CRAY CHANNELS

VOLUME 16, NUMBER 2 • A CRAY RESEARCH, INC. PUBLICATION

powering the automotive industry

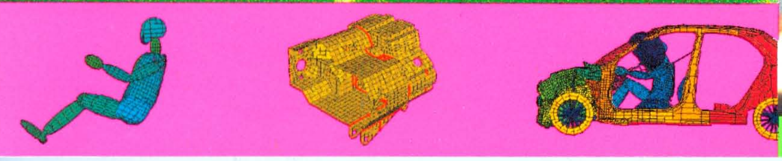

Announcing the CRAY J916 computer system and HEXAR software auto

C O N T E N T S

I N T H I S I S S U E

This issue of CRAY CHANNELS presents recent examples of supercomputer applications in the automotive industry—applications that evolved from leading-edge research and development tools to today's full-fledged production design tools. Companies in the automotive industry are rapidly increasing their use of simulation to design new, high-quality products and manufacturing processes—faster and at lower cost, providing a much-sought competitive edge.

Today's automotive manufacturers also increasingly require suppliers to design the systems and components they furnish. As a result, suppliers now undertake many high-performance simulations previously carried out by manufacturers. To obtain the necessary computing capability, they often turn to Cray Research for supercomputer systems and the expertise developed from years of working with the automotive industry. In addition, most automotive manufacturers have Cray Research systems—and some require their suppliers to maintain compatibility.

This issue contains two articles from automotive suppliers. An article from Altair Computing discusses design optimization, and an article from Alcan International describes their sheet metal forming simulations. Several additional articles illustrate the use of simulation by manufacturers, including articles from Chrysler Corporation (defroster duct airflow), Ford Motor Company (crashworthiness simulation), the BMW motorcycle division (noise and sound optimization), and Hyundai Motor Company (crashworthiness simulation).

A final note: Cray Research is now on the World-Wide Web. Check out our http://www.cray.com/ location with your favorite web browser.

F E A T U R E S

D E P A R T M E N T S

Cover Photograph: Liaison International

C R A Y C H A N N E L S

CRAY CHANNELS is a publication of the Cray Research, Inc. Marketing Communications Department. Published three times per year, it is intended for users of Cray Research computer systems and others interested in the company and its products. Subscription inquiries and address changes should be sent to Department D, CRAY CHANNELS, Cray Research, Inc., 655D Lone Oak Drive, Eagan, Minnesota 55121. Other comments and queries may be emailed to channels@cray.com. For more information about Cray Research and its products, visit http://www.cray.com/ on the World-Wide Web.

Volume 16, Number 2

Editorial staff
Chris Kruell, editor
Molly Sachse, associate editor
Rich Brueckner, Cindy Bielke, Heidi Hagen, and Ken Jopp, contributing editors

Design and production
Barbara Cahlander, electronic production artist
Eric Hanson, art direction and graphic design

New CRAY J916 system aligns automotive suppliers with manufacturers

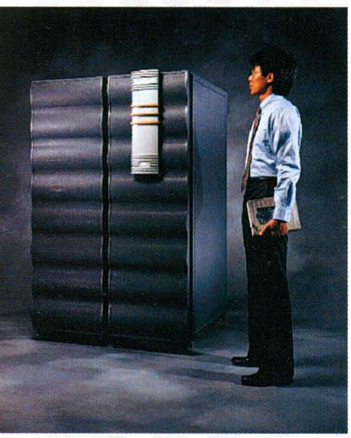

CRAY J916 supercomputer system

Supercomputer technology is a competitive must-have for the world's leading carbuilders. Now more than ever, the companies that make and supply automotive manufacturers with subsystems from engines to air conditioners are addressing their own needs for high-performance simulation. Compatibility of hardware and software is a major issue. In this environment, Cray Research's cost-effective, easy-to-use, air-cooled supercomputers are proving to be an ideal way for automotive suppliers to align themselves technically with their customers.

The CRAY J916 system is the latest air-cooled supercomputer from Cray Research. With prices beginning at under $300,000, CRAY J916 systems provide the highest throughput in their price range with up to 16 CPUs and large central memories. Binary compatible with more powerful Cray Research supercomputers, the CRAY J916 system runs the most widely used scientific and engineering applications. And it runs them fast—without modification. With up to 12 times the performance at the same price points as previous systems, the CRAY J916 system is ideal for budget-conscious departments and organizations. It does not require a computer room environment or special power or cooling arrangements. And with its open software environment, it integrates easily into existing distributed computing networks.

The CRAY J916 system has the memory and I/O bandwidth needed to sustain parallel workload throughput. In addition, as problem size and complexity increase, optimized applications allow users to take full advantage of the CRAY J916 system's parallel-vector processing (PVP) capabilities. Some large manufacturers, seeking full-service vendors that can design and engineer parts with the same technology and speed that they have, see the CRAY J916 system as an opportunity for their suppliers to share the high-performance road.

> *"We've got all these hot rod workstations, but the Cray is more like a bulldozer. We just keep stacking the work up in front of it, and the Cray plows through it when the workstations get bogged down."*
>
> *Jim Brancheau*
> *Director of Engineering*
> *Altair Engineering*

Manufacturers—setting the pace

Almost every major carmaker in the world owns one or more Cray Research supercomputers. Cray Research's automotive customers include:

- ☐ Audi
- ☐ BMW
- ☐ Chrysler
- ☐ Daihatsu
- ☐ Daimler-Benz
- ☐ Fiat
- ☐ Ford
- ☐ General Motors
- ☐ Hino Motors
- ☐ Honda
- ☐ Hyundai
- ☐ Isuzu
- ☐ Kia Motors
- ☐ Mazda
- ☐ Mitsubishi
- ☐ Nissan
- ☐ Opel
- ☐ Peugeot
- ☐ Porsche
- ☐ Renault
- ☐ Toyota
- ☐ Volkswagen

For years these companies and others have put the power of Cray Research's PVP systems to work on a host of applications that began with structural analysis, crash simulation, and computational fluid dynamics (CFD) but today also include metalforming, molding and extrusion, casting, engine combustion simulation, and acoustic and mechanical simulation.

Rules of the road

Being first to market, controlling costs, improving quality, and complying with stringent federal requirements regarding environmental factors and safety are the driving forces in today's automotive industry. Additional pressures include increasing numbers of concurrent car development programs and increased consumer demand for something fun to drive! To remain competitive, manufacturers must reduce the cycle time from new product design to delivery. Every development phase—design, engineering, component manufacturing, and final production—must be compressed. In some cases, automotive engineers are so confident in the designs and visualizations they get from their Cray Research systems that they are eliminating a majority of physical prototyping. That saves time and money. Despite a tough economy, all three U.S. auto manufacturers—Chrysler Motor Corporation, Ford Motor Company, and General Motors Corporation—made investments between 1992 and 1994 in upgrading their Cray Research systems; so did international customers Peugeot S.A., Daihatsu Motor Company, Ltd., Mazda Motor Corporation, and Régie Nationale des Usines Renault.

Ford, the first commercial company to order a CRAY C916 system, ordered its second CRAY C916 system in 1993. The system is used in part for new supercomputing applications in manufacturing simulation, including the simulation of sheet-metal forming of car components.

Fabricating dies to stamp sheet-metal body panels requires one of the longest lead times in an automotive development program. Each attempt can cost several hundred thousand dollars. Supercomputer analysis, replacing trial and error, can predict the metalforming characteristics of certain die geometries, considerably shortening the time required to design a die set that produces the right sheet-metal part the first time.

In addition to PVP, Cray Research offers customers and prospects a massively parallel supercomputer. A major automotive company was one of the first customers to order a CRAY T3D system, the world's first truly scalable heterogeneous MPP system, with a mix of PVP and MPP capabilities to address specific problems with unrivaled speed and efficiency. Since introducing the CRAY T3D system in September 1993, Cray Research has moved aggressively to expand the number of commercial and industrial application programs available for CRAY T3D systems. For example, General Motors Power Train Group is a participant in a $52 million, cost-shared program between Cray Research and the U.S. Department of Energy to develop scientific and commercial software for MPP systems in collaboration with the DOE's Los Alamos and Lawrence Livermore National Laboratories and 14 other leading U.S. industrial firms.

> *"Conducting design and analysis work on our own Cray Research system strengthens our credibility with customers and prospects in the automotive industry who often are Cray Research customers themselves."*
>
> *Colin Cox*
> *Senior Principal Engineer*
> *Hawtal-Whiting*

Calling all suppliers

"We work with many suppliers," said Richard Rossio, General Manager, Scientific Laboratories and Proving Grounds, Chrysler Corporation. "In fact, one of our suppliers has just handed us a model of a windshield with new material that we will work together to do validation testing on. I believe the automotive supplier market will be an important one for Cray Research. We are interested in working with Cray Research to develop this market to bring the suppliers up to speed with this technology. The new low-cost supercomputers from Cray Research are well-suited for this market."

A provider of services and custom software for all major American automobile companies, Altair Engineering keeps its four-processor CRAY EL system busy with crash and bumper analysis, structural analysis of vibrations and passenger fatigue, and structural optimization for various car shapes and sizes. This Troy, Michigan-based automotive engineering company installed a CRAY XMS system in 1991, upgraded to a CRAY EL system a year later, and advanced to the four-processor CRAY EL supercomputer in 1993. Altair will upgrade next to a CRAY J916 system. Jim Brancheau, Altair's director of engineering, says the company consistently chooses Cray Research systems because of their technical computing strength and the support Cray Research provides for Altair's high-end software. He also cites performance as a key factor. "The CRAY EL lets you run the huge problems. We've got all these hot rod workstations, but the Cray is more like a bulldozer," says Brancheau. "We just keep stacking the work up in front of it and the Cray plows through it when the workstations get bogged down."

Hawtal-Whiting, an engineering design consulting firm in the United Kingdom, has a CRAY EL98 supercomputer at its Basildon headquarters and is scheduled to upgrade to a CRAY J916 system in 1995. Hawtal-Whiting offers structural, acoustic, and crash analyses of new vehicle bodies. The CRAY EL98 system allows engineers at the firm to analyze substantially larger models in significantly less time. "Conducting design and analysis work on our own Cray Research system strengthens our credibility with customers and prospects in the automotive industry who often are Cray Research customers themselves," said Colin Cox, Hawtal-Whiting senior principal engineer.

Increasing the mechanical strength and safety of automotive seats, belts, and headrests is a primary application for the CRAY EL system at Keiper Recaro, a supplier of these components to the German auto industry. As cars become smaller and more fuel efficient, safety requirements become more challenging. Engineers at Keiper Recaro use MARC, a general purpose, finite element program for linear and nonlinear analyses, to simulate product performance during dynamic crash tests and to optimize the designs. Keiper Recaro design engineer Werner Welter said the fast CPU speed and large memory available with the CRAY EL system has significantly decreased product development time for new auto seats.

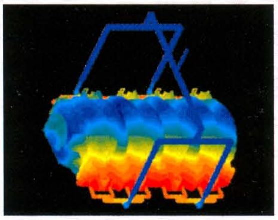

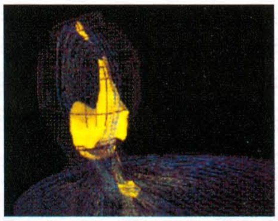

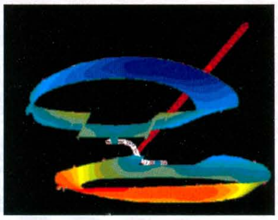

From top, a Kia Motors Corporation impact simulation with dummy; a MOLDFLOW filling and temperature analysis for an automotive intake manifold; CRI/*Turbo*Kiva analysis of a Nissan Motor Company engine; MOLDFLOW filling and warpage analysis of an automobile instrument cluster bezel.

