We used the lab's TAURUS code[4] to provide graphical displays of the computational results. TAURUS draws contours joining points that have the same value as any physical variable. Examples include stress in the x direction (or y or z), pressure, total effective stress, and effective plastic strain (Figure 1). The complexity of Figure 1 dramatically illustrates why researchers previously were unable to find a mathematical formulation for the strain distribution far from the particles.

We also examined the stress and strain states in individual elements around the particles and how they varied with time. We determined these variables averaging the values at each of the eight nodes at the corners of the element in question.

Results and interpretation

Our preliminary analysis of the results has produced several new and interesting observations on the stress and strain distributions close to and far from the particles. We found that the maximum amount of plastic strain occurs between the particles in the direction of the applied load. We also found regions of low strain between adjacent particles where one might expect the softer material to yield more.

One advantage of finite-element simulations is that they allow us to divide a problem into parts and review them separately. Figure 2 shows the Von Mises effective stress at maximum elongation in the formerly spherical particles; Figure 3 shows the corresponding effective stress in the matrix. The particles deform less and carry a higher effective stress because they have a higher yield strength than the matrix. Moreover, the variation in stress and strain from point to point within a given particle is small, as was predicted almost 30 years ago on theoretical grounds for elastic deformation.[5]

One might expect the vertical channel between the two large particles (in Figure 1) to undergo a large strain to compensate for the lack of plastic flow in the particles. Instead, we find that the large, strong particles must carry a higher-than-average stress to maintain equilibrium conditions. The vertical stress component, taken over the horizontal area of the planes passing through the particles, must balance the distant applied vertical loads that are stretching the unit cell; therefore, the matrix region between the particles must carry a lower-than-average stress to compensate for the higher-than-average stress in the particles. In effect, the particles shield these regions of the matrix from the full force of the applied loads.

Figure 1 also shows that the maximum strains occur above and below the particles, with the largest effective plastic strain midway between the largest particles and their reflected counterparts in the unit cells above and below (not included in Figure 1). In these top and bottom horizontal planes, one might expect to find strain higher than average because no particles exist to share the load. However, very little plastic strain exists at the center of these planes.

To understand the reason for the low plastic strain at the center of the top and bottom horizontal planes, we can go back to a 30-year-old stepwise computing method developed to find the strain field around an isolated single particle in an infinite matrix. The first step is to remove particles mathematically from the matrix, leaving a set of holes. Then we strain the matrix, in this case by 33 percent, pulling the holes out of shape. Because we know that the matrix and the particles must deform plastically, we must find a set of tractions that will bring holes and particles to the same final shape when applied to the matrix along the inner surfaces of the elongated holes and to the surfaces of the undeformed particles. Having restored both shape compatibility and stress equilibrium, we can set the particles back into the holes and add the stresses and strains produced by the tractions to the general strain field of the matrix.

Looking at this transformation process in detail, we see that at the top and bottom of each hole the matrix must be stretched back inward to conform to the stiffer particles, while around the sides of the particles it must be pushed outward. Clearly, this produces tensile stresses above and below each particle and compressive stresses around its sides. However, it is less obvious how these stresses vary with distance from the particle. In particular, we want to know why the maximum effective plastic strain is greatest at a distance from the largest particles (Figure 1), especially

Figure 4. Axial stress (the vertical component of the Von Mises stress shown in Figure 3) in the matrix at maximum displacement. These stresses range from 400.1 to 680.7 MPa with fringe levels set at 443, 492, 540, 589, and 638 MPa.

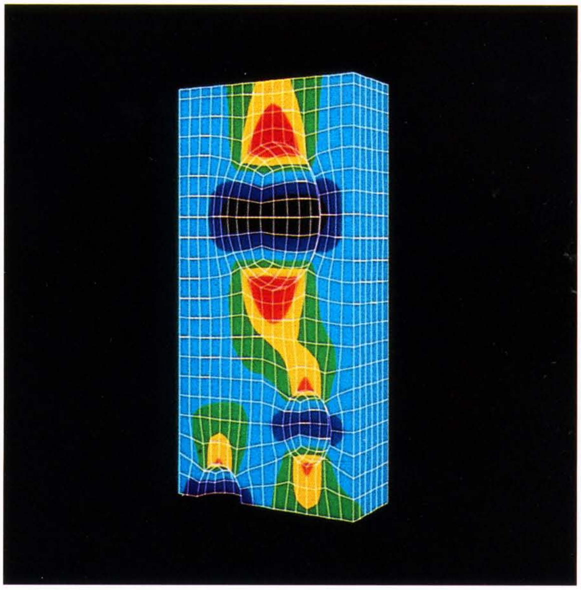

since the vertical tensile stress (Figure 4) is at a maximum directly above and below the particles.

In these regions of maximum vertical tensile stress, the tensile stress in the horizontal (x and z) directions is also very large. Thus, the matrix has a very high hydrostatic tensile stress, which plays an important role in void formation at the interface and correctly predicts where void nucleation will take place. But it produces relatively little strain. The deviatoric components, which depend on the difference between the vertical and horizontal components of tensile stress, and therefore contribute directly to plastic strain, are larger at a distance from the particle where the horizontal tensile stresses are less intense.

Although this analysis has helped to explain where the maximum plastic strain occurs, we still need to consider the cause of a region of low plastic strain at the middle of the upper face in Figure 1. The region lies in the channel between the two large particles and their reflected counterparts in the cells immediately above them (and similarly, in the bottom plane). Once again, this is a natural consequence of the requirement for stress equilibrium. Because extra vertical stress is needed at the maximum strain region to stretch the matrix back to fit the particle, the stress elsewhere in the same horizontal plane must be lower to maintain a balance with the applied vertical loads. Thus, although no particles are in either the upper or lower plane, the particles above and below the plane still can affect the strain distribution dramatically.

The most important general effect found by these finite-element calculations is that regions of high stress extend both up and down in the direction of the applied loads for about one particle diameter. When the particles occupy more than about 20 percent of the volume of the unit cell, their elevated stress fields effectively link up in the loading direction during tensile deformation.

Thus, the stress distribution after considerable plastic strain resembles that found in materials reinforced with long fibers, even though the particles are still nearly round. The elevated stresses above and below the particles, together with the requirement that the average vertical stress in any plane must balance the applied vertical load, creates low-stress channels between the particles also aligned in the loading direction. This explains, at least in principle, why equations based on the "law of mixtures" have been relatively useful in describing the properties of such alloys.

Although this somewhat cursory examination of the major effects has been quite revealing, much still needs to be learned by further analysis of the details of these finite-element calculations. We will continue to work on the details of the calculation, comparing the results with those predicted analytically. We also plan to incorporate the effects of prestress produced by thermal incompatibilities and transformation strains.

Summary

Using the Lawrence Livermore Laboratory's CRAY X-MP/416 computer system, we have performed a unique three-dimensional, finite-element simulation for large plastic strain in a material composed of two ductile solid phases of different yield strengths. The stronger component is present as relatively large and widely spaced grains occupying less than 10 percent by volume of the material. Our most significant result is the mapping of plastic strain variation throughout the matrix. These results, not intuitively obvious before the calculation, make sense in view of equilibrium and compatibility constraints. We have gained new insight into the reasons that calculations based on the "law of mixtures" work as well as they do in materials of this nature.

About the authors

Elane Flower works in the advanced engineering analysis group at the Lawrence Livermore National Laboratory and has applied finite element modeling to a variety of material deformation problems. She received B.S. and M.S. degrees in metallurgical engineering from the Colorado School of Mines.

Dan Watt is a professor of engineering in the Engineering Materials Group of the Mechanical Engineering Department at the University of Windsor. He received a B.A.Sc. degree in metallurgical engineering from the University of Alberta and a Ph.D. degree in metallurgy from McMaster University. He did postdoctoral work at the Cavendish Laboratory at Cambridge University.

References

1. Hallquist, J. O., "NIKE3D: An Implicit, Finite Element Code for Analyzing the Static and Dynamic Response of Three-Dimensional Solids," Lawrence Livermore National Laboratory report UCID-18822, Revision 1, 1984.
2. Stillman, D. W. and J. O. Hallquist, "INGRID: A Three-Dimensional Mesh Generator for Modeling Nonlinear Systems," Lawrence Livermore National Laboratory report UCID-20506, 1985.
3. Hughes, T. J. R., R. M. Ferencz, and J. O. Hallquist, *Computer Methods in Applied Mechanics and Engineering*, Vol. 61, pp. 215-248, 1987.
4. Brown, B. E. and J. O. Hallquist, "TAURUS, An Interactive Post-Processor for the Analysis Codes NIKE3D, DYNA3D, TACO3D, and GEMINI," Lawrence Livermore National Laboratory report UCID-19392, Revision 1, 1984.
5. Eshelby, J. D., *Proceedings of the Royal Society*, A 252, p. 571, 1959.

Large-scale solutions in structural analysis

Dawson Deuermeyer, Cray Research, Inc.

Although Cray Research computer systems are valuable tools for accelerating work that might be carried out on other systems, some computational problems are so large that only Cray systems provide a practical solution. In such cases, no other systems can provide a solution quickly enough to be useful in a typical engineering environment. Several such problems recently were solved by application specialists in Cray Research's Industry, Science & Technology Department using a CRAY Y-MP system and the MSC/NASTRAN structural analysis program from the MacNeal-Schwendler Corporation. These problems were not performance exercises, but involved typical structures of interest to scientific agencies and manufacturing corporations. The techniques used are applicable to all varieties of structural analysis problems.

Structural analysis problems can pose computational difficulties for several reasons:

- ☐ The physical phenomena involved can be complex and difficult to simulate. Such a situation, however, may tax the skills of the engineer more than it taxes the computer being used.
- ☐ The phenomena might be simple to simulate, but nonetheless be very demanding of computer resources. The use of a very fine computational mesh on a large structure is an example of such a situation.
- ☐ The type of analysis to be performed largely determines the complexity of the computational problem. A static analysis with 250,000 degrees of freedom (DOF) might be much less demanding than a 2500-DOF normal modes analysis or a 250-DOF nonlinear analysis. Each type, or subtype, of analysis has its own practical size limit. The ability to expand that limit can have a dramatic impact on engineering productivity.

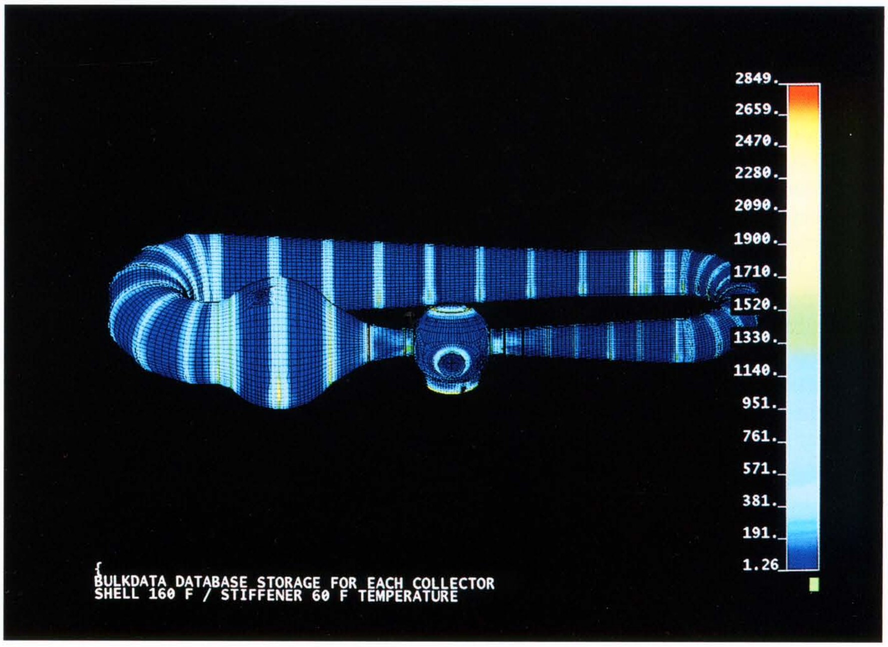

Figure 1. Wind tunnel model from NASA's Ames Research Center, showing temperature induced stresses.