To help customers capitalize on their supercomputing investments, Cray Research actively supports a wide variety of application software available from the leading commercial application vendors. This is a list of popular software in use on Cray Research systems in the automotive industry.

- ☐ ABAQUS
- ☐ ABAQUS/Explicit
- ☐ ADAMS
- ☐ ANSYS
- ☐ AVL-FIRE
- ☐ BEASY
- ☐ C-FLOW
- ☐ CFD ACE
- ☐ CM-Logic
- ☐ CRI/*Turbo*Kiva
- ☐ CSA/NASTRAN
- ☐ DADS
- ☐ EnSight
- ☐ FIDAP
- ☐ FLOW-3D
- ☐ FLUENT
- ☐ HEXAR
- ☐ KIVA (I, II, III)
- ☐ LS-DYNA3D
- ☐ MARC
- ☐ MOLDFLOW
- ☐ MSC/DYTRAN
- ☐ MSC/NASTRAN
- ☐ NIKE2/3D
- ☐ OPTRES
- ☐ PAM-CRASH
- ☐ PAM-STAMP
- ☐ PROCAST
- ☐ RADIOSS
- ☐ RAMPANT
- ☐ RASNA
- ☐ SPEED
- ☐ STAR-CD
- ☐ SYSNOISE
- ☐ UAI/NASTRAN

Automotive climate control simulation using CFD

Mark E. Gleason and Richard L. Sun,
Chrysler Corporation, Auburn Hills, Michigan
John M. Tripp, Cray Research, Inc.

Executive summary

Computational fluid dynamics (CFD) is presently the primary tool at Chrysler Corporation for prototyping the ducting of climate control systems. Computational modeling on Cray Research systems helps Chrysler engineers better understand flow characteristics and optimize designs to work more efficiently. This modeling process has improved product quality and performance while reducing design cycles and prototyping costs.

When Chrysler began exploring computational fluid dynamics (CFD) in 1985, the original charter was to use this technology to develop the aerodynamic qualities of the external surfaces of our cars and trucks. An examination of the available tools at that time made it quite clear that the external aerodynamics problem was impractical to solve in the near future due to the size of the problem and the need to maintain a quick turnaround time to keep pace with the design cycle. Therefore, our attention turned to the fluids and heat flow problems on our vehicles that were amenable to the then current level of CFD tools.

One area of great opportunity that we examined was that of the climate control system, including the passenger compartment of the automobile. At that time, the design of the ducting systems for heating, air conditioning, and defrosting was primarily a process of soft prototyping and observing the resultant performance. This process was time-consuming and produced highly variable

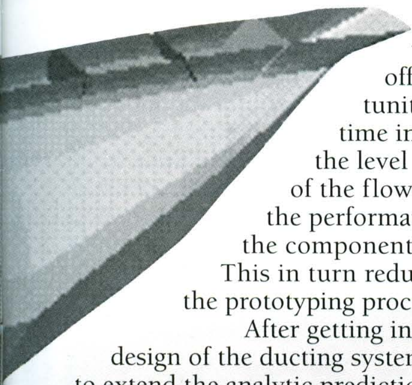

results. CFD offered the opportunity to reduce the time involved, increase the level of understanding of the flows, and improve the performance qualities of the components and system. This in turn reduces the cost of the prototyping process.

After getting involved in the design of the ducting systems, we were asked to extend the analytic prediction process to the defrosting of the windshield. This meant that we would be given a specific geometry of the defroster ducting and the passenger compartment (including the windshield and instrument panel) and be asked to predict the defroster ice-melting pattern and how it changes with time. This entire process would need to be done without the aid of any physical prototypes. This article presents some of the results of our initial studies.

Technical approach and tools

The first example is an isolated defroster duct for a minivan application. This project's objective was to improve the flow distribution, reduce the pressure drop, and minimize the wind rushing noise of the duct. The second example is a complete windshield de-icing simulation. The objective of this project was to validate a new and innovative prediction methodology for windshield de-icing. The model simulated is a 1993 LH sedan (Dodge Intrepid, Chrysler Concord, and Eagle Vision) passenger compartment, complete with windshield and defroster ducts.

Defroster duct simulation

We used the PATRAN code from PDA Engineering to generate the defroster duct geometry and finite element mesh. We used the FIDAP code from Fluid Dynamics International for the analysis. The interface between PATRAN and FIDAP was readily available and proven at the time this project started.

The geometry flexibility inherent in FIDAP's finite element algorithm is well-suited for complex shapes. To simulate the effects of flow turbulence, we used a two-equation turbulence model coupled with the special wall elements in the turbulent boundary layer. The model contained 31,000 elements. We used the segregated solution algorithm in FIDAP. For the solution phase, we ran FISOLV on a single CPU of Chrysler's CRAY Y-MP 8I system.

The EnSight code (formerly MPGS) from Computational Engineering International was used to display and animate the dataset obtained from the 3-D analysis. Dynamically scanning the 3-D flow field was very helpful for understanding the results and locating the areas of potential improvement. EnSight was run in a distributed mode with the front-end (graphics rendering) running on a Silicon Graphics Personal IRIS and with the back-end (compute intensive) running on the CRAY Y-MP 8I system.

Figure 1 (left). Minivan defroster duct geometry for the original model design.

Figure 2. Velocity vector distribution (side views) and velocity magnitude contours (top views) for (a) original design and (b) improved design.

Windshield de-icing simulation

We used the Control Data Corporation ICEM-CFD code to generate the multidomain mesh. Geometry was imported through the IGES graphics format from CAD surface data generated by the CATIA CAD application. The model contained 110,000 cells and comprised 50 domains.

The CFD-ACE code from CFD Research Corporation was used to solve the steady-state flow field and the transient ice-clearing process. CFD-ACE has been adapted to simulate the ice-clearing process by its customized capabilities in handling the complex 3-D transient, multimedium, multiphase conjugate heat transfer process.[1] To provide accurate results, we considered the airflow and heat transfer within the complete passenger compartment in the simulation. We implemented the multidomain finite volume discretization scheme in the code with a fully conservative and implicit interface treatment between domains.

We again used EnSight to postprocess the steady-state flow field and the ice-melting patterns at various times. EnSight can easily animate the transient temperature data obtained from the simulation. We then recorded the results of that animation on videotape. All of the figures in this article were generated with EnSight.

Defroster duct results

For the defroster duct design project, we modeled an isolated duct from a minivan defroster system. Figure 1 shows the original model geometry. The final design was arrived at after four design iterations, all of which were carried out on the computer. The original design and the final, improved design are compared in Figure 2. Note the highly nonuniform flow distribution at the defroster duct exit in the original design. The high-velocity flow is colored red, and the low-velocity flow is colored blue. Also note the dramatically improved flow distribution in the new design. Flow exiting the defroster duct is much more uniform, and regions of separated flow are almost eliminated.

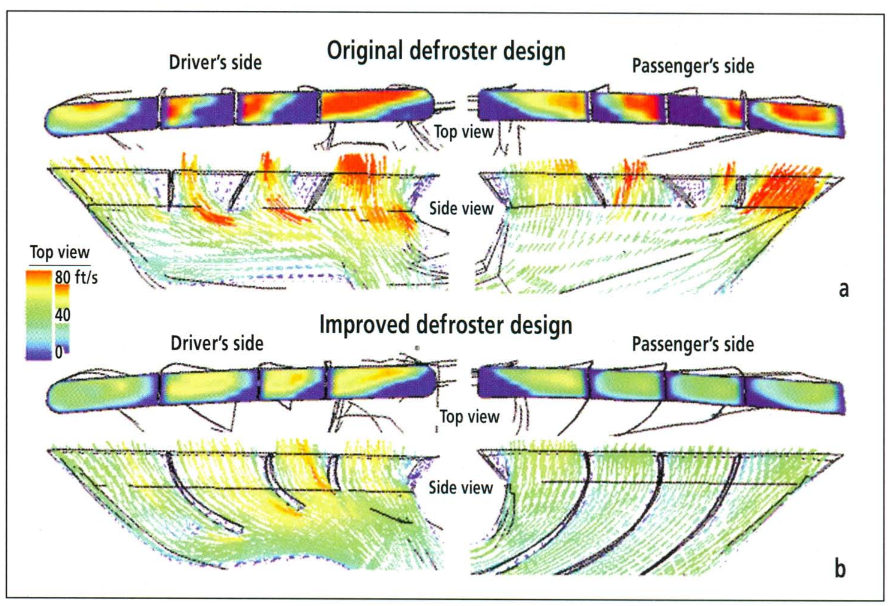

Defroster Flow Split		
Original Design		
	Test	CFD
Driver	49%	48.9%
Passenger	51	51.1
Final Improved Design		
	Test	CFD
Driver	53%	53.1%
Passenger	47	46.9

Table 1. Experimental and CFD-predicted flow split between driver and passenger sides.

Performance Improvement		
	Original	Final
Pressure drop	865 Pa	500 Pa
Maximum turbulence	100 m^2/s^2	75m^2/s^2

Table 2 (left). Predicted performance levels of original design and final, improved design.

Figure 3 (below). LH passenger compartment showing defroster system airflow. Blue-green-yellow areas show velocity contours just off the inside surface and red lines show particle traces released in the defroster duct.

The CFD-predicted flow splits are compared with experimental values in Table 1. Excellent agreement with the experimental data was obtained for both manifolds, original and improved. The driver's side of the improved design also has a larger percentage of the flow. This is a desirable characteristic that aids in meeting federal regulations. Table 2 shows the predicted levels of performance improvement. The pressure drop through the duct was reduced 40 percent, and the maximum turbulence levels were reduced 25 percent. The reduced pressure drop was later confirmed by experimental testing, and the reduced turbulence levels and reduced flow separation resulted in a quieter defroster system.

Windshield de-icing results

To model the windshield de-icing process properly, the entire passenger compartment was modeled. This technique accounted for the effect of flow recirculation in the passenger compartment. Figure 3 shows the complete passenger compartment of the 1993 LH modeled in this project. The ice-clearing simulation is actually a two-step process.[1] First, the steady-state flow field is obtained by iterating the Navier-Stokes equations until convergence is reached. After that, the flow field is held fixed, and the heat transfer equations are solved for the conjugate heat transfer through the composite windshield (glass-plastic-glass) and ice layer. The windshield is considered cleared once the ice has reached its melting point.

Figure 4 depicts the time history of the ice-clearing process. To model the de-icing process more accurately, the temperature of the defroster duct airflow varies with time. The air is initially at ambient temperature, and the temperature is increased with

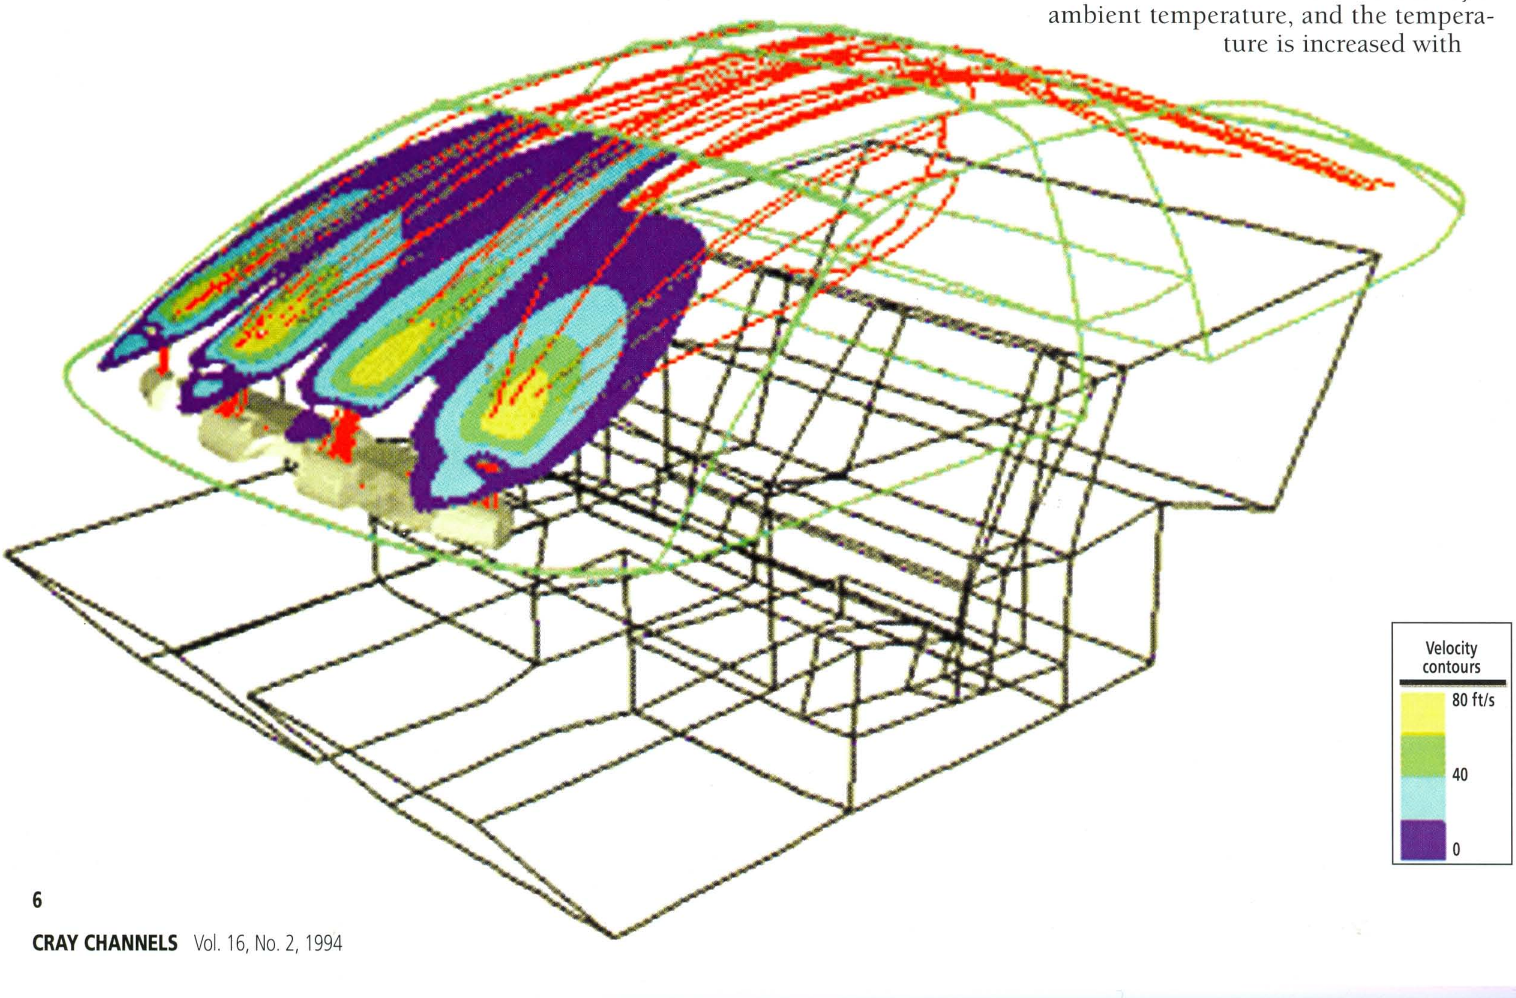

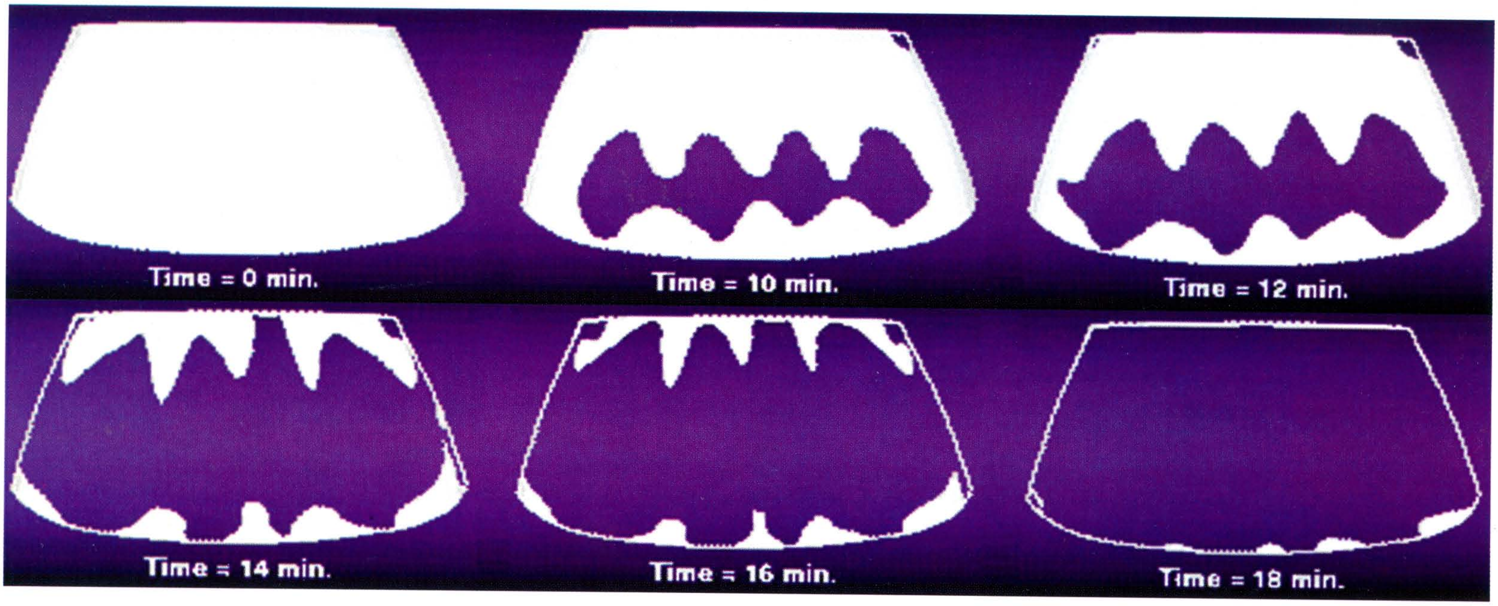

Figure 4. Ice-melting pattern predicted with CFD-ACE for 1993 LH sedan.

time. The air temperature as a function of time is known from engine warmup characteristics. The ice-melting time history agrees very well with experimental results.

One aspect of the validation is to compare the velocity distribution just off the surface of the inside of the windshield. Figure 5 shows predicted velocity distributions at a distance of 2 to 7 mm from the glass surface. This agreement is very reasonable considering the accuracy of measurements taken off the windshield glass and variations in the computational mesh.

Conclusions

CFD will play an increasing role as automotive companies continue to push for faster product development. CFD technology has been demonstrated as an effective design and development tool for automotive climate control systems at Chrysler. Rapid CFD job turnaround on Cray Research supercomputers can make an impact on the ever-shortening product development cycle by reducing prototype development time and costs. The wealth of flow field information that CFD provides enables engineers to make better design decisions earlier in the development cycle, and that leads to a higher quality product.

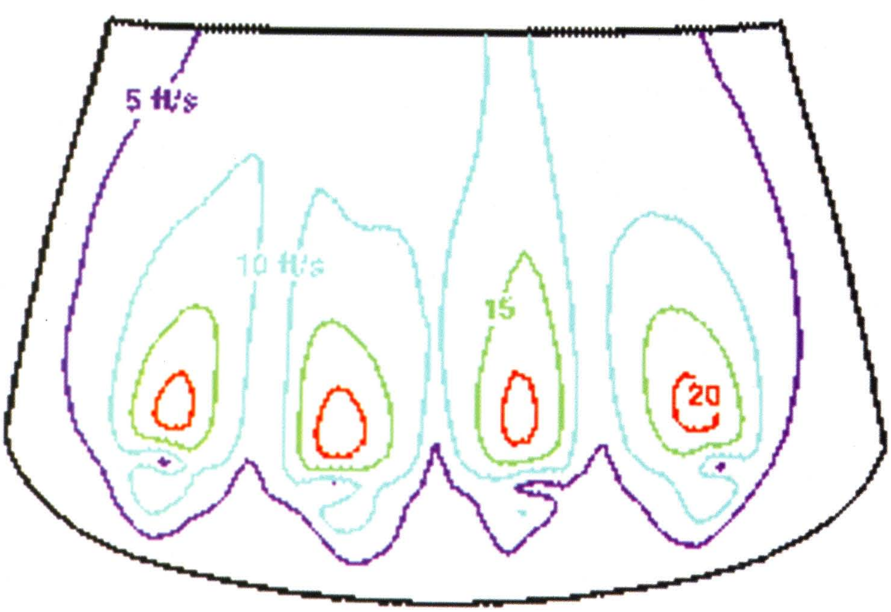

Figure 5. Predicted air velocity values, at 2 to 7 mm from glass surface.

Acknowledgments

The authors thank Hsien Lee of Bquad Engineering, Inc. for his efforts on the minivan defroster duct and windshield de-icing simulations, Merban Sioshansi of Chrysler Corporation for his efforts in the ice-melting simulation presented in this article, and John Hirshey and the Technical Computing Center staff at Chrysler for their technical assistance.

About the authors

Mark Gleason is the supervisor of the Aerodynamics and Fluid Dynamics Group at Chrysler Corporation. He has served 20 years in the automotive industry as an aerodynamics and fluids specialist working with Ford Motor Company and Automotive Aerodynamics, Inc. before joining Chrysler. His current responsibilities include supervising both experimental aerodynamics and the core group of computational fluid dynamics at Chrysler. Gleason received B.S. and M.S. degrees in automotive engineering from the University of Kansas.

Richard Sun is a CFD specialist in the Aerodynamics and Fluid Dynamics Group at Chrysler Corporation. Prior to joining Chrysler he worked for American Computing Corporation writing CFD software for automotive applications. He received his Ph.D. degree from Imperial College of Science and Technology, London.

John Tripp is a Cray Research CFD specialist who works onsite at the Chrysler Technology Center. He received an M.S. degree from the University of Arizona and a B.S. degree from the University of Wisconsin. Prior to joining Cray Research, he was a group specialist in CFD at General Dynamics, Fort Worth Division.

References

1. Lee, J.-G., Y. Jiang, A. J. Przekwas, and M. Sioshansi, "Automotive Windshield Ice-Clearing Analysis," SAE International Congress and Exposition, SAE Paper 930289, Detroit, Michigan, March 1993.

Numerical analysis of vehicle CRASHWORTHINESS in various configurations

James C. Cheng and Jiamaw Doong,
Light Truck Division,
Ford Motor Company, Detroit, Michigan

Executive summary

Using CAE analysis on a CRAY C90 system in the development of new vehicles helps the Ford Motor Company maintain a competitive edge and adhere to strict safety requirements. The enhancement of finite element software and the evolution of new-generation supercomputers, such as CRAY C90 systems, allow design engineers at Ford to use crash simulation as a practical tool that saves time and money.