The solution of the problems described here and similar problems demonstrates one of the main benefits of large-scale computing: it redefines what is possible and makes solution strategies routine that otherwise would be too difficult to use in production environments.

Wind tunnel

Application specialists at Cray Research recently ran an analysis on a pressurized wind tunnel model from NASA's Ames Research Center at Moffet Field, California (Figure 1). The model comprised

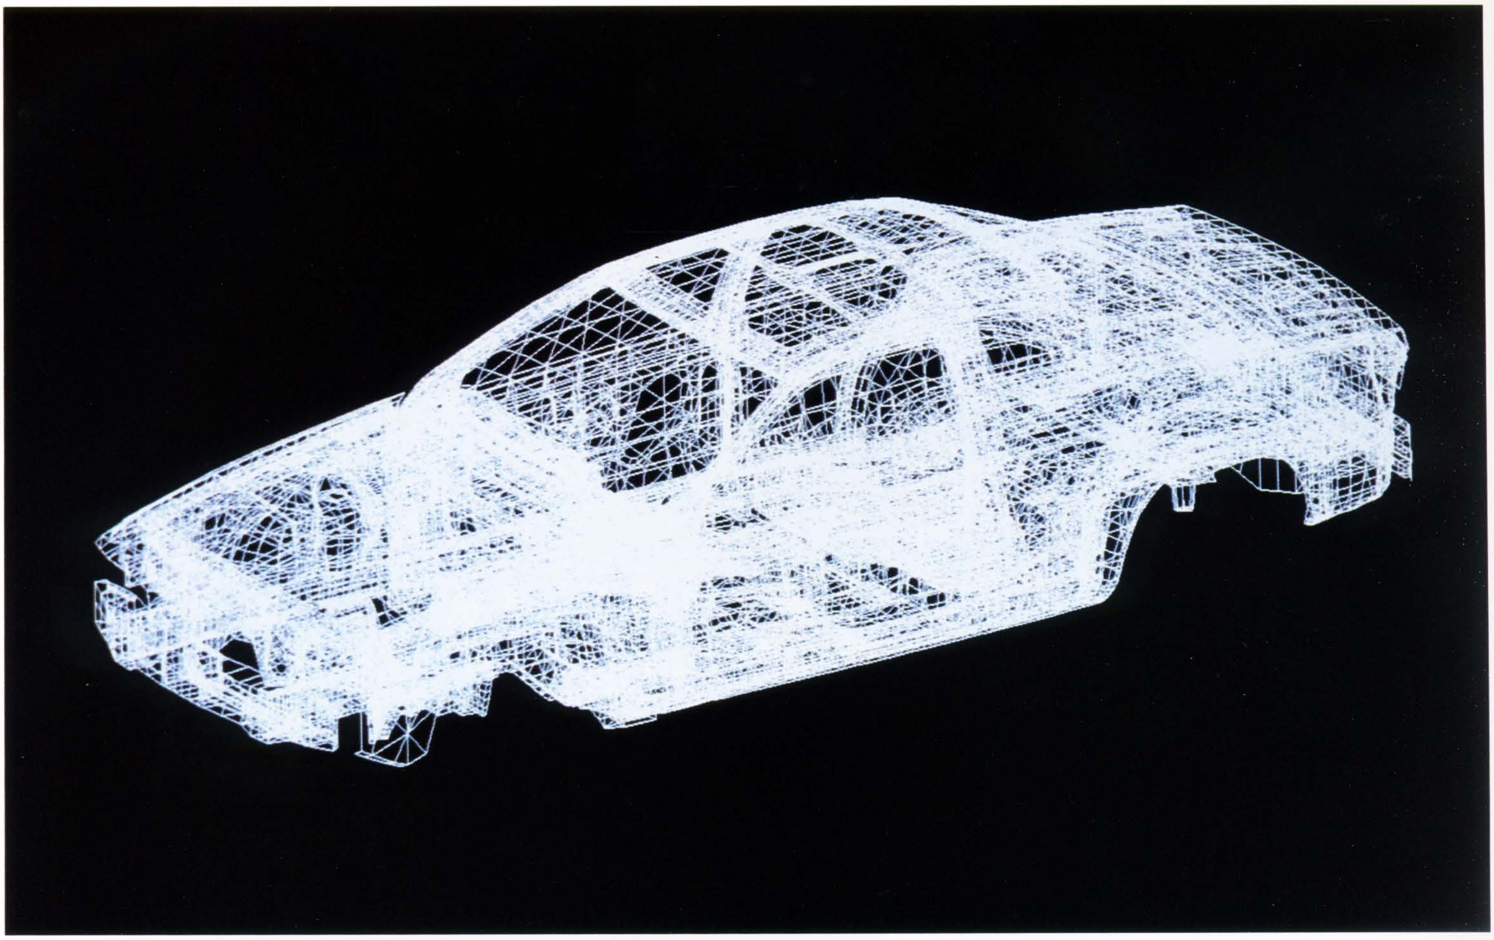

36,519 grid points and had 40,099 finite elements, most of which were two-dimensional, and 217,918 degrees of freedom. A static solution for 18 load cases conducted by Norman Engineering, a NASA contractor, required about two CPU days to complete on an IBM 4341 computer. The same analysis required about one-half hour of wall-clock time on a single processor of a CRAY Y-MP system at Cray Research. For static analyses, wall-clock times might run to 30 percent more than CPU time alone and up to 50-70 percent more for normal modes analyses. Cray Research systems are able to cut this overhead difference effectively due to two principal features: Cray Research's SSD storage device, which eliminates most of the I/O wait time, and the Cray operating systems UNICOS and COS, which deal very efficiently with heavy production workloads.

To run on the IBM system, the wind tunnel model had to be subdivided into 53 sections, called superelements. The use of superelements makes the computational task less demanding because each part of the structure can be analyzed separately. Appropriate measures then can be taken to ensure that computer resources are not exceeded. For example, data from different superelements can be off-loaded to tape and the individual jobs scheduled for execution. Although superelements can provide a useful option for the analysis of large models, they add complications to the overall process. For example, the engineer must decide how the model is to be subdivided to minimize "bookkeeping," and the subdivision process itself can introduce errors to the result. In contrast to the IBM system, the Cray system's data storage and I/O capabilities enabled application specialists at Cray Research to solve this model in its entirety, a capability that greatly simplifies the engineer's task.

The performance obtained on the wind tunnel model was made possible by the combined capabilities of the CRAY Y-MP system and enhancements included in the latest version (66A) of the MSC/NASTRAN program. These enhancements include

☐ Removal of the 65,535 degree-of-freedom limit for a single structure or a single superelement

Figure 2. Finite element mesh used in the model of the Ford Motor Company automobile body.

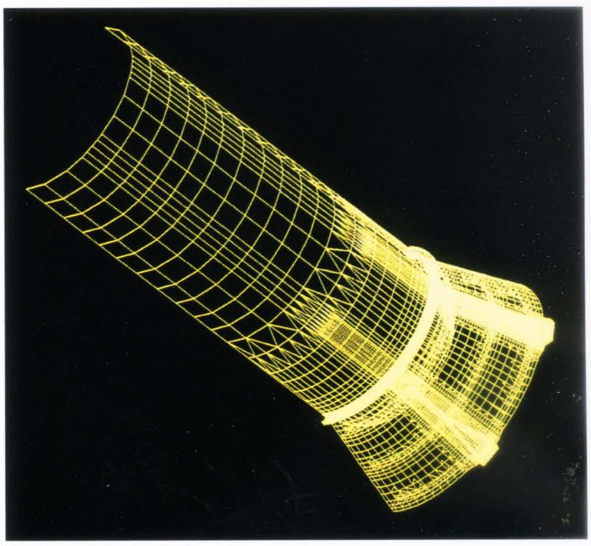

Figure 3. Finite element mesh used to model the space shuttle solid rocket motor aft skirt.

- ☐ Incorporation of design optimization technology
- ☐ Redesign of the executive system
- ☐ Introduction of new types of finite elements, in particular QUADR elements

Auto body

Another large structural analysis problem recently solved was a fully trimmed automobile body model from the Ford Motor Company (Figure 2). The model comprised 45,087 grid points and had 51,059 finite elements, most of which were two-dimensional, and 270,522 degrees of freedom. Because of its size and complexity, this model was divided into 42 superelements for easier computation. Although this method adds complications, as noted above, it was useful in this case because the overall model was heterogeneous in geometry and in material properties. Superelements enable engineers to isolate sections of a model, which then can be tested and redesigned independently of the model as a whole. Redesigned superelements then can be integrated into the rest of the model. Moreover, the superelement method enables engineers to use the parallel processing capabilites of Cray systems, because several superelements can be distributed among multiple processors and solved concurrently. A normal modes solution for this model, which seeks to find the vehicle's vibrational frequencies, required 3800 seconds of wall-clock time using up to seven CPUs on the CRAY Y-MP system. The total CPU time was 7100 seconds.

Aft skirt

Another problem from NASA involved work conducted in conjunction with the space shuttle redesign that followed the Challenger accident (Figure 3). The problem came from NASA's Marshall Space Flight Center in Huntsville, Alabama, and involved evaluation of the solid rocket motor aft skirt. The skirt is the flared base of the solid rocket motor, and the entire weight of the fully loaded space shuttle rests on the skirts of the two solid rocket motors. Four large bolts run through each skirt to help secure the shuttle on the launch pad. A model of the skirt was created to study stresses around the bolts because cracks were found to form around the bolt holes. The model comprised 12,598 grid points and had 10,101 two- and three-dimensional finite elements and 45,361 degrees of freedom. This problem was computationally intensive not only because of its size, but also because nearly half of the elements in the model were three-dimensional, and an exceptionally fine grid was required around the bolt holes. A static solution for one load case required only 863 seconds of CPU time on a single processor of a CRAY Y-MP system.

Flap actuator

The largest of the set of problems was a model of an airplane flap actuator for the Airbus A330/A340 airplane (Figure 4). The actuator is a product of the Zahnradfabrik Friedrichshafen Aerospace Division in West Germany. The model comprised 51,537 grid points and had 153,967 degrees of freedom. The problem was noteworthy for several reasons in addition to its size. Of its 33,619 finite elements, 32,883, or almost 98 percent, were three-dimensional elements, a situation that significantly increases the number of computations required to solve the problem. Of the remaining 736 elements, 701 were gap elements. The use of gap elements can increase the computational difficulty of a problem by an order of magnitude because they must be treated as nonlinear elements. The solution of this model required about 22.5 CPU hours to perform on an IBM 3090/180 computer without a vector facility and took about one wall-clock hour on the Cray system using up to five CPUs. Superelements were used to solve this model, enabling the analysts to take advantage of parallelism.