Vehicle safety has received significant attention from consumers, governments, and vehicle manufacturers during the past decade. To emulate real-world accidents, various test configurations have been proposed and adopted. Designing the "Best in Class" truck remains our top goal at Ford Light Truck. The next-generation Ford minivan, Windstar, has to pass not only U.S. safety regulations and European export requirements, but also the more stringent Ford internal corporate guidelines. CAE analysis of Ford trucks on a CRAY C90 system plays an important role in the vehicle development process, encompassing safety requirements and competitive issues.

Although various tools, such as the occupant kinematic code,[1] spring-mass code,[2] and simplified beam code,[3] are used within Ford to evaluate overall crash performance, this article covers only finite element analyses using RADIOSS.[4] RADIOSS is a transient dynamic code that considers both material and geometric nonlinearity. An explicit time integration scheme is adopted. A typical

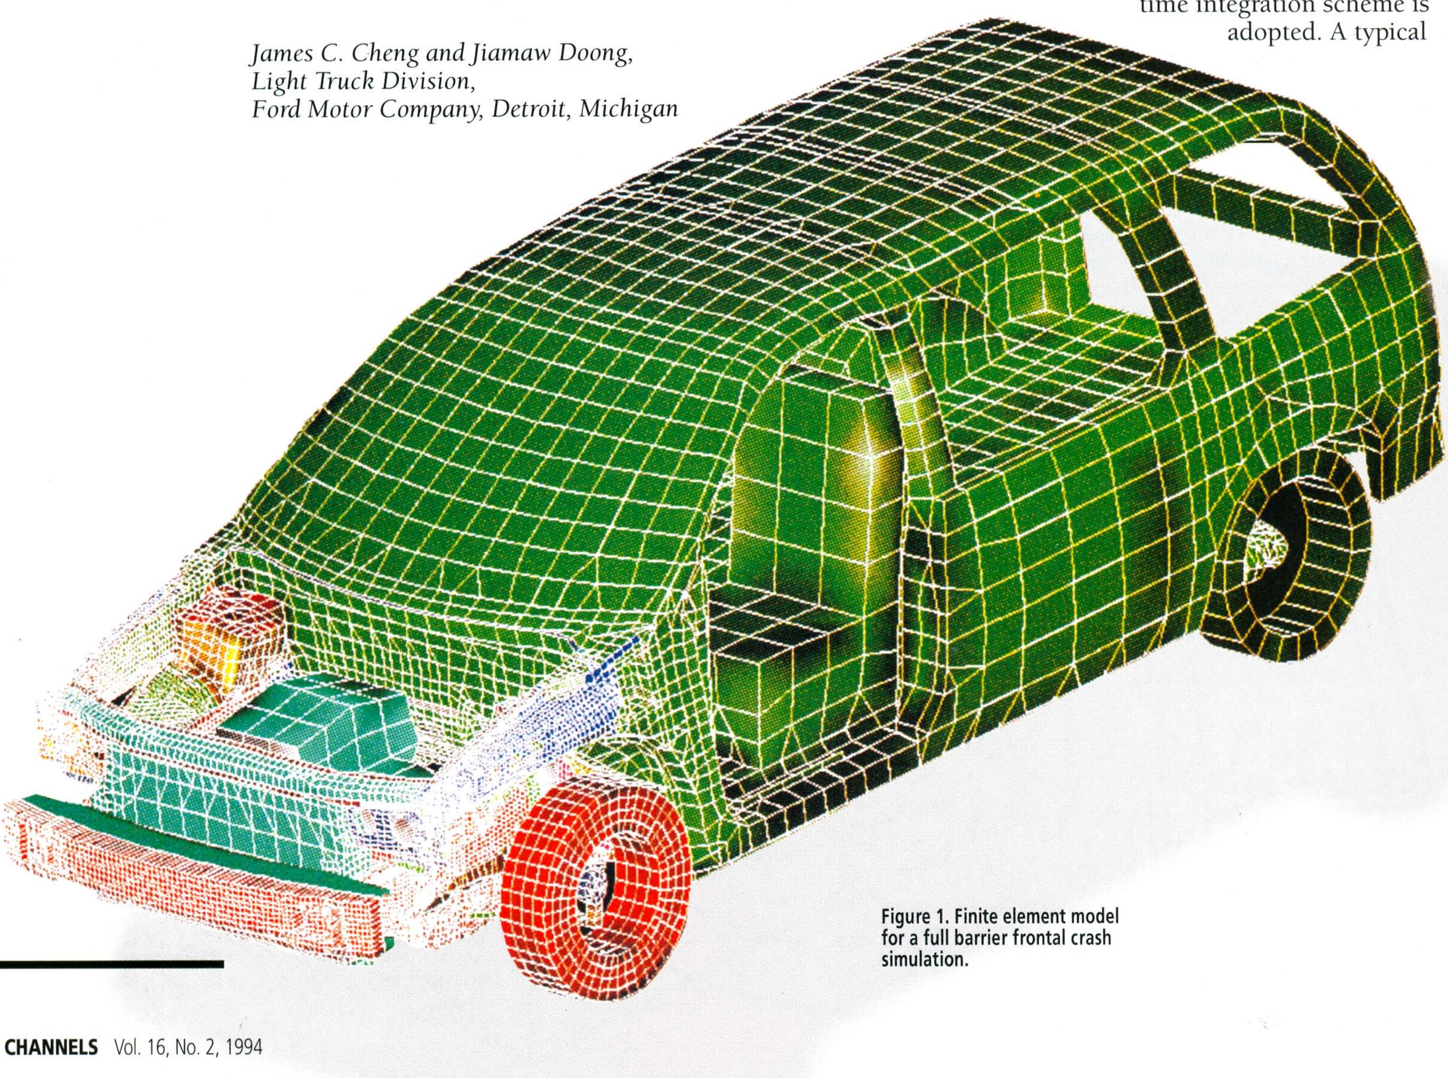

Figure 1. Finite element model for a full barrier frontal crash simulation.

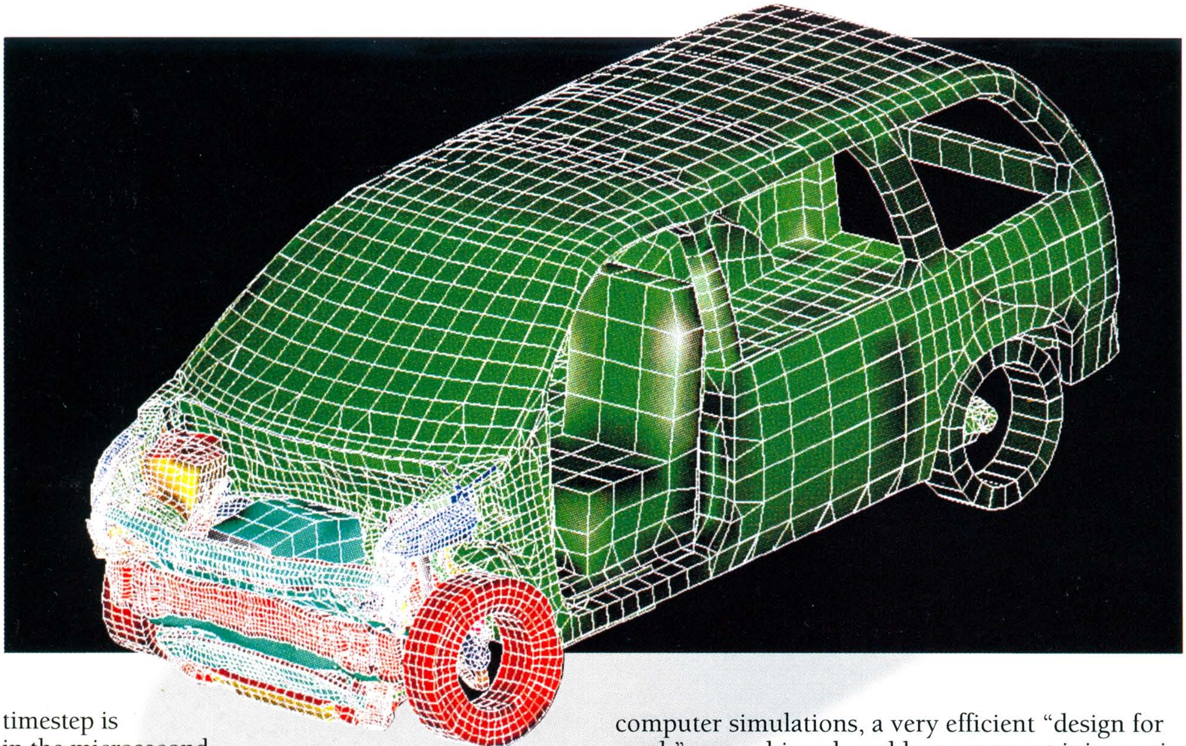

Figure 2. Crash mode of a full barrier crash at 35 mph.

timestep is in the microsecond range. Minimum timestep control is necessary to conserve computer processing time.

The models are built on workstations by using Ford's in-house software, PDGS, or commercial software such as Altair's Hypermesh or PDA Engineering's PATRAN. Thin shell and solid elements compose the majority of the model. Nonlinear spring and truss elements are used to model the nonmetallic parts, or to simplify the model. Finer mesh is necessary in the area of server deformation. Because of the different test configurations described above, an analyst either can build different models for different analyses or use a single model of finer mesh and sustain severe penalty of extra CPU time. Ford uses the first approach.

The number of elements for a typical crash model ranges between 20,000 and 60,000. The jobs were run on a CRAY C90 supercomputer located at the Ford Engineering Computer Center. Between 30,000 and 50,000 CPU seconds are required to complete a full system simulation, the results of which are shown in Figure 1.

Full barrier frontal crash

All manufacturers must pass the 90 degree full barrier Federal Motor Vehicle Safety Standard test (FMVSS) 208. In terms of vehicle structural design, it is crucial to generate a "soft pulse" to satisfy the occupant injury criteria. The kinetic energy is absorbed through the collapsing and bending of longitudinal members. All front-end structures such as shotgun, front rail, and engine cradle are modeled in detail. The engine is modeled as a rigid body with proper geometric definition. The rigid wall option in RADIOSS is used to model the barrier, and all possible contact zones are enforced. Figure 2 shows the deformed shape of the Windstar vehicle for a 35 mph New Car Assessment Program test. The displacement and velocity at the b-post are compared with the test as shown in Figure 3. With the aid of computer simulations, a very efficient "design for crash" was achieved, and low occupant injury criteria were observed from the test.

50 percent offset frontal crash

The 50 percent offset frontal crash test configuration was suggested through some European manufacturers and consumer magazines.

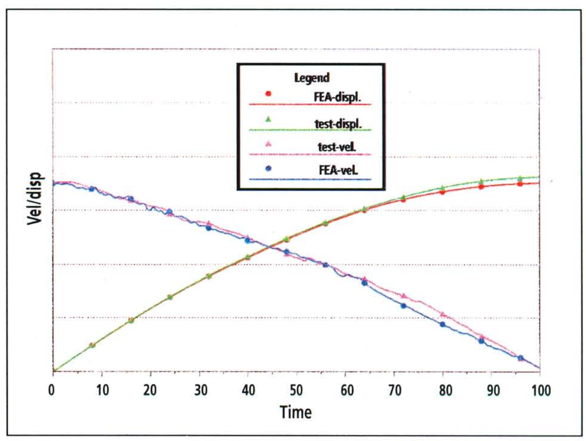

Figure 3 (left). B-pillar displacement and velocity comparisons between a full barrier crash and simulation at 35 mph.