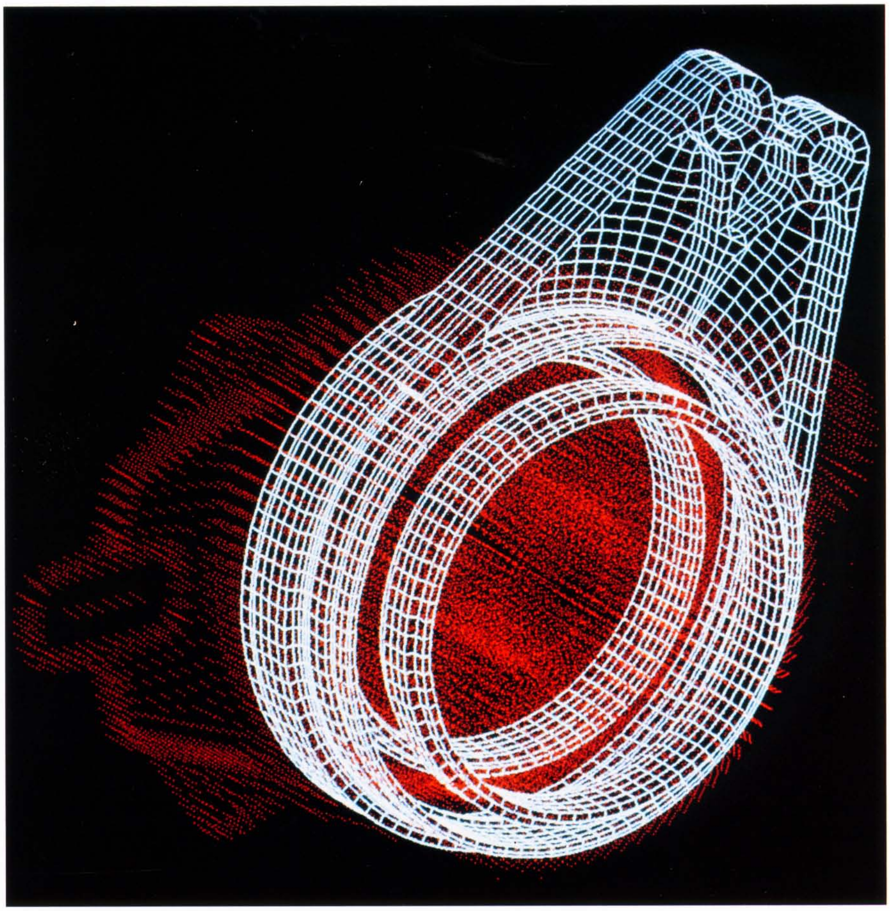

Figure 4. Grid points for the aircraft flap actuator model. The mesh for one superelement is shown.

These examples demonstrate the potential of Cray Research supercomputers to solve problems that would not be practical to solve on other systems. Although not all structural analysis projects will require the kinds of modeling capabilities described here, engineers using Cray Research's high-end systems can tackle large problems knowing that they are not likely to be constrained by computing capacity. Cray Research hardware and software systems are designed to help engineers maximize their productivity by ensuring that computing capacity is not a bottleneck in engineering environments.

About the author

Dawson Deuermeyer is a senior applications specialist in the Engineering Applications Group of Cray Research's Industry, Science & Technology Department. He received a B.S. degree in physics from St. Cloud State University in Minnesota in 1976.

A new measure of supercomputer performance

Results from the Perfect Benchmarks

Charles M. Grassl and James L. Schwarzmeier
Cray Research, Inc.

The development of the Perfect Benchmarks is an ongoing effort to evaluate supercomputer performance. This article describes the Perfect Benchmarks, characterizes the benchmarks in terms such as degree of vectorization and types of algorithms included, and presents performance results for the baseline and optimized versions of the codes for a variety of supercomputer systems.

Traditional means of measuring supercomputer performance do not reflect accurately how real application codes will perform. To answer the need for a more comprehensive and realistic way to evaluate supercomputer performance, a group of industry and university affiliates of the Center for Supercomputing Research and Development (CSRD) at the University of Illinois collected an industrywide set of 13 application programs for measuring supercomputer performance. The Perfect (Performance Evaluation for Cost-Effective Transformations) Benchmarks can complement standard benchmarks to create a more useful hierarchy of performance tests.

Members of the Perfect Committee, an academic and industrial group that is developing and evaluating this set of benchmarks, represent the California Institute of Technology, CSRD, Cray Research, Inc., HNSX Supercomputers, the Houston Area Research Center, IBM, the John von Neumann Center for Supercomputing, Princeton University, and the University of Houston. Members submitted and selected 13 programs to represent four application areas: fluid dynamics, physics and chemistry modeling, engineering design, and signal processing.

Background

The 13 programs are composed of 500 to 19,000 lines of Fortran statements and are written and maintained to run when compiled with any Fortran 77 compiler. Each program reports the CPU time, elapsed (wall-clock) time, and number of floating-point operations per second, and includes a verification output file. The operation count is determined by the number of floating-point adds, multiplies, and divides as reported by the hardware performance monitor on a CRAY X-MP system at the University of Illinois. The programs in the Perfect Benchmarks generally are run individually, and reported results consist of individual program times, MFLOPS, and various means or averages.

The Perfect Committee has reported performance results of the Perfect Benchmarks on 26 computer systems, including supercomputers from Cray Research, Fujitsu, Hitachi, IBM, and NEC. Baseline and optimized versions of the Perfect Benchmarks were run on one, two, four, and eight CPUs of a CRAY Y-MP8/832 computer system. The results exhibit many high-performance features of the CRAY Y-MP computer system, including the effectiveness of the compiling system for automatic vectorization and parallelization and the effectiveness of manual optimization, including parallelization, for application programs.

Benefits to the industry

The results of the Perfect Benchmarks can be used by compiler writers, computer architects, and application developers, all of whom can benefit from the documented experiences of the Perfect Committee. For example, compiler writers can strive to design compilers that might implement automatically some of the benchmarkers' optimizations. Computer architects can use the Perfect Benchmarks as a set of test programs to monitor the performance of specific architectural features, such as cache sizes or numbers of vector registers. Application developers can choose algorithms and data structures similar to those that perform well on certain classes of supercomputers, such as shared-memory vector processors.

As previously mentioned, the Perfect Benchmarks can complement standard benchmarks such as the LFK, the LINPACK tests, and the NAS Kernels. Together, these benchmarks form a useful hierarchy of tests. The most specific CPU-intensive test is the LFK test, which focuses on individual Fortran DO-loops that were chosen to be representative of DO-loops in programs used at the Lawrence Livermore National Laboratory. The LINPACK tests generally consist of a Fortran program for solving 100-by-100 and 1000-by-1000 systems of equations. The smaller-sized problem primarily tests the performance of SAXPYs, whereas the larger problem tests for parallel processing efficiency and peak deliverable MFLOPS. The NAS Kernels are seven mathematical subroutines that perform specific algorithms, including matrix multiplication, fast Fourier transforms (FFTs), Cholesky decomposition, and pentadiagonal matrix inversion. The NAS Kernels test the performance of entire subroutines and algorithms.

The Perfect Benchmarks are next in this hierarchy of benchmarks. All programs in the suite are complete applications; they have input and output and perform useful calculations. One of the programs in the suite has significant I/O, a few have modest I/O, and the rest have negligible I/O.

Because the Perfect Benchmarks are the first attempt to establish an industrywide set of representative application codes, a few areas can be strengthened. One major deficiency is problem size. By scaling down mesh sizes to allow for reasonable execution time on workstations, supercomputer performance is penalized with artificially small vector lengths, granularity, and I/O. Also, the suite lacks some important application codes such as particle-in-cell codes, which are important to the national laboratories, and structural codes, which are important in the automotive industry for crash analysis. Another major flaw with the suite is the lack of a throughput test. Most supercomputer users compute in a batch environment with many other users, and the ability of the supercomputer to process many jobs simultaneously is very important to its overall cost effectiveness.

Characteristics of the Perfect Benchmarks

Table 1 lists several characteristics of the 13 Perfect Benchmarks, including program size, memory size of the executables, and the number of I/O words transferred. Notice that MG3D is the only program with large I/O.

By comparing the internal characteristics of the baseline Perfect Benchmarks to the optimized versions, several trends are apparent. The amount of vectorization of the baseline programs ranges from 1 percent for program QCD, to nearly 100 percent for programs ARC3D, FLO52, BDNS, and TRFD. The optimized codes clearly show a higher percentage of vectorization and longer vector lengths than the baseline codes. Perhaps these results are not surprising for optimizations that are effective for a vector machine like a Cray Research system, but the results do indicate that many so-called "scalar" codes can be modified to be predominantly vector codes. In total, the optimized programs performed 9 percent fewer floating-point operations and 13 percent fewer memory accesses than the baseline programs. The decrease in memory traffic in program DYFESM is dramatic. This means that the optimized loops use data more efficiently by operating from vector registers rather than from memory.

A few general observations can be made about the codes on the application-type level. The fluid dynamics programs make extensive use of floating-point operations. These codes perform a small set of simple operations on the grid of data values and therefore have high numbers of floating-point operations per memory reference. The waiting times for vector registers and functional units for each of the Perfect Benchmark codes were monitored. For the fluid dynamics codes, the baseline and optimized waiting times are similar, suggesting that the baseline fluid dynamics codes already performed well. The biggest differences between baseline and optimized waiting times were found in ADM and OCEAN. Of the fluid dynamics codes tested, these two programs experienced the largest gains through optimization. Because the fluid dynamics programs generally were vectorized well, they would

Program	Original source system	Lines of Fortran	Memory size (Mwords)	I/O transfers (Mbytes)
Fluid dynamics				
ADM	IBM 3090	6105	0.320	0.230
ARC3D	CDC 7600	3607	1.250	1.250
FLO52	CRAY-1	1987	0.500	0.120
OCEAN	CRAY-1	4343	0.490	0.010
SPEC77	CYBER 205	3888	1.300	6.736
Chemistry and physics				
BDNA	IBM 3090	3980	0.540	3.184
MDG	IBM 3090	1238	1.350	0.510
QCD	MARK 1	2327	2.330	0.002
TRFD	IBM 3090	485	1.100	0.001
Engineering design				
DYFESM	CRAY X-MP	7608	0.021	0.044
SPICE	CDC 6600	18523	0.510	0.038
Signal processing				
MG3D	CRAY X-MP	2760	1.410	477.392
TRACK	MARK III	3790	0.200	0.184
Total		60641	—	489.701

Table 1. Origins and characteristics of the Perfect Benchmarks.

benefit most from additional functional units or additional CPUs. These programs benefitted greatly from additional CPUs (parallel processing).