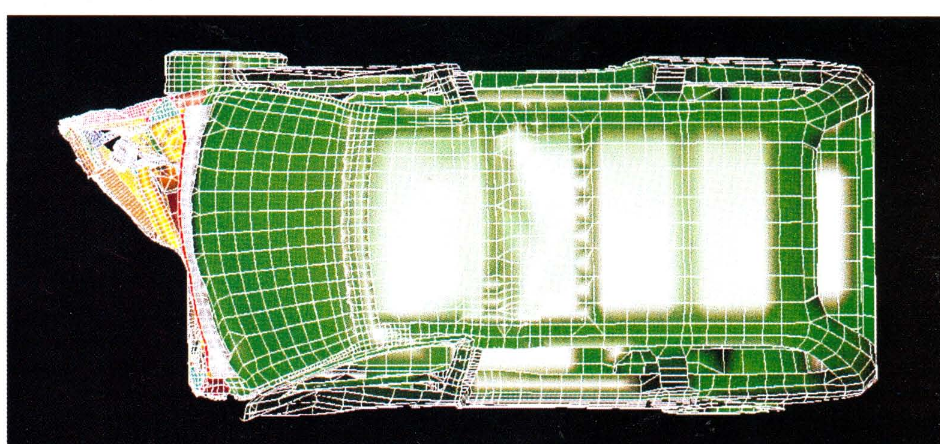

Figure 4 (below). Crash mode of a 50 percent offset barrier test at 35 mph.

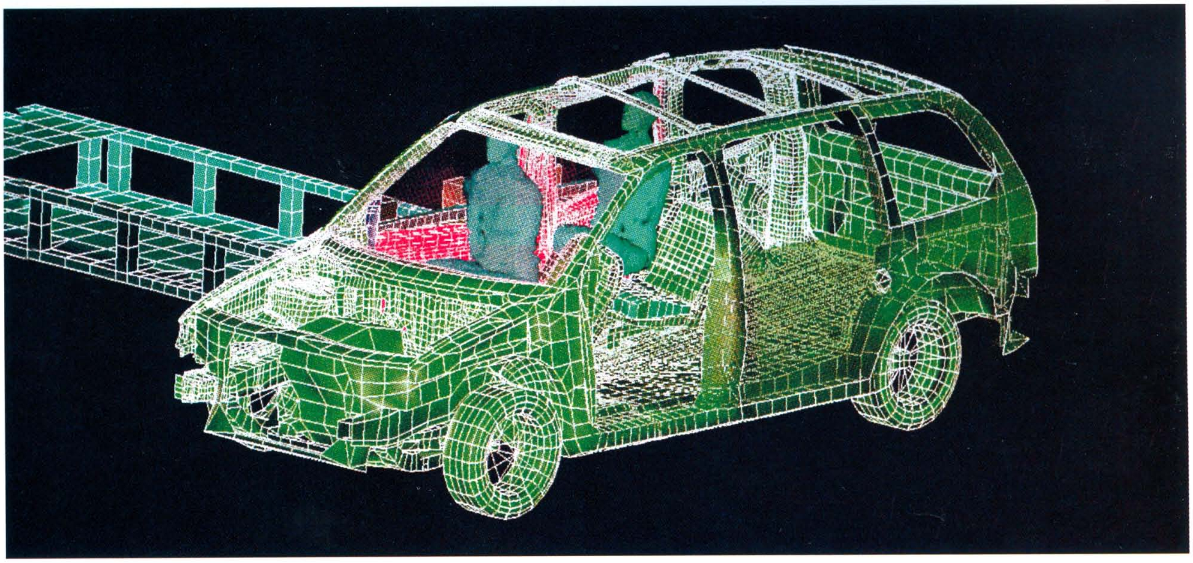

Figure 5. Crash mode of a dynamic side impact test with two SID dummies.

This test might be considered supplemental to the full barrier test in future rulemaking. Unlike the full barrier test, energy absorption is concentrated on the driver's side. Because more crush distance is experienced in the driver's side, the left side of the model is refined all the way through the b-pillar area. The front floor and dash also are modeled in detail to assess the engine intrusion. Figure 4 shows the post-crash deformation.

Dynamic side impact with deformable moving barrier

Although no dynamic side impact requirement exists for light truck vehicles, Ford has decided that a Windstar will comply with regulation FMVSS 214, which currently applies only to cars. A deformable barrier with a honeycomb front face hits the target vehicle at a crab angle of 27 degrees. National Highway Traffic Safety Administration (NHTSA) side impact dummies (SID) are placed in the front and rear seats of the struck side of the target vehicle. As shown in Figure 5, the side structures such as the door, pillars, rocker, and underbody cross members are modeled in detail to obtain the proper structural response. A special honeycomb material model was developed in RADIOSS to properly represent the property of aluminum honeycomb. Two finite element SID dummies were placed in the model. Figure 6 gives comparisons of barrier motion between CAE and the test. Windstar passed the car dynamic side impact standard.

The scope of analysis

The following list outlines the major analyses performed by Ford Light Truck Crash CAE.

Analysis	Purpose
30 mph fixed barrier frontal crash	FMVSS 208 regulation
35 mph fixed barrier frontal crash	NCAP test
50 percent offset fixed barrier frontal crash	European regulatory agencies/publications
30 degree angular fixed barrier with and without anti-slide device	FMVSS 208 and European regulatory agencies/publications
Dynamic side impact	FMVSS 214 regulation; European recommendation
Rigid moving barrier rear impact	FMVSS 301 regulation
50 mph car-to-car side impact	Corporate guideline
50 mph car-to-car rear impact	Corporate guideline
Roof crush	FMVSS 216 regulation

Car-to-car impact

Car-to-car impact requirements are part of Ford's corporate safety guidelines. To investigate fuel leakage during and after a crash, a Taurus is used as a "bullet car" that crashes into a stationary Windstar at 50 mph. The test modes include side impact, rear impact, and rear offset impact. Although the current CAE technology has difficulty predicting fuel tank rupture or fuel leakage,

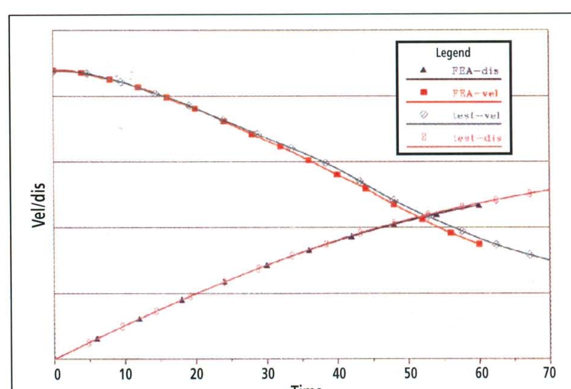

Figure 6. Barrier CG motion comparisons between test and simulation.

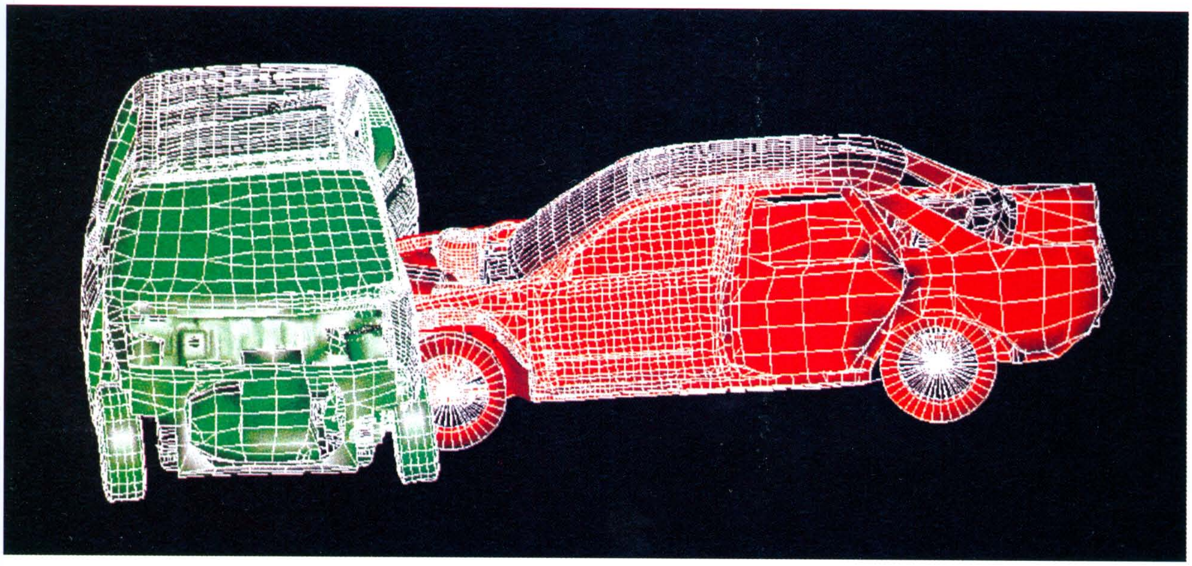

Figure 7. Car-to-car side impact simulation.

tank intrusion can be easily identified by inspecting the crash mode. Because the accuracy of the Taurus model was crucial for the analysis, it underwent careful validation through correlations with full barrier and 50 percent offset barrier tests. Figure 7 shows the crash model of a side impact analysis at 65 ms; no fuel tank intrusion can be observed.

Roof crush analysis

Regulation FMVSS 216 requires that light truck vehicles have resistance to more than 1.5 times their unload vehicle weight (UVW) within 5 inches of roof crush distance. The force was applied to both front edges of the roof through a flat plate with tilt angle. Because this is a pseudotest, a very slow velocity was applied to the flat plate to reduce the dynamic effect. Significant CPU time is required if the enforced speed is too slow. Special features were developed by RADIOSS—such as a mass scaling, inertial relief, and dynamic relaxation—to enhance the code efficiency. Because the windshield dominates the roof crush resistance, special care was taken in modeling the windshield. Figure 8 shows the crush mode of Windstar. Comparison of the peak force between the test result and analysis, shown in Figure 9, showed the difference to be within 5 percent.

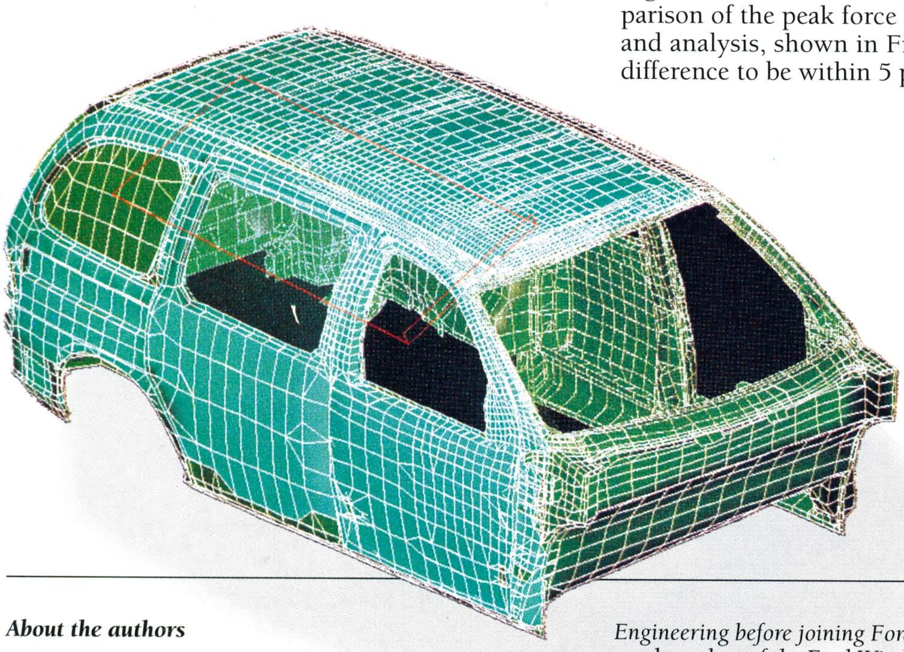

Figure 8 (left). Roof crush analysis.