The baseline chemistry and physics programs were not well-vectorized and spent little time waiting for vector functional units. On the other hand, the optimized versions of these programs wait a high percentage of time for vector registers and functional units. This is indicative of DO-loops that comprise many lines and encompass many variables in complicated arithmetic expressions.

The engineering design programs DYFESM and SPICE have complicated structures and use large amounts of indirect addressing. The baseline versions of these programs often require nearly two memory references to obtain a data value, especially in SPICE. These codes have relatively high values for the number of memory references per floating-point operation (1.46 and 1.74, respectively). This leads to a relatively poor MFLOPS rate, as shown in Table 3. The two codes differ significantly in one respect: the baseline version of DYFESM was relatively highly vectorized (95 percent), while the baseline version of SPICE was relatively scalar (7 percent vector instructions).

The two signal-processing programs MG3D and TRACK have a vector/scalar relationship similar to DYFESM and SPICE. The optimized version of program MG3D is dependent largely on functional units and vector registers. This dependency originates in the Cray Research SCILIB libraries, whose use constituted the dominant optimization for this code. TRACK is similar to SPICE in its use of sparse matrix techniques and many gather/scatter operations. Results for baseline tests and optimized versions of the Perfect Benchmarks were run on one, two, four and eight CPUs of a CRAY Y-MP8/832 computer system. CSRD has provided the partial results for supercomputers from Fujitsu, Hitachi, IBM, and NEC.

The median run time for the baseline programs is 22 seconds on one CPU of a CRAY Y-MP computer system, and the median optimized time is 11 seconds. The median times for baseline and optimized programs on the IBM 3090-600S/VF system are 389 seconds and 61 seconds respectively. On a VAX-11/785 computer, the median run time for the baseline programs is one hour. Thus, it is difficult to scale the problem size for each program so that it runs in a reasonable amount of time for even a small number of systems. The MG3D program, which ran for the longest time on most systems, generally was not run on the smaller systems due to CPU time and disk space limitations.

Table 2 lists the results of the baseline Perfect Benchmarks for one, two, four, and eight CPUs of a CRAY Y-MP8/832 system. Results also are shown for a Fujitsu VP-100 system, a Hitachi S-820/80 system, an IBM 3090-600S/VF system, and an NEC SX-2 system. These tests were run without manual optimization. The Cray CFT77 compiling system was used for both vectorization and parallelization.

The IBM system achieves its best relative performance on scalar code, such as SPICE and TRACK, whereas it is not a strong performer on vector code. The Hitachi S-820/80 and NEC SX-2 computers perform relatively well compared to a single CPU of a CRAY Y-MP computer system for a stride-one vectorized program like FLO52. However, for nonunit-stride vector operations, such as those in ARC3D and OCEAN, and for gather/scatter operations, such as in BDNA, the Hitachi and NEC computers fall short of even a single CPU of the Cray Research system. This is somewhat of a surprise, considering that the Hitachi and NEC systems have four vector pipes compared to the CRAY Y-MP system's one vector pipe. Evidently, good performance on the Hitachi and NEC computers is tied closely to stride-one accesses, although neither computer uses a vector cache. The Fujitsu system has the poorest scalar performance and only moderate vector processing power. The Fujitsu VP-100 system is a one-pipe system and so is not directly comparable to the four-pipe Hitachi S-820/80 and NEC SX-2 systems.

	Computer system							
Program	CRAY Y-MP/1	CRAY Y-MP/2	CRAY Y-MP/4	CRAY Y-MP/8	Fujitsu VP-100	Hitachi S-820	IBM 3090	NEC SX-2
Fluid dynamics								
ADM	19	19	19	19	7	13	8	15
ARC3D	140	193	252	291	51	—	9	68
FLO52	109	182	285	329	57	226	17	129
OCEAN	32	34	35	36	7	—	5	17
SPEC77	36	36	36	36	15	21	10	26
Chemistry and physics								
BDNA	84	96	113	121	17	39	14	57
MDG	17	17	17	17	7	15	8	14
QCD	13	13	13	13	5	9	6	9
TRFD	56	56	56	56	30	—	11	55
Engineering design								
DYFESM	47	52	59	59	28	69	18	58
SPICE	6	6	6	6	3	5	4	4
Signal processing								
MG3D	23	23	23	23	11	—	11	33
TRACK	8	8	8	8	8	7	5	5

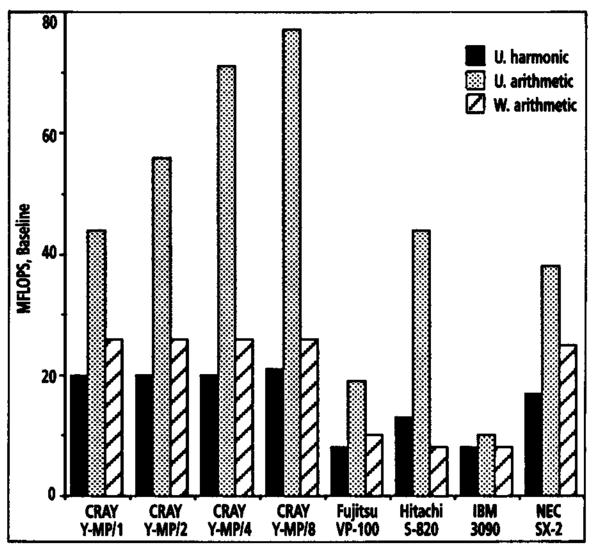

Table 2 and graph. Baseline performance in MFLOPS for various large computer systems.

On the accompanying bar graphs, the unweighted harmonic mean (U. harmonic) represents the processing rate for all programs combined if each program had the same number of operations. The unweighted arithmetic mean (U. arithmetic) represents the processing rate for all programs combined if each program ran for the same amount of time. The weighted arithmetic mean (W. arithmetic) represents the processing rate for all programs combined with the current run times and operations and run as one program.

Results for optimized tests

Table 3 lists results of the optimized Perfect Benchmarks for one, two, four, and eight CPUs of a CRAY Y-MP8/832 system. Results also are shown for an

IBM 3090-600S/VF system and an NEC SX-2 system. These programs were optimized by Cray Research, IBM, and NEC representatives respectively. Hitachi and Fujitsu did not provide data for this comparison.

Program	CRAY Y-MP/1	CRAY Y-MP/2	CRAY Y-MP/4	CRAY Y-MP/8	IBM 3090	NEC SX-2
			Computer system			
Fluid dynamics						
ADM	62	75	86	91	8	15
ARC3D	148	289	548	989	59	121
FLO52	111	187	300	347	32	129
OCEAN	124	194	249	275	5	17
SPEC77	101	191	327	543	10	31
Chemistry and physics						
BDNA	142	184	242	288	62	57
MDG	80	161	310	595	73	136
QCD	39	74	144	250	6	9
TRFD	76	148	272	440	11	55
Engineering design						
DYFESM	136	209	275	295	32	64
SPICE	20	20	20	20	8	5
Signal processing						
MG3D	191	369	676	1094	15	33
TRACK	18	26	33	39	5	5

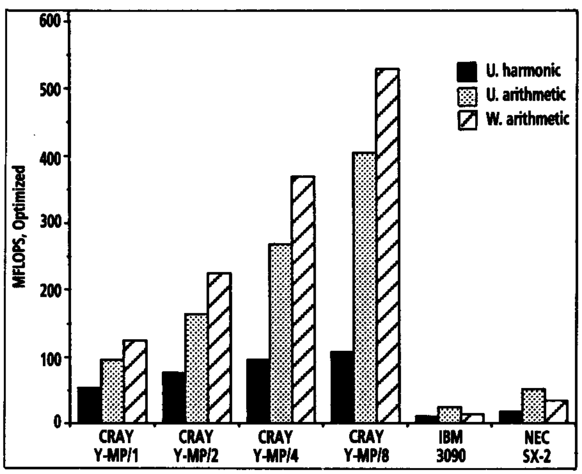

Table 3 and graph. Optimized performance in MFLOPS for various large computer systems.

For an explanation of the accompanying bar graph, see Table 2 caption.

Not all programs were well-vectorized even after significant optimization effort. The baseline versions of programs QCD, SPICE, and TRACK were almost entirely scalar before optimization. SPICE originally was written in the 1960s for small memory machines, and much of its scalar nature originates from its memory management scheme. QCD is a Fortran translation of a C program written for hypercubes, and thus the baseline version did not map well onto a shared-memory vector architecture. After vectorizing parts of each of these programs, much of the computation still remained in integrating the data structures. That is, much of the computational time is spent searching algorithms, which is not reflected in MFLOPS ratings, and in gather/scatter operations. With parallelization, most of the programs showed significant speedup for multiple Cray CPUs. The average program speedup with eight CPUs was a factor of four. In spite of not being well vectorized, TRFD and TRACK exhibited significant parallelization. ARC3D, MG3D, SPEC77, QCD, DYFESM, and TRACK were parallelized at a relatively high level. For instance, in the case of TRACK, the parallelization was implemented four subroutine levels up from the actual computations. This level of parallelism would be very difficult to detect with automatic tools or compilers. However, the expert programmer or knowledgeable scientist can exploit this parallelism manually.

The NEC system also shows speedups due to vectorization, but they are not as dramatic as those achieved on the Cray Research system. The NEC SX-2 system derives much of its speed from its multiple functional units, or pipes. Without very long stride-one vectors and high vectorization levels, it is difficult to utilize a significant fraction of the peak performance of the NEC SX-2.

The optimizations did not show as many performance benefits on the IBM system as on the other vector processor systems. Not one of the optimized programs ran at a speed faster than 75 MFLOPS. The fastest running optimized program on the 3090-600S system, program MDG, ran at 73 MFLOPS. The optimization for this program includes vectorization and parallelization. Very few programs were able to utilize parallel processing on the IBM system.

Summary of performance results

Three out of 13 baseline application programs demonstrated significant automatic parallelization with the Cray Research Autotasking compiler system, CFT77. With further optimization, all programs but one were parallelized significantly. The average eight-CPU speedup due to parallelization was 4.1. The optimized version of MDG had a speedup of 7.4, which indicates the Cray Research system's efficiency in parallel processing.

Both baseline and optimized versions of the Perfect Benchmarks were run on the CRAY Y-MP series of computer systems and compared with results provided by other vendors. The CRAY Y-MP system clearly delivers the best performance of all the computer systems over the cross section of codes, ranging from scalar, to stride-one vector, to nonunit-stride vector programs. Optimization is most effective on the CRAY Y-MP system, resulting in greater performance gains than on the other computer systems tested. This demonstrates the CRAY Y-MP system's very attractive price-performance ratio for real application codes.

About the authors

Charles Grassl is manager of Cray Research's Benchmarking Department. Before joining Cray Research in 1984, he was employed by Sperry Semiconductor, where he was a process simulation engineer. Grassl received a Ph.D. degree from the University of Wisconsin at Madison.

James Schwarzmeier is a member of Cray Research's computational chemistry group. He formerly was a member of the system and performance analysis group. He received a Ph.D. degree in plasma physics from the University of Wisconsin at Madison in 1977. After serving as a postdoctoral fellow at the Courant Institute of Mathematical Sciences, New York University, he worked at the Los Alamos National Laboratory for six years.