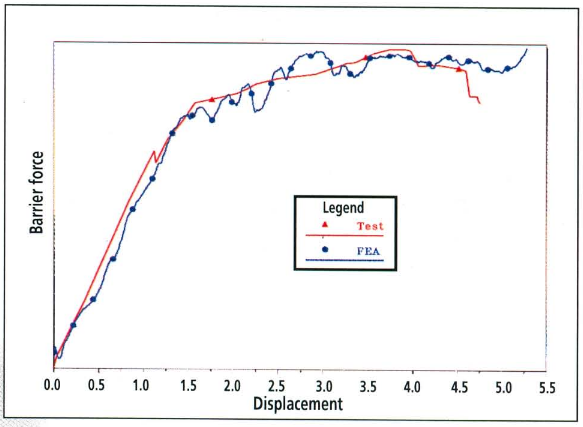

Figure 9 (below). Roof crush resistance comparison between test and analysis.

About the authors

James C. Cheng is supervisor of the Ford Light Truck CAE section. He received an M.S. degree from the University of California-Berkeley in 1982 and a Ph.D. degree from Purdue University in 1986. He joined Ford in 1988 and since has concentrated on crash CAE analysis of light trucks.

Jiamaw Doong is a crash technical specialist in the Ford Light Truck CAE section. He received an M.S. degree from the University of Missouri-Rolla and a Ph.D. degree from The Ohio State University. He was senior engineer at EASI Engineering before joining Ford in 1991. He was chief CAE crash analyst of the Ford Windstar project.

References

1. MADYMO user's manual, TNO Road-Vehicle Research Institute, The Hague, The Netherlands.
2. CRUSH 2 user's manual, Ford Motor Company, Detroit, Michigan.
3. BEAT user's manual, Ford Motor Company, Detroit, Michigan.
4. RADIOSS user's manual, Mecalog S.A.R.L., France.

MOTORCYCLE simulation at BMW aims at *quieter world*

Pierre Guisset, Numerical Integration Technologies NV, Leuven, Belgium
Alfred Irrgang, BMW AG, Munich, Germany

Executive summary

Faced with the dual challenge of increasingly stringent pass-by noise requirements, such as the European standard ISO 362, and consumer demand for powerful, smooth "superbikes," BMW is changing its basic approach to motorcycle design. The traditional sequential development approach that often resulted in costly retrofitting at the prototype testing stage late in the design process is being replaced with a concurrent vehicle development and optimization process based on computer simulation. Using the acoustic radiation prediction software package SYSNOISE by Numerical Integration Technologies (NIT), BMW engineers can test design alternatives at the soft-prototype stage to predict the acoustic radiation patterns of a motorcycle.

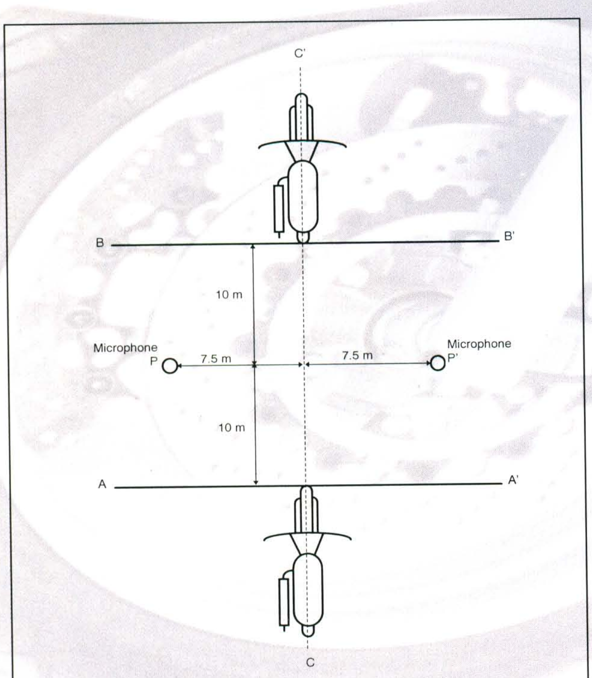

Figure 1. European standardized test procedure for motorcycle pass-by noise measurements.

Anyone whose sleep has been disturbed by the noisy roar of a motorcycle or the shrill whine of a moped passing through the night will appreciate the efforts of design engineers, software developers, and Cray Research to make motorcycles significantly quieter. BMW Motorrad GmbH, the German automotive manufacturer's motorcycle division, is renowned for its silky-smooth "superbikes"—powerful motorcycles that already provide a level of acoustic refinement and sound quality that few can match. To improve its products even further and to comply with increasingly stringent pass-by noise requirements, BMW is working with the acoustic radiation prediction software package SYSNOISE and its developers at Numerical Integration Technologies (NIT).

During the development of a vehicle one of the main challenges for design engineers is the resolution of vibro-acoustic problems such as overall noise and vibration levels, harshness, and production variability. Because of the difficulty of modeling dynamic performance to date, undesirable dynamic performance often went undiscovered until final assembly, necessitating extensive testing of the physical prototype, or a range of prototypes, to identify the underlying cause of the problem. Generally, this "test-analyze-fix" approach results in extensive and costly additions for acoustic control. Compounding the challenge, government legislation is calling not only for a dramatic reduction in pass-by noise levels but also a reduction in the weight of the vehicle to reduce CO_2 emissions. This conflict in design options is made even more challenging by current consumer preferences that are shifting the acoustic target from a mere reduction in overall sound levels to shaping the sound to meet subjectively desirable characteristics.

This has resulted in a refocusing of approaches and technologies used in acoustical design, from sequential development and a "refinement pipeline" approach to engineering desirable vibro-acoustic profiles up front that are then applied during a concurrent vehicle development and refinement process. The goal of today's methodology is to enable engineers to systematically and consistently optimize vibrational and acoustic behaviors from concept to completion. The ability to link design variables to interpretable acoustic phenomena allows engineers to understand the effect of design alternatives at the soft-prototype stage. A key factor in this process, and the focus of this article, is the prediction of the acoustic radiation patterns of a motorcycle at the design stage.

Meeting the ISO 362 test for pass-by noise

Legislation calls for vehicles to comply with strict pass-by noise levels that must be measured under normalized test conditions. In Europe, the standard is ISO 362, which is similar to testing procedures in force in other areas of the world. During this test the driver must drive to line A-A´ at a constant speed of 50 km/h (see Figure 1), then accelerate at full throttle to line B-B´, first in second gear and then a second time in third gear. During the acceleration from A-A´ to B-B´, the sound pressure levels are measured at the two microphones P and P´. Today the noise level may not exceed 77 dB (A)

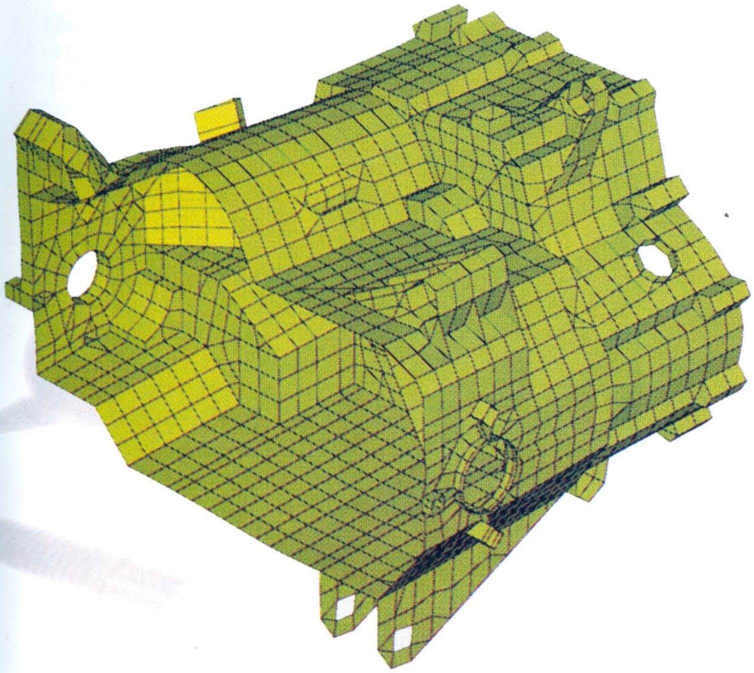

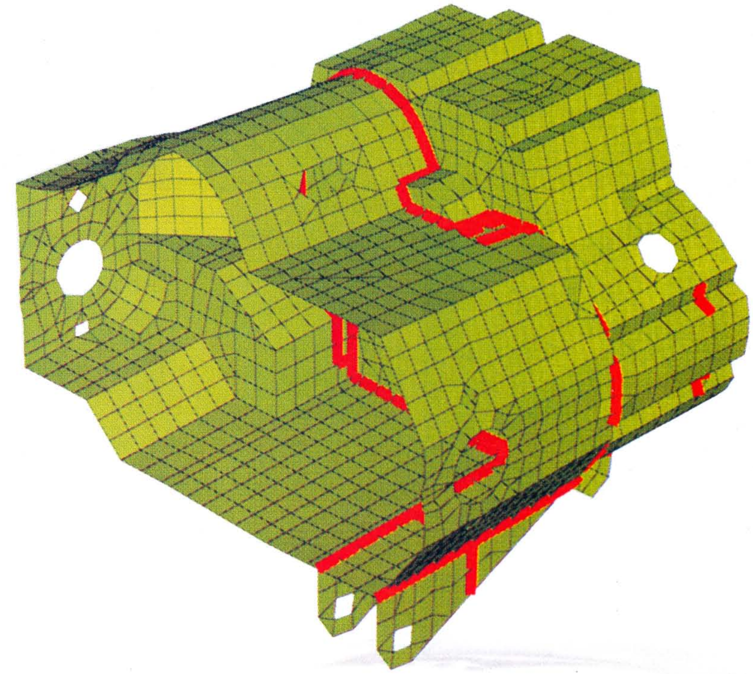

at any time during the test drive, but by 1997 this level will be even more stringent, reducing it by a factor of two, or –3 dB (A). The challenge for vehicle manufacturers is clear: to stay in business they must optimize the design of the dominant noise sources and minimize overall noise levels.

Figure 2 (left). Structural finite element model of the gearbox, defined using I-DEAS and computed with MSC/NASTRAN.

Figure 3 (right). Acoustical boundary element (BE) model showing junction lines in red.

Modeling the physical phenomenon

Testing indicates that the gear housing of the motorcycle is normally the dominant sound source, especially between 700 Hz and 1400 Hz. This is due mainly to the high dynamic response of the gears and to the open bell-shaped casing in particular. The contact forces at the gear teeth and the rotation of the gear wheels induce vibrations in the gear shaft that are transmitted to the case, which then radiates sound energy. This behavior is known as one-way coupling: the structural vibrations generate pressure waves (noise) in the surrounding air, which can be reasonably assumed not to modify the structural behavior of the source. Accordingly, the structural and acoustic models can be decoupled during the overall modeling process and become relatively straightforward. The first stage is to model the vibration of the housing by a structural analysis, typically a finite element (FE) model. The second stage is to use the predicted velocity boundary conditions as input to the acoustic modeler, which, in this case history, is based on the boundary element (BE) method.

Predicting vibration levels using the FE model

The first step consists of a structural dynamic analysis to determine the vibration levels in the frequency range of interest. An FE model is created using the I-DEAS program (see Figure 2) while the analysis is performed with MSC/NASTRAN. The loading conditions for the analysis were taken as harmonic excitation at the shaft mounts, which is a reasonable representation of the actual dynamic behavior of the gearbox. A modal frequency response (solution sequence 111) was performed to find the structural vibrations (punch velocity file). Because the quality of the FE model is critical to the success of the subsequent modeling processes, the FE model was previously validated by correlating the analytically derived eigenmodes and eigenfrequencies with those measured by experimental techniques.

Predicting acoustic radiation pattern using the acoustic BE model

The BE model is built from all radiating surfaces of the structural FE model. Because of the computational power of the CRAY Y-MP system, the mesh did not need to be coarsened, eliminating a tedious and time-consuming step. The structural dynamic behavior defines the velocity boundary conditions for the acoustic radiation analysis performed with the SYSNOISE software. The model characteristics (involving junctions and openings) led to selecting the indirect BE method as shown in Figure 3. This method leads to a symmetric system of equations so that the required memory and computational times are reduced.

Computational times

The analyses were performed on the CRAY Y-MP 8/4128 system at the BMW computing center. The original BE model involved 1827 nodes and 1899 elements. Because the BE method involves fully populated matrices, special attention has to be paid to computational times. In a BE analysis the computing time is typically spent during the assembly of the coefficient matrix (which has to be repeated for every frequency in the desired acoustic spectrum) and the solution of a linear system of equations. Several actions can be taken to improve the model.

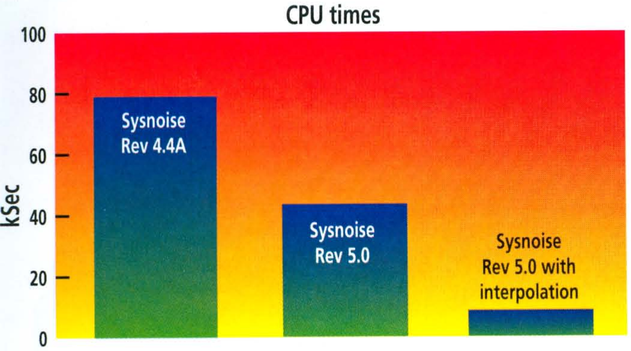

Figure 4. Computing times comparison for a 100 frequencies analysis on BMW's CRAY Y-MP system.

First of all, it is useful to know that the indirect BE model can handle meshes with inconsistent normal vector orientation by adding artificial junction relations—although clearly at the cost of additional nodal degrees of freedom. In this case, the automatic junction relation generation of SYSNOISE built 981 junction relations, leading to 2884 degrees of freedom. Reference computing times with SYSNOISE Rev 4.4A were 667 seconds for matrix assembly and 102 seconds for a system solution.

SYSNOISE Rev 5.0 provides an automatic optimizing algorithm for the element normal vector orientation consistency, which reduces the number of junction relations to 184 and the model size to 2172 degrees of freedom. This allows a CPU-time reduction of 21 percent for the matrix assembly (525 seconds) and of 59 percent for the system of equations solution (42 seconds). Further code improvements (especially vectorization) reduce these compute times by another 23 percent for matrix assembly (404 seconds) and by 33 percent for system solution (28 seconds).

As can be seen, the matrix assembly is the most time-consuming operation in a typical BE analysis running on the CRAY Y-MP system. Therefore, a spectacular improvement can be obtained by using new methods of matrix frequency interpolation. The user defines master frequencies, in which the coefficient matrix is explicitly computed, and slave frequencies, in which the coefficient matrix is ingeniously interpolated between the two bounding master frequencies. If the master frequencies are reasonably selected, this methodology produces very accurate results. It provides a dramatic reduction in computational time, with only a small penalty due to I/O operations. For a slave frequency, the matrix assembly typically uses 14 seconds of CPU time.

A realistic analysis in the relevant frequency range (500 to 1500 Hz) will include about 100 frequencies. The original model (SYSNOISE Rev 4.4A) then will need 100 x (667 + 102) = 76,900 seconds of CPU time. Model and source code optimization (SYSNOISE Rev 5.0) reduced the required CPU time to 100 x (404 + 28) = 43,200 seconds (44 percent reduction). Matrix frequency interpolation with 10 master frequencies (leading to insignificant result difference) still reduces CPU time to 10 x (404 + 28) + 90 x (14 + 28) = 8,100 seconds—an order of magnitude faster (see Figure 4).

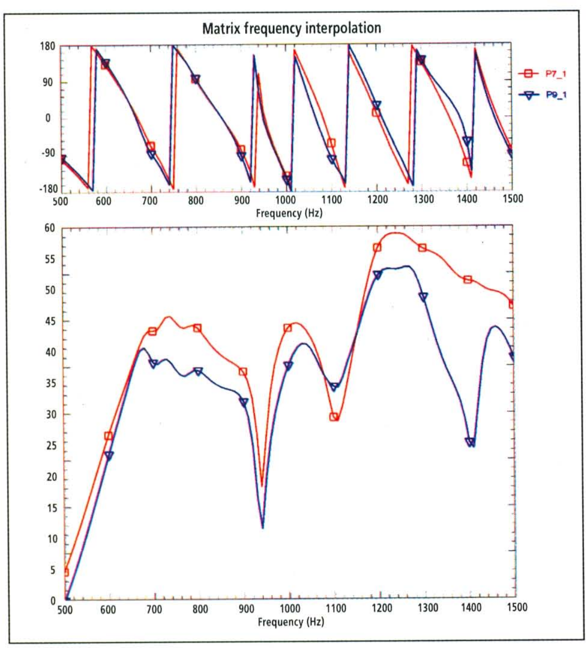

Figure 5 (below left). Frequency response functions (dB (A)) at microphone P, with the standard matrix evaluation procedure (red curve with square markers) and with the matrix interpolation algorithm, using 11 master frequencies (blue curve, with triangle markers). Markers show the master frequencies. Results are almost identical.

Virtual pass-by noise analysis results and design optimization

Two field points are defined on each side of the gearbox to represent the pass-by noise test microphones; a hemispherical field surface and a plane are described to help visualize the gearbox's 3-D acoustic pattern. Results are shown either as frequency response functions (see Figure 5) or as a contour plot of the acoustic radiation pattern on the field surface (see Figures 6a and 6b). As can be seen from the frequency response functions, the radiation shows a peak between 1200 and 1300 Hz, which is also confirmed by experiments.

Determining a good, even optimized, structural modification for improved acoustic radiation is not an easy task.

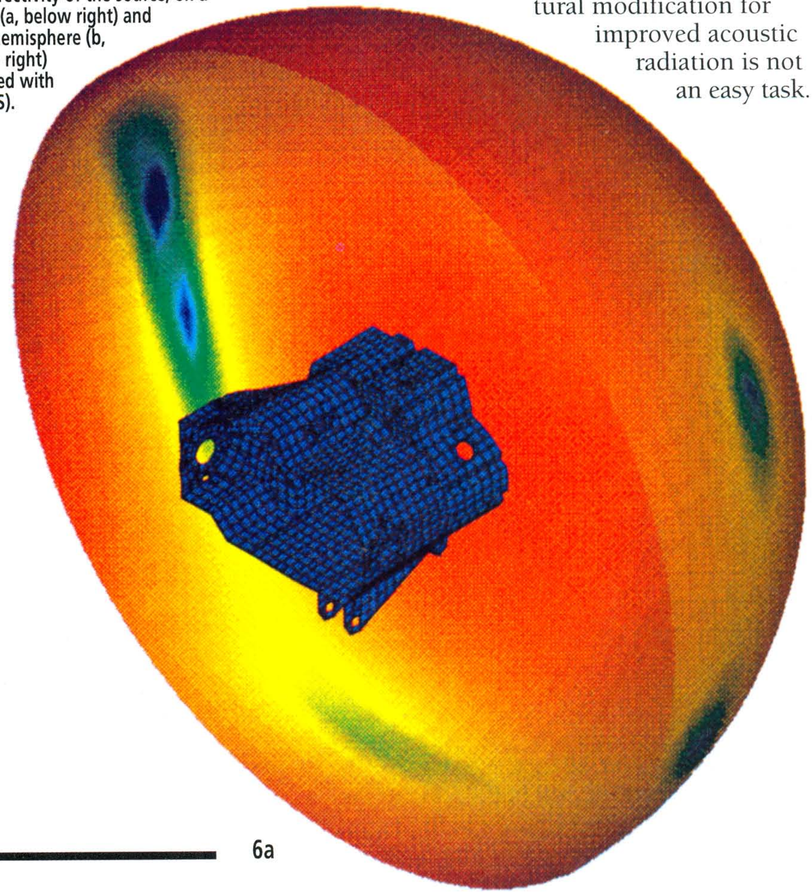

6a

Figure 6. Acoustic radiated field at 1240 Hz (dB (Lin)), showing the directivity of the source, on a plane (a, below right) and on a hemisphere (b, above right) (plotted with I-DEAS).

The most convenient procedure is to define the possible structural modifications (for example, the thickness of some parts) and to link a structural design sensitivity analysis (how the vibration changes if the structural design parameter changes) to acoustic sensitivity analysis (how the sound pressure level changes if the surface velocity changes) in a so-called global acoustic sensitivity analysis.

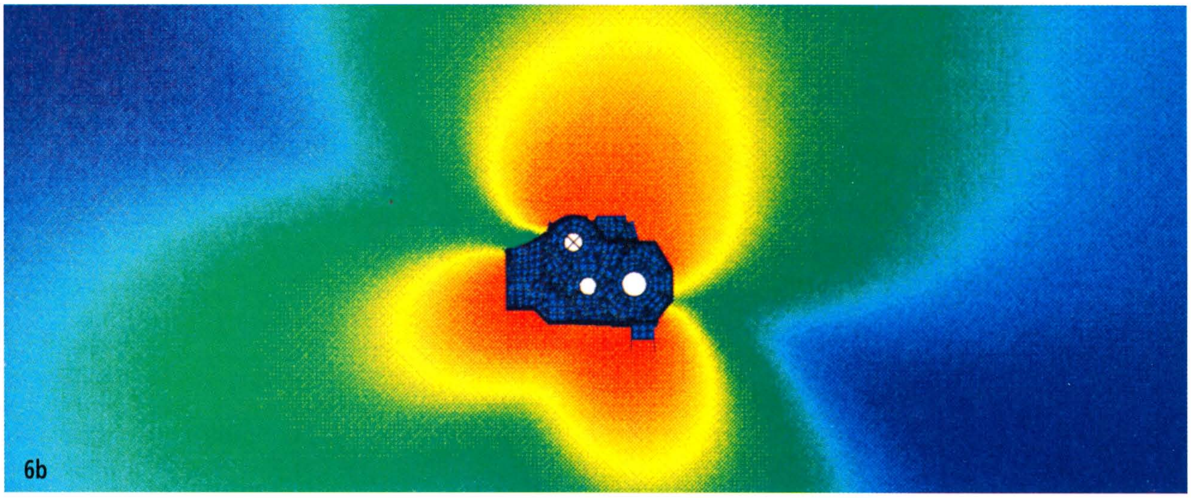

From these results, structural modifications can be proposed and tested. In this particular case, because of the dynamic behavior of the gear housing, a good modification was to add a few stiffeners on the circumference of the open end. Figure 7 shows a noise reduction of more than 3 dB (A) at the microphone locations.

Conclusion

Numerical simulation tools enable the prediction of acoustic patterns of new products very early in the design process. Design engineers can then make well-informed concept-level decisions, examine more candidate designs, and produce quieter, more refined products more efficiently. The ultimate goal is to avoid the flawed concept altogether. Acoustic radiation prediction is now a mainstream tool, complementing and building on existing tools for structural analysis in the designer's list of options. The stage is now set for the integration of these tools into an automatic multicriteria optimization procedure, using already available design sensitivity information and algorithms.