CORPORATE REGISTER

Environmental organizations worldwide select Cray systems

The National Oceanic and Atmospheric Administration (NOAA) has ordered two eight-processor CRAY Y-MP supercomputer systems, which will provide the United States with numerical weather prediction and climate research capabilities that are among the best in the world. The purchased systems will be installed in the first half of 1990 at the National Meteorological Center (NMC), a division of the National Weather Service, in Suitland, Maryland; and the Geophysical Fluid Dynamics Laboratory (GFDL) in Princeton, New Jersey. Both systems will run Cray Research's UNICOS operating system. The new supercomputer at NMC is part of an overall modernization program for the National Weather Service. Five weather forecasting models are run regularly on their computer systems. These include 2- and 10-day weather forecasts and a 72-hour aviation forecast. The new Cray Research supercomputer will improve both the accuracy and speed of the forecasts. The GFDL applies supercomputer modeling to atmospheric and oceanic research. The laboratory is involved in long lead-time investigations in climate dynamics and weather forecasting.

The European Centre for Medium-range Weather Forecasts (ECMWF) has ordered an eight-processor CRAY Y-MP system, which will be installed in the third quarter of 1990 at Reading, Berkshire, England. This is the fourth time ECMWF has selected a supercomputer from Cray Research. With the installation of a CRAY-1A system in 1978, ECMWF became Cray Research's first European customer. A CRAY X-MP/22 system was installed in 1983 and was followed by a CRAY X-MP/48 system in 1986. "We are confident that the installation of the CRAY Y-MP system will enable the center to maintain its leading role in numerical weather prediction," said Lennart Bengtsson, director of ECMWF. "We plan to use the power of the CRAY Y-MP supercomputer to increase substantially the resolution of our operational forecast model and to improve the handling of the physical processes. The more accurate weather forecasts that will be produced will be of immediate benefit to the meteorological services of the member states." ECMWF is supported by 18 European nations.

The National Center for Atmospheric Research (NCAR) in Boulder, Colorado, has ordered an eight-processor CRAY Y-MP system, which will be installed in the second quarter of 1990. The system will more than double NCAR's computing capacity. "The CRAY Y-MP system will enable NCAR researchers to expand their models for greater accuracy as they take advantage of the additional processors and parallel architecture," said John Rollwagen, Cray Research chairman and chief executive officer. Computational researchers at NCAR use supercomputers in research areas including climate simulation, ocean simulation, atmospheric chemistry, cloud and storm behavior, and solar effects.

Deutscher Wetterdienst (DWD), the West German weather service, has installed a four-processor CRAY Y-MP supercomputer at its office in Offenbach, Federal Republic of Germany. "We are familiar with Cray Research systems from our affiliation with the European Centre for Medium-range Weather Forecasts (ECMWF)," said Reiner Lamp, director of the computer center at DWD. "This new system will provide the power we need to improve still further our numerical weather prediction and climate-modeling capabilities." The DWD produces short-range forecasts covering one to three days and coordinates efforts with ECMWF in medium-range weather forecasting, which covers three to ten days. Both organizations use supercomputers to maintain and improve the accuracy of their forecasts.

Direction de la Météorologie Nationale (DMN), a new customer for Cray Research, has installed a CRAY-2/4-256 supercomputer at Ecole Polytechnique in France. DMN plans to use the system for day-to-day weather forecasts as well as for climate research. The Cray Research system will enable DMN to implement higher resolution models and thereby obtain more accurate forecasts.

Taiwan National University has installed a CRAY X-MP/14se supercomputer, the first Cray Research system to be installed in the Republic of China. The purchased system was installed at the university's computer center in Taipei, during the first quarter of 1990. The university will use the system for computational support of scientific research projects there and at other universities and research institutes in the Republic of China. The university also will use the system to educate and train the country's academic community in the use of supercomputer technology. "This order is especially exciting for us because it is our first from the Republic of China, and it signals continued success for our marketing efforts in the industrialized countries of Asia," said Marcelo A. Gumucio, president and chief operating officer of Cray Research.

The Saudi Arabian Oil Company (Saudi ARAMCO) installed a CRAY-2 computer system during the fourth quarter of 1989 at its computer facility in Dhahran, Saudi Arabia. Saudi ARAMCO, a Cray Research customer since 1984, is using the CRAY-2 supercomputer primarily for reservoir simulation. The company is using Cray Research's UNICOS operating system.

The ARCO Oil and Gas Company has installed a four-processor CRAY Y-MP computer system with SSD solid-state storage device at its computer center in Plano, Texas. ARCO, a division of the Atlantic Richfield Company, was among the first petroleum companies to acquire a Cray Research supercomputer when they installed a CRAY-1S system in 1981. ARCO will use the system to support petroleum exploration and production.

Florida State University (FSU) installed a four-processor CRAY Y-MP8 system in the first quarter of 1990. The supercomputer will be used to help the Department of Energy develop efficient, innovative solutions to complex energy-related problems and to develop new algorithms. "Developments in computational science are applications driven," said Robert M. Johnson, vice president of Research and Graudate Studies at FSU. "We improve computational methods as a natural sequence of doing computationally intensive scientific problems on state-of-the art supercomputers. We are looking forward to the added power provided by the CRAY Y-MP system."

The **National Aeronautics and Space Administration** (NASA) Goddard Space Flight Center has ordered a four-processor CRAY Y-MP8 system for Goddard's Space and Earth Sciences Computing Center in Greenbelt, Maryland, and a four-processor CRAY Y-MP8 system for NASA's Lewis Research Center in Cleveland, Ohio. Both systems will be installed in the second quarter of 1990. "This is the first time that NASA has placed an order with us for two systems," said John Rollwagen, chairman and chief executive officer of Cray Research, "Once the CRAY Y-MP system is installed at Goddard, all major NASA research and operational centers will have Cray Research supercomputers. We are pleased to be the supercomputer supplier for all of their centers." The CRAY Y-MP system at Lewis will supplement a CRAY X-MP system that was installed in January 1986. NASA Lewis focuses on research, development, and project activities related to aeropropulsion, space propulsion, and power systems. Researchers at the Goddard Center will use the new supercomputer for research in areas such as global warming, earth observing systems, space physics, astronomy and solar physics, and terrestrial and ocean modeling.

The **Atomic Weapons Establishment** (AWE) in Aldermaston, United Kingdom, has ordered an eight-processor CRAY Y-MP supercomputer, which will be installed at the AWE facility during the first half of 1990 and will use the UNICOS operating system. The new system will replace a CRAY-1 system that has been in continuous use at the establishment for more than 10 years. "The Cray Research systems at Aldermaston have become a cornerstone for our work here, and we are all very excited by the leading edge R&D environment the new CRAY Y-MP system will offer our scientists," a Ministry of Defense spokesperson said.

The **University of Trieste,** a new customer for Cray Research, has installed a CRAY X-MP/14 system at the Universita' Degli Studi Di Trieste Computer Center. The Italian university will use the system for research in medicine, engineering, physics, computer science, and astronomy.

The **Catalonian Research Foundation** has ordered a CRAY X-MP/14se computer system to serve a network of universities in Spain. The supercomputer will be installed at the Polytechnic University in Barcelona, Spain in the third quarter of 1990, pending export license approval. Through a network, the supercomputer will serve three academic institutions: the Polytechnic University, the Autonomous University, and the Central University. This will be the third Cray Research system installed in Spain.

UNICOS 5.1 includes new features

Release 5.1 of the Cray Research operating system UNICOS, which runs on any model Cray Research supercomputer, adds new features and contains many enhancements to UNICOS 5.0. The UNICOS 5.1 release reflects the continuing commitment of Cray Research to provide its users with a stable, portable, and powerful software environment for both batch and interactive processing.

UNICOS 5.1 supports all Cray Research computer systems (with the exception of either CRAY-1 systems or CRAY X-MP systems without an I/O Subsystem and with less than two million words of memory). All peripheral storage units and communications hardware products supported by UNICOS 5.0 are supported by UNICOS 5.1, as is the Cray Operator Workstation.

UNICOS 5.1 supports the following languages on all Cray computer systems: CFT77 3.1 or later, Cray C 4.1 and 5.0, Cray Standard C 1.0 or later, Pascal 4.0 or later, Cray Ada 1.1, Allegro Common Lisp 1.0, and CAL version 2 release 3.2. The UNICOS 5.1 version of CFT2 is supported on CRAY-2 computer systems. CFT 1.16 is supported on CRAY Y-MP, CRAY X-MP EA, CRAY X-MP, and CRAY-1 computer systems.

The following new features of UNICOS are included with this release:

- ☐ A disk quotas feature allows system administrators to control the amount of disk space and the number of files used by various clients. Limits can be set by account, by user, or by group.
- ☐ Access to Storage Technology Corporation's STK4400 Autoloader is provided through Cray Research VM station software by way of a UNICOS Station Call Processor (USCP) interface. (An autoloader is a robotic device that mounts and dismounts tape cartridges without operator intervention.) The interface is being tested for Cray Research MVS station software, and, with a small additional modset, support will be available for a TCP/IP path to Sun Workstations running the UNIX operating system.
- ☐ The Sun Microsystems yellow pages distributed data lookup service allows multiple systems in a networked environment to maintain current information for their password files. This feature is especially useful when combined with the Network File System (NFS) feature of UNICOS.
- ☐ The disk file accounting feature is an enhancement to a basic capability introduced in UNICOS 5.0 that allowed an account ID to be associated with a file. UNICOS 5.1 provides more convenient user access to disk-accounting information.

In addition to adding new features, release 5.1 adds many enhancements to existing features of UNICOS. Enhancements have been made to the following operating system components: tape subsystem, security, data migration, and Network Queuing System. Enhancements also have been made to the following software products and communications protocols: the CDBX debugger, SEGLDR, TCP/IP (including support for the Simple Network Management Protocol), NFS, USCP, and on-line diagnostics.

For more information about UNICOS 5.1 contact the nearest Cray Research sales office.

APPLICATIONS UPDATE

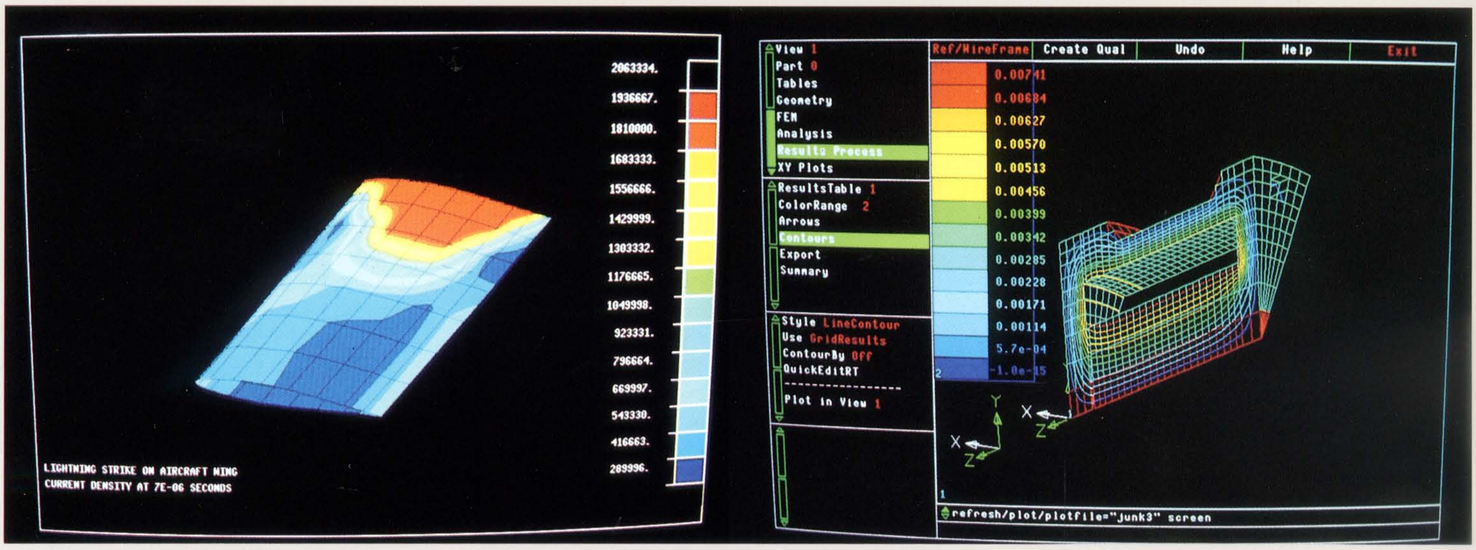

Lightning strike on an aircraft wing (left) and a satellite mirror positioning system (right) simulated with MSC/EMAS.

MSC/EMAS provides electromagnetic analysis

MSC/EMAS is a new, general-purpose, two- and three-dimensional finite element program for solving electric and magnetic field problems for linear, nonlinear, and anisotropic materials. The program, which is based on MSC/NASTRAN, runs on all CRAY X-MP and CRAY Y-MP computer systems under the UNICOS operating system. The program was developed by the MacNeal-Schwendler Corporation's Engineering/Electromagnetic Applications Department.

The electromagnetic analysis system can be used by electrical analysts, designers, engineers, researchers, and physicists who need to determine electromagnetic fields and field effects quickly and accurately. MSC/EMAS analyzes a full range of electromagnetic behavior, from electrostatics and nonlinear magnetostatics to eddy currents and wave propagation. MSC/EMAS applications include

- ☐ electromechanics
- ☐ power conversion
- ☐ sensors
- ☐ solenoids
- ☐ electrostatic discharge
- ☐ magnetic bearings
- ☐ MRI
- ☐ microwave cavities
- ☐ wave guides
- ☐ lasers

The program includes DC, AC, and transient solution capabilities and encompasses eddy currents and resonant mode analysis. MSC/EMAS includes an interactive graphics interface based on MSC/XL, a program that provides an efficient interface to finite element analysis. This closely integrated pre- and postprocessor combines a modern menu structure, powerful command language, advanced graphics, and integration with MSC's family of finite-element tools.

MSC/EMAS is based on a new unified formulation of Maxwell's equations. The program enables users to mix one-, two-, and three-dimensional elements with scalar elements. This capability enables users to address analysis requirements using simplified models while retaining high solution accuracy.

In addition, circuit analysis can be included in the model by using scalar elements such as resistors, capacitors, and inductors. Users then can include driving point impedance with a voltage source, enabling circuit models to be solved simultaneously within the electromagnetic solution.

For more information about using MSC/EMAS on Cray computer systems, contact Edwin J. Fabiszak, Jr., The MacNeal-Schwendler Corporation, Engineering/Electromagnetics Applications Department, 9076 North Deerbrook Trail, Milwaukee, WI, 53223; telephone: (414) 357-0323; or Greg Clifford, Cray Research, Inc., 1333 Northland Drive, Mendota Heights, MN, 55120; telephone: (612) 681-3658.

LS-DYNA3D optimized for Cray Research supercomputers

LS-DYNA3D is an explicit three-dimensional finite element program designed to help engineers and scientists analyze the large deformation response of solid, shell, and beam structures. LS-DYNA3D has been optimized for Cray computer systems and offers users many new capabilities not available in the public domain version of DYNA3D. The program is available for use on CRAY-1, CRAY X-MP, and CRAY Y-MP computer systems running under the Cray operating system UNICOS.

LS-DYNA3D is a product of Livermore Software Technology Corporation, founded by John O. Hallquist, the author of DYNA3D. The company was formed to enhance, support, and maintain the program.

In the past DYNA3D has been used extensively in the automotive, metalworking, aerospace, nuclear, and defense fields. Where comparisons with experimental results have been possible, a remarkably high degree of correlation often has been achieved.

The package includes new capabilities, including

- ☐ dynamic relaxation data base
- ☐ arbitrary node and element numbering
- ☐ extensive I/O options
- ☐ rivet and spotweld capability
- ☐ eroding contact
- ☐ torsional springs and dampers
- ☐ computation of energy dissipation due to hourglass control of zero energy modes
- ☐ fabric model for air bags and seat belts
- ☐ spatial variation of viscous and Coulomb friction
- ☐ nodal rigid bodies
- ☐ advanced data base for composite analysis
- ☐ fully supported thick-shell elements
- ☐ energy dissipative moving and fixed rigid walls
- ☐ force variation over rigid walls
- ☐ 35 consecutive models

The following enhancements provide up to a 50 percent increase in speed over public domain DYNA3D: right-hand side vectorization, single point constraint and nodal boundary condition vectorization, and vectorization of the contact algorithms.

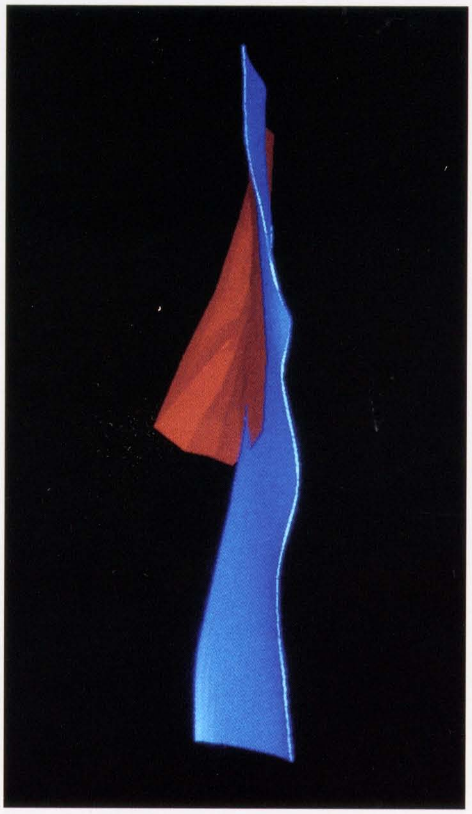

Simulation of a bird striking a turbine blade. Analysis was performed with LS-DYNA3D running on a CRAY X-MP system.

For more information about using LS-DYNA3D with Cray computer systems, contact John O. Hallquist, Livermore Software Technology Corporation, 2876 Waverly Way, Livermore, CA, 94550; telephone: (415) 449-2116; or Mike Long, Cray Research, Inc., 1333 Northland Drive, Mendota Heights, MN, 55120; telephone: (612) 681-3656.

GPSS/C available on Cray Research systems

GPSS/C (General Purpose Simulation System) from Simulation Software Ltd. is a language designed for discrete event modeling. The language collects statistics, produces tabulated results, and performs other tasks, enabling users to concentrate on the important issues of model development. The language is able to interface with user-written subroutines, a capability that may be used to access graphics packages, data bases, or routines that are performing scientific or engineering calculations. GPSS runs on Cray Research computer systems under the UNICOS operating system.

GPSS features an interactive source-level debugging capability that allows the setting of breakpoints in the model, single-stepping, and examination of attributes of model entities. Internally compiled code and optimizations shorten the execution times for production model runs. Direct access to online help is available. GPSS is used to teach simulation and modeling at universities and in industrial applications such as materials handling, logistics, transportation, industrial engineering, and communications.

GPSS provides a set of abstract components of various types and a set of operators called *blocks* that perform certain actions on the individual components. The *transaction* is the component that moves through a sequence of blocks that has been designed to model the system being studied. The state of the components of the model determines the details of the way in which a block of a given type will operate. For example, a block that permits a transaction to take control of a piece of equipment will not allow the transaction to proceed if the equipment is already at maximum capacity.

Several types of equipment components are available in GPSS. A *facility* is an entity that is either available for use or is in use by no more than one transaction at a time. A *storage* is similar but has a capacity that may be specified to suit the needs of the model builder. Finally, a logic switch is a simple ON/OFF element that may be set and tested to modify the path of a transaction through the blocks of the model.

These components are abstractions. In a model of a factory operation, for example, transactions could represent units being assembled; facilities might represent robot welders; and a logic switch could be used to simulate a machine breakdown. In a model of a high-speed communications network, transactions might correspond to messages, facilities to transmission lines, and storages to message buffer space. Whatever the application, GPSS entities provide a natural parallel to the parts of the system being modeled.

As transactions move through blocks that operate on equipment entities, GPSS collects equipment usage and transaction behavior statistics such as average contents, average occupancy time, and maximum contents, for inclusion in a report produced automatically when the model run completes. GPSS includes many additional measurement, reporting, and control capabilities. For more information about using the GPSS package on Cray Research computer systems, contact David Martin, president, Simulation Software Ltd., 760 Headley Drive, London, Ontario, Canada, N6H 3V8; telephone: (519) 657-8229; or Doug Petesch, Cray Research, Inc., 1333 Northland Drive, Mendota Heights, MN, 55120; telephone: (612) 681-3654.

Call for entries: Cray Research 1990 Gigaflop Performance Awards

In 1989, Cray Research initiated the Gigaflop Performance Award Program to highlight achievements in computational science. In its first year, the program recognized 20 individuals and teams who had solved problems using Cray Research computer systems running in excess of one gigaflop — that is, over one billion (Giga-) FLoating Point Operations Per Second.

This year the program award target has been raised to 1.5 GFLOPS. If you are conducting important work and achieving performance at this level, Cray Research encourages you to submit an entry for evaluation. Entries for this year's program are being accepted through September 1, 1990. A panel of Cray Research scientists and engineers will review each entry to confirm that it meets the award criteria. Awards will be presented in New York City at a reception and dinner the week of November 12, 1990, during the IEEE/ACM Supercomputing '90 conference.

The Cray Research Gigaflop Performance Award Program is conducted to recognize scientists and engineers who are working at the leading edge of their respective disciplines while furthering the science of supercomputing. For an entry form or more information, contact Vicky Frank, Gigaflop Performance Award Program administrator, Cray Research, Inc., 1333 Northland Drive, Mendota Heights, MN, 55120.

Computer cooling problem solved

Application specialists in Cray Research's Industry, Science & Technology Department recently solved a problem for the Hewlett-Packard Company that should help the computer maker improve the efficiency of its design-cycle. The problem involved managing the temperature in Hewlett-Packard's 850 series of minicomputers. The 850 series systems are air cooled, and Hewlett-Packard engineers traditionally have used a combination of experimental and personal-computer-based computational methods to determine the systems' cooling requirements.

But as the component density inside the systems has grown, determining the needed cooling power has come to require greater precision. Specifically, a precise determination of the air-pressure drop across the system is needed to determine the systems' air flow characteristics. The pressure drop is a measure of the air resistance against which the systems' cooling fans must work. This information then enables the PC-based program to compute the temperatures of the chips and other components.

"Traditionally, the pressure drop and flow characteristics of computers have been measured experimentally on prototypes," explains Vivek Mansingh, a research and development engineer in Hewlett-Packard's Systems Technology Division. "But an accu-

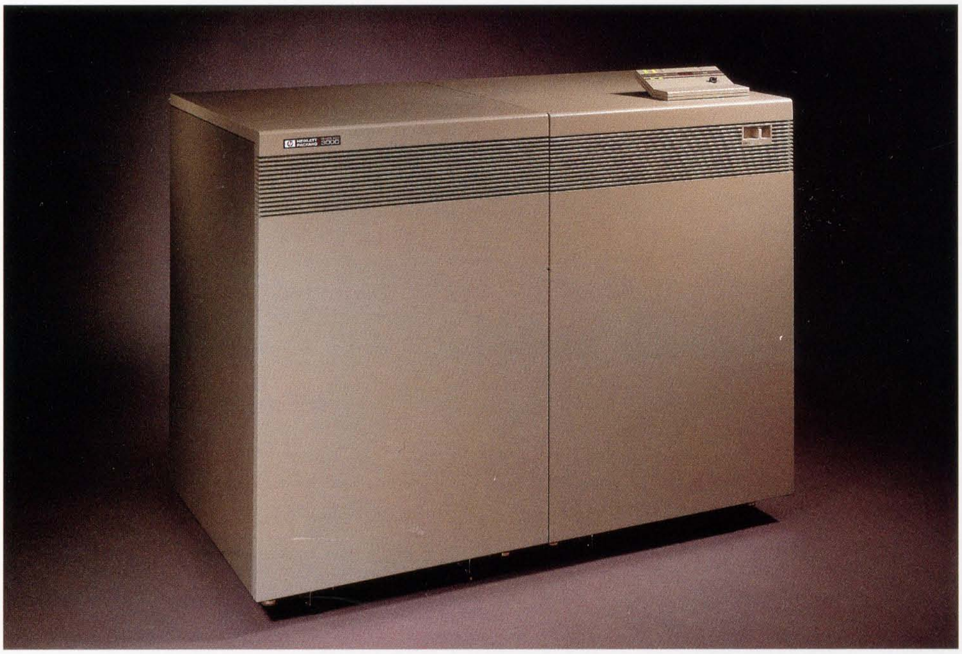

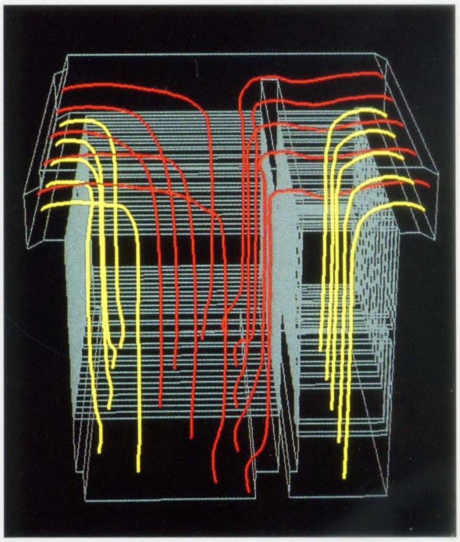

A minicomputer from Hewlett-Packard's 850 series (left). Air flow through the computer was modeled on a CRAY Y-MP system to help determine the best way to keep the computer cool (above). The colored lines are particle traces that follow the downward flow of air.

rate prototype is available only after all of a system's components have been designed. Therefore, if the pressure drop in a real system turns out to be excessive, or the air flow characteristics are found to be different than in the prototype, some major changes might be required at the end of the design cycle. And these design changes can significantly delay the development of a product. Numerical modeling is useful not only to help to avoid these changes, but it also enables us to simulate effects that are difficult or expensive to create experimentally, such as high altitude or microgravity."

Kent Misegades, manager of Cray Research's computational fluid dynamics (CFD) applications group explained that the project was, among other things, a test to see how well finite-element CFD programs calculate air flow and pressure drop. The program used in this research was the FIDAP program from Fluid Dynamics International. Misegades and Mansingh recognized from the outset that detailed three-dimensional modeling from the component level to the system level was impossible due to the complexity of the problem. Therefore, because the main focus of the modeling effort was on characteristics at the system level, some component-level simplifications were made in the geometry. Specifically, printed circuit board components were modeled as volumetric flow blockages.

The model may have been the largest FIDAP model ever run; it comprised more than 60,000 elements and had more than 240,000 degrees of freedom. The model ran initially at about 10 hours on one processor of a CRAY Y-MP system, but with programming suggestions from the software vendor, the run time was reduced to about 8 hours.

The model verified experimental measurements made by Hewlett-Packard engineers, and successfully predicted a significant cross-flow of air from the back to the front compartments of the system. "The particle traces show some extremely interesting flow characteristics that could not easily have been determined otherwise," notes Mansingh. "And despite the simplifications used in the model, it provided reasonably accurate predictions of the pressure drop and of the flow velocity through the computer's boards." The model predicted the air velocities through the CPU board slots to be about 1.5 to 2.5 meters per second, whereas experimental measurements found velocities of about 1.5 to 2 meters per second.

This research not only solved a specific problem, but also demonstrated the ability of finite-element CFD programs to predict the air flow characteristics inside computer systems that have complex arrangements of internal components. The ability to determine computationally the pressure drop and air flow characteristics for a given computer will enable designers at Hewlett-Packard to devise the best strategies for keeping their systems cool. The information gained from large-scale modeling efforts can be used to determine early in the design cycle the optimal fan size, the number of fans, and the board and component layout that will provide the best thermal management for the systems. And this capability will become increasingly important as the company makes use of new chips that have advantages over older ones but that also produce more heat.

Supercomputing hypersonics

Aircraft that travel faster than the speed of sound — supersonic aircraft — are commonly used around the world and have taken their place as established products of the aerospace industry. But few of these craft attain hypersonic speeds — speeds faster than five times the speed of sound — and no hypersonic commercial aircraft is yet in operation.

The development of hypersonic commercial aircraft promises to be of great benefit to commercial aviation, and research efforts in various countries are aimed at the design and engineering of such aircraft. Commercial hypersonic flight could cut intercontinental flight times significantly. Los Angeles-to-Tokyo trips, for example, could be completed in less than three hours. Hypersonic flight also holds the promise of single-stage-to-orbit flight vehicles. The evaluation of hypersonic designs requires testing capabilities beyond those provided by wind tunnels. Such design work is an exacting technical challenge being met largely through the use of computational models and supercomputers.

Researchers Taras M. Atamanchuk and Jean P. Sislian at the University of Toronto's Institute for Aerospace Studies are using the university's CRAY X-MP/24 computer system to investigate shock-induced combustion ramjet engines as potential propulsion systems for hypersonic flight vehicles. Shock-induced combustion is not as well understood as diffusive combustion, which is an alternative and more easily implemented type of combustion, but it offers some compensating advantages. During shock-induced combustion, the temperature and pressure of a combustible mixture are raised sufficiently, by means of a shock, to initiate combustion. This type of combustion results in a very short combustion zone and reduced thermal loads within the combustor.

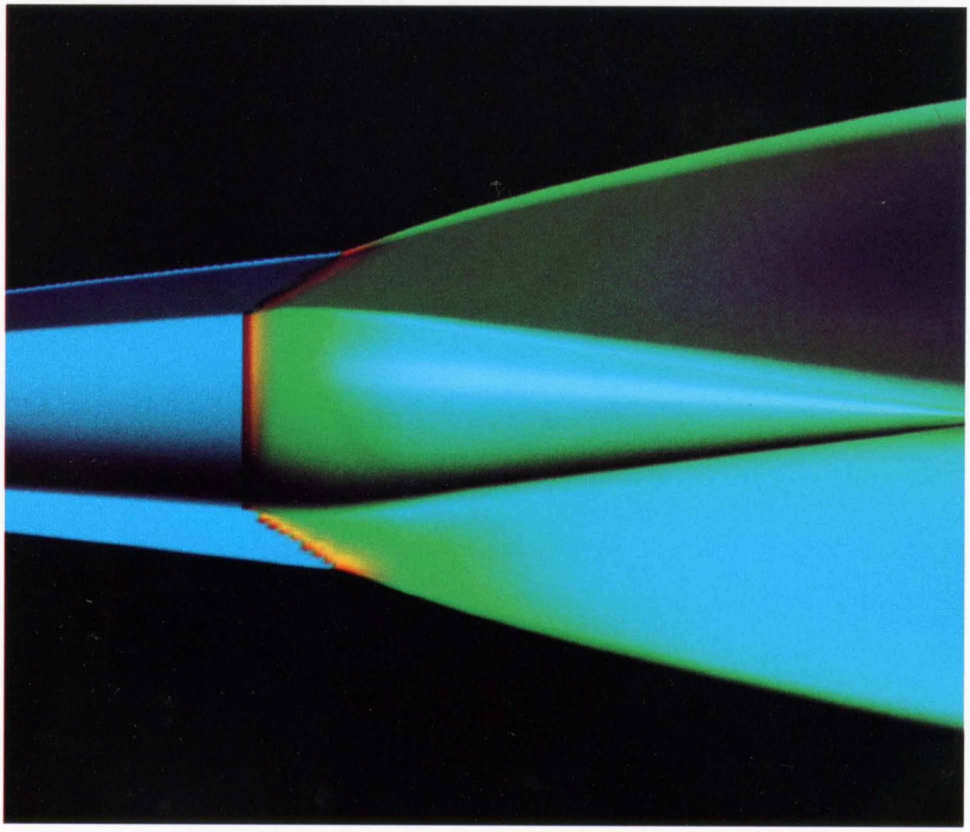

Model showing the underside of a hypersonic flight vehicle. Colors correspond to air pressures. The vehicle is shown operating at Mach 10.

As with all aerospace design efforts, delivering a favorable thrust-to-drag ratio is a central goal in the design of hypersonic craft. The thrust-to-drag ratio of a vehicle operating at a given Mach number is dependent on two factors: the amount of heat added to the flow around the body and the body geometry. For vehicles operating at design conditions, the body geometry is not known beforehand and actually is generated as part of the solution. As a result, a maximum cycle temperature restriction is required to solve such problems. The researchers developed a simplified detonation wave combustion model that revealed that this maximum temperature is dependent primarily on the amount of heat added to the flow. The determination of the value of heat addition leaves the vehicle geometry as the only independent variable from which to obtain a corresponding thrust-to-drag ratio. However, this ratio is not known a priori and it is necessary to iterate on the vehicle geometry until the desired thrust-to-drag ratio is obtained. This procedure must be repeated for a range of flight Mach numbers as well as for several different body configurations, such as with and without a cowl, which can be planar or axisymmetric. The calculations are two-dimensional, and each run typically does not require a large amount of CPU time, the average time per run being 240 CPU seconds on the Cray system. However, the number of runs is on the order of several thousand. Consequently, a supercomputer is needed to complete all of the runs in a reasonable time.

The researchers also are investigating integrated three-dimensional lifting-propulsive bodies, known also as waveriders. These bodies are created by replacing flow streamlines with solid surfaces, a procedure that introduces another degree of freedom to the problem, because the thrust-to-drag ratio obtained from the two-dimensional calculations is unlikely to correspond to that obtained in three dimensions. Designers of lifting-propulsive bodies are interested primarily in performance at cruise conditions, where the thrust-to-drag ratio equals 1.0. Cruise conditions, however, depend on not only the two-dimensional flow field but also the three-dimensional body itself. In other words, for every planar or axisymmetric body and flow field and its corresponding thrust-to-drag ratio, there exists a range of possible three-dimensional bodies and their associated thrust-to-drag ratios, one of which may be equal to 1.0.

Images for this study were generated by transferring raw pressure and geometry data from the Cray system to an Apollo network for further processing and graphical analysis. The visualization tools used were developed at the University of Toronto's Institute for Aerospace Studies, in part by researcher Vince Pugliese, and are built on Apollo's Graphics Metafile Resource 3-D subroutine package. Final rendering was carried out on an Apollo DN590 graphics workstation with a 24-bit frame buffer.

The image on page 27 shows the computed pressure distribution acting on the undersurface of a hypersonic lifting propulsive body at Mach 10, based on an axisymmetric flow field. Pressures are color-coded with blue representing the lowest pressure, through yellow, to red, which represents the highest pressure. The flight vehicle is operating at cruise conditions. The inlet section is predominantly blue followed by a very short combustion zone, in red, generated by a bump on the body. The high pressure flow then expands through a half-open nozzle section producing thrust. This blended wing-body design represents the complete integrated flight vehicle and produces both lift and thrust. The results obtained so far indicate that this mode of combustion is a promising means of hypersonic propulsion.

Scientist wins award for supercomputer research

Gregory J. McRae, associate professor of chemical engineering and engineering and public policy at Carnegie Mellon University, has been chosen as the first recipient of the Forefronts of Large-Scale Computing Problems Award. The SIAM Institute of the Advancement of Scientific Computing created the award to provide a distinguished prize in computational science equivalent to the Turing award in computer science.

An eight-member panel selected McRae for his innovative computer modeling of large atmospheric systems. He has used Cray computer systems to develop a detailed model of pollution in the Los Angeles area. Legislators there used the model to help formulate a new pollution control law, which was passed last year. McRae received the award at the IEEE Supercomputing '89 Conference in Reno, Nevada. The conference and the new award promote interdisciplinary excellence and innovation in computational science.

McRae has been using the Cray supercomputers at the Pittsburgh Supercomputing Center since 1986. He cites their capacity as the key technological factor that enabled him to do his research. "To simulate processes of this scale would have taken several years on conventional computers," says McRae. "At the Pittsburgh Supercomputing Center, it took 200 hours computing time, or a few weeks work." In addition to atmospheric research, McRae uses the supercomputer to collaborate with computer scientists and engineers at Carnegie Mellon to study parallel and distributed processing.

McRae joined the Carnegie Mellon faculty in 1983, after receiving a Ph.D. degree in environmental science from the California Institute of Technology. McRae has received a number of other honors, including the National Science Foundation's Presidential Young Investigator Award, the George Tallman Ladd Research Prize, and an Environmental Science Fellowship from the American Association for the Advancement of Science. McRae is on the editorial board of the International Journal of Supercomputer Applications, and he is member of several government and White House task forces on environmental issues.

Cray Research France contributes to high-energy physics award

The 1989 High Energy and Particle Physics Prize of the European Physical Society has been awarded to Georges Charpak, an internationally known researcher who has contributed to advancements in the field of particle physics. The award was presented at the International Europhysics Conference held in Madrid, Spain, in September 1989. The prize was established with contributions from the European Physical Society and seven European industries, including Cray Research France. Charpak uses Cray Research systems for his research at the European Center for Nuclear Research (CERN), where he has been a staff member since 1963.

Charpak was the first recipient of the award, which was established to honor researchers who have made important contributions to theoretical or experimental particle physics. The selection committee chose Charpak after consulting 130 high-energy physicists from around the world. Charpak is recognized for his outstanding contributions to the development of detectors in particle physics. He has made significant contributions to the development of the multiwire proportional chamber, the drift chamber, and other gaseous detectors, which are used in high-energy physics, medicine, and biology.

Charpak currently is a senior physicist in the experimental physics division of CERN. He completed graduate studies at the Nuclear Chemistry Laboratory at the College of France in Paris. He also was the director of research at the National Center for Scientific Research at the laboratories of the College of France.

GALLERY

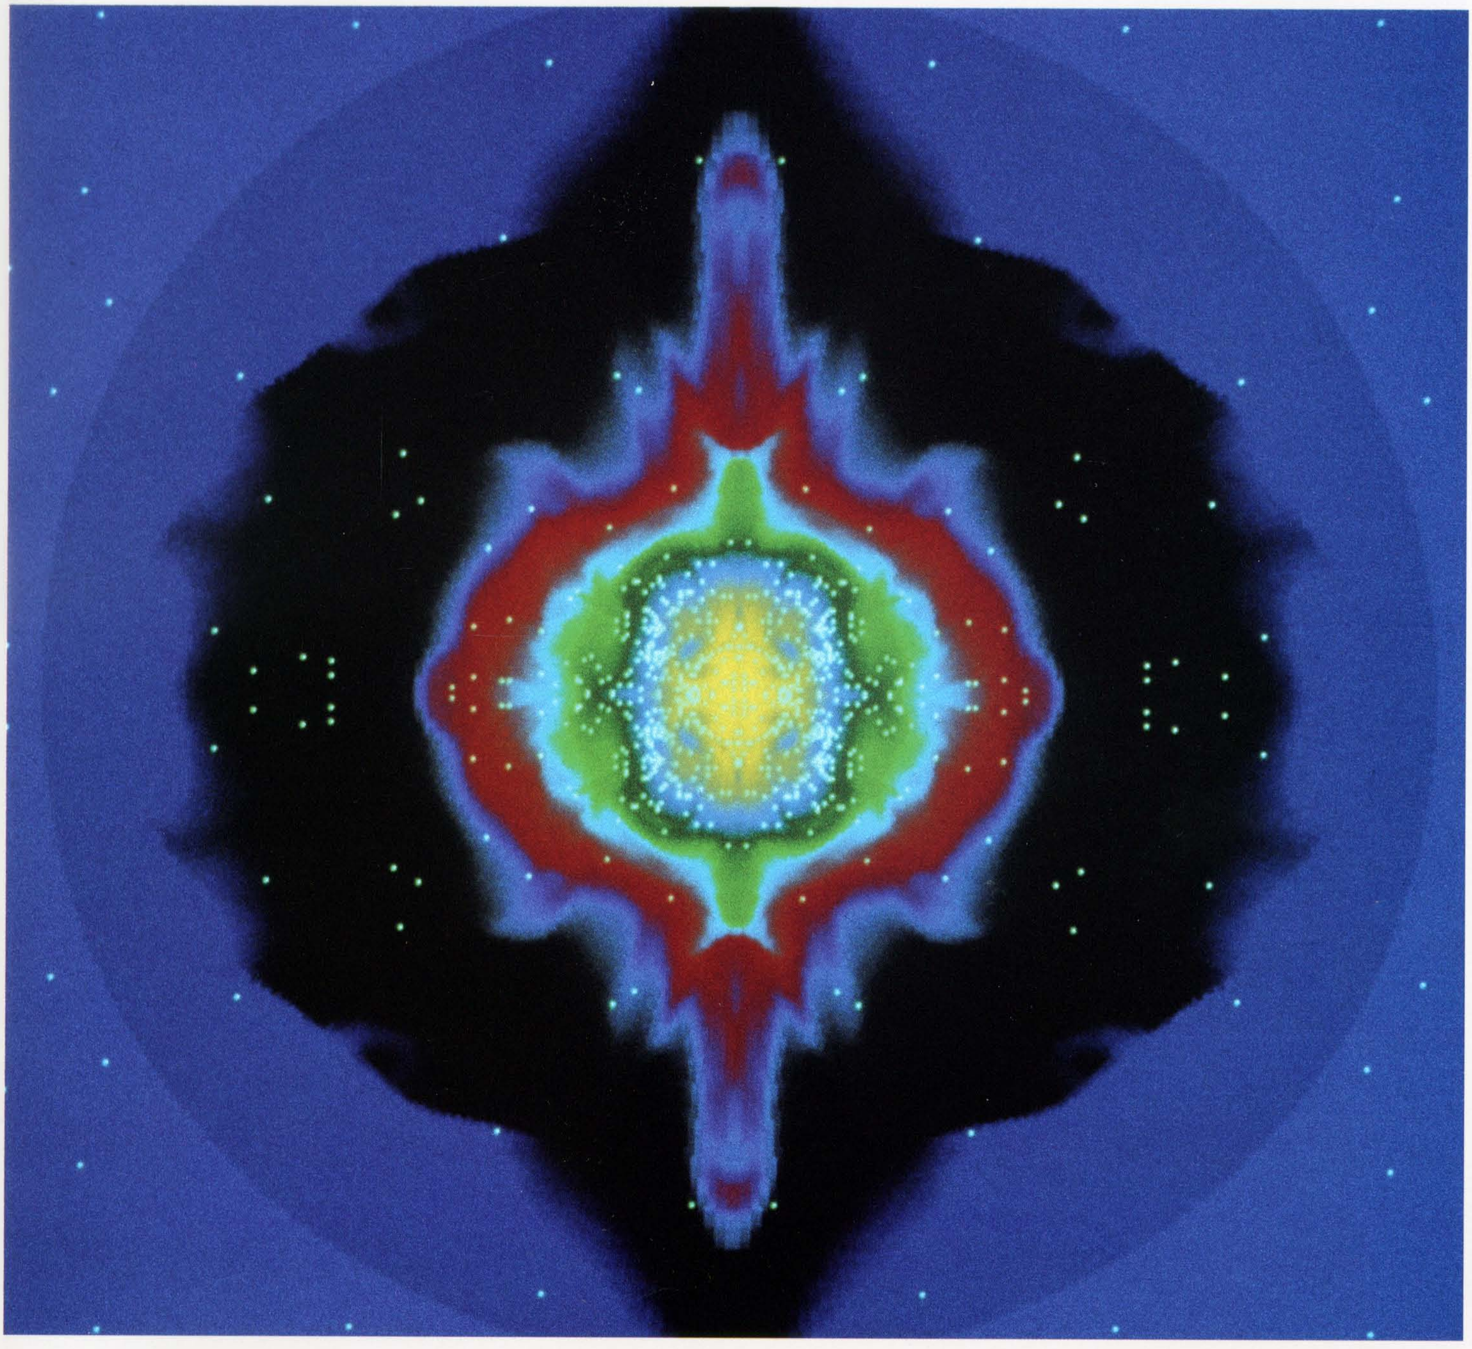

Two neutron stars collide in this image rendered on a CRAY X-MP system by Donna Cox and Charles Evans at the National Center for Supercomputing Applications, Urbana-Champaign, Illinois.

CRAY CHANNELS welcomes Gallery submissions. Please send submissions to the address inside the front cover.

CRAY RESEARCH, INC.
Marketing Communications
1333 Northland Drive
Mendota Heights, MN 55120
612/681-3438