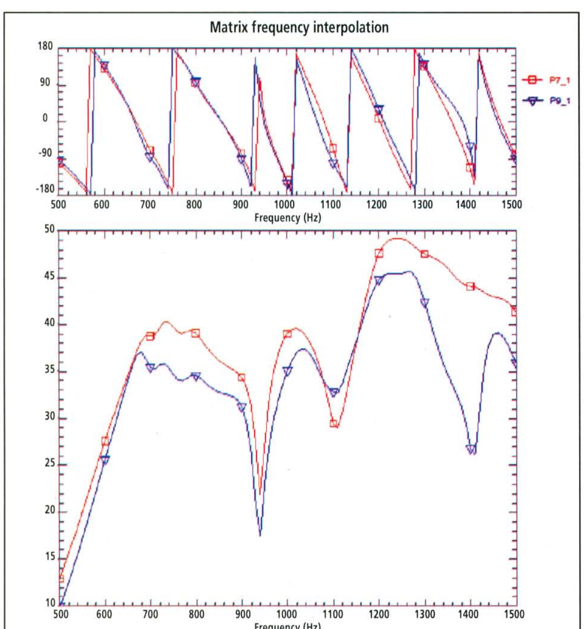

Figure 7. Frequency response functions (dB (Lin)) at microphone P, comparing the original design (red curve with square markers) and the modified design (blue curve with triangular markers), showing an improvement of more than 3 dB (A) at the peak frequency.

However, the sheer scale and complexity of the computation problem cannot be underestimated. Optimization runs for complex structural systems such as aircraft, cars, or motorcycles will be possible only on powerful supercomputers such as CRAY Y-MP systems.

Acknowledgments

The authors thank BMW Motorrad GmbH for giving permission to publish this study and for providing all of the materials, including experimental and numerical models. They also thank Cray Research, Inc. for its continuous support in improving the numerical simulation tools.

About the authors

Pierre Guisset specializes in numerical simulation methods and design optimization algorithms. From 1980 to 1986 he was research assistant at the University of Louvain, Belgium. He started working on the development of the SYSNOISE software in 1987 and has been SYSNOISE product manager at Numerical Integration Technologies since 1991.

Alfred Irrgang studied mathematics and computer science with special focus on numerical mathematics at the Technical University of Munich. Since 1987 he has been working on numerical simulation methods and focusing on acoustic modeling at the BMW Research Center.

References

Coyette, J. P. and J. L. Migeot, "An Overview of Boundary Element Approaches for Acoustic Modelling," Inter-Noise 93, Leuven, Belgium, August 24-26, 1993.

Guisset, P. and J. P. Coyette, "A Global Acoustic Sensitivity Procedure for Structural Optimization," Inter-Noise 93, Leuven, Belgium, August 24-26, 1993.

Migeot, J. L., P. Guisset, and S. Sarti, "Using SYSNOISE for Solving Practical Car Acoustics Problems," Second International Conference on Vehicle Comfort: Ergonomic, Vibrational, Noise and Thermal Aspects, Bologna, Italy, October 14-16, 1992.

Von Estorff, O., L. Belke, and A. Irrgang, "Boundary Elements for Noise Radiation Analysis of Vehicle Components," 6th VDI Congress "Numerical Analysis in Automotive Engineering," September 21-23, 1992.

Supercomputer simulation of automotive sheet-metal forming

Chris Galbraith and Mark Finn,
Kingston Research and Development Centre, Alcan International Ltd., Kingston, Ontario, Canada

Executive summary

To help automakers "get it right the first time" when forming sheet-metal parts, Alcan engineers have developed computational models of the metalforming process using Cray Research supercomputers. By predicting wrinkles and splits in aluminum car fenders and the physical response of parts when removed from the forming equipment, Alcan engineers are able to assist their automotive customers in designing manufacturing equipment and processes from the outset. This efficient design methodology eliminates the need for expensive and time-consuming retooling. Moreover, as automakers move from steel to aluminum for many parts, computational modeling could emerge as an essential tool for assessing the adjustments to manufacturing equipment needed to accommodate the new material.

In recent years, North American automotive companies have undertaken aggressive efforts to improve product quality and decrease the time to market for their products. Manufacturing simulation, in which computers are used to simulate manufacturing processes such as sheet-metal forming, is one methodology with the potential to achieve these dual goals.

By simulating the manufacturing process and rendering the simulation predictions graphically, engineers are better able to understand the sheet-stamping process. This improved understanding can lead to process improvements that in turn yield higher quality parts. Gaining such an understanding is particularly important when parts are to be formed from nontraditional materials. In the past, steel has been used to form virtually all stamped components in cars, but the drive to reduce vehicle weight (and hence fuel consumption and emissions) has led to the substitution of aluminum for steel in an increasing number of applications. Because die designers have much less experience in designing tooling for forming aluminum components, manufacturing simulation can be particularly effective for facilitating sheet-forming operations for aluminum.

This article presents the results of a finite element simulation of an automotive sheet-stamping operation.[1] As part of a pilot project investigating the use of aluminum for sheet stampings, Alcan worked with the Ford Motor Company to develop stamping tools for producing aluminum fenders. The fender geometry was that of the 1992 Mercury Sable. The original tooling geometry (used to form steel fenders) was modified to form fenders from aluminum.

In conjunction with this practical exercise, Alcan engineers at the Kingston Research and Development Centre attempted to simulate the forming of the aluminum fenders to evaluate the effectiveness of simulation as an engineering tool. The evaluation was based on achieving an acceptable solution time for a model with more than 100,000 elements and the ability to analyze successfully the simulation results and present them to engineers and designers. Additional criteria included the ability to predict the major problems found in sheet stampings, namely splitting, wrinkling, and the part's overall shape after springback.

Using the finite element (FE) code LS-DYNA3D[2] on CRAY Y-MP and CRAY C90 supercomputers, the engineers were able to model the first draw operation. This work demonstrated the ability to predict accurately the location and

Figure 1. Contour of Von Mises stress during binder wrap: (a) Contact occurs at point [1] causing a stress wave to march up the dogleg (b). (c) Secondary contact occurs at point [2], causing wrinkles [3] and [4], shown in (d).

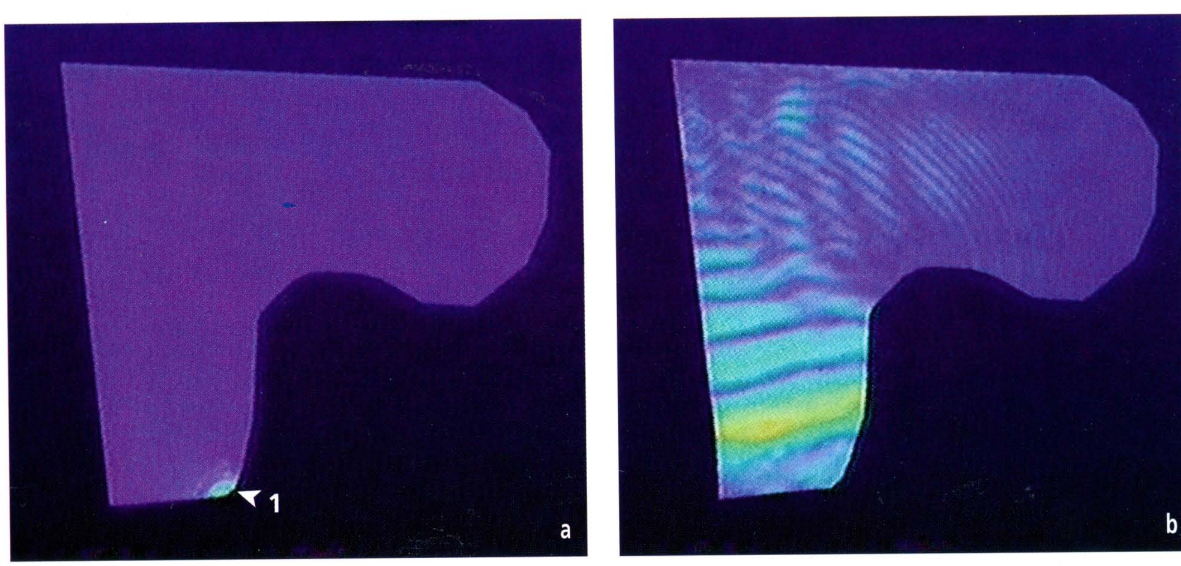

size of wrinkles, the presence or absence of splits, and the final fender shape after springback. From the results of this work, it is clear that the FE method is a useful tool for aiding in the design of aluminum automotive sheet stampings.

CAD to code: generation of an FE deck

For applying the FE technique to real-world manufacturing scenarios, Alcan has developed a philosophy of modeling the actual process as closely as possible. In the case of sheet-metal stampings this means replicating the real tooling geometry as accurately as possible. In accordance with this philosophy, care is taken to ensure that the model description of the tooling geometry is identical to the geometry of the actual forming tools. Automotive panels are typically complex shapes with many fine geometric features, and these features must be included in the model because they strongly influence the forming outcome.

Within Ford, the Product Design Graphics System (PDGS) is used to generate CAD surfaces of the forming tools.[3] These surfaces are used to generate cutter paths for manufacturing the tools on a numerically controlled (NC) milling machine, and these surfaces were used in the forming process simulation to generate an FE mesh of the tooling shape.

Ford's PDGS CAD database was converted to a neutral VDA format, imported into ICEM CFD/CAE, and converted to ICEM's internal format.[4] Graphical examination of the surfaces revealed several flaws. For instance, adjacent surfaces that should have shared a common edge often overlapped or were separated by a gap. Upon completion of surface cleanup, the mesh generation began. Mesh regions called subfaces were created between curves that lay on the tooling surfaces. When the subfaces have been defined, the mesh is projected onto the CAD surfaces. ICEM CFD/CAE lets analysts interrogate the mesh for a wide range of problems, including degenerate elements, reversed element normal orientation, warpage, and high aspect ratio. Bad elements are shaded in a different color and hence can be easily distinguished.

Once the mesh is projected on the tooling surfaces, an FE input deck can be created. The deck must include descriptions of the aluminum material's behavior under load, its interaction with the tools, and movements of the tools (loads and displacements). Once the deck is created, the FE model is ready to run.

Analysis: running the deck using LS-DYNA3D

LS-DYNA3D is a commercially available FE code marketed by the Livermore Software Technology Corporation. The code features robust contact algorithms and is highly vectorized to run efficiently on Cray Research supercomputers. Unlike many general purpose FE codes, LS-DYNA3D is a dynamic explicit code that uses an explicit time integration scheme based on the central difference method. This numerical method has several implications for the analysis. The use of four-noded quadrilateral elements makes it possible to diagonalize the mass and stiffness matrices, leading to a very efficient means for solving the governing equations. Essentially, the equations become decoupled, removing the need for inverting matrices or iterating to a solution. To maintain the accuracy lost by moving from higher- to lower-order elements, it is necessary to use a more refined mesh than for the high-order element case.

The FE model used for the fender-forming simulation comprised 130,000 nodes and a comparable number of elements. The tooling elements were sized to define the tooling geometry accurately. Because the tools are relatively rigid compared to the aluminum blank, they were modeled with rigid (nondeforming) elements.

The aluminum blank was modeled using 30,000 Belytschko-Tsay shell elements with six degrees of freedom per node and five integration points through the thickness. The material model chosen to represent the aluminum came from the LS-DYNA3D material library. Material type 24 (piecewise linear isotropic plasticity) was chosen as a reasonable approximation to the behavior of the AA6111 alloy used to form the fender. The flow curve was determined from a hydraulic bulge test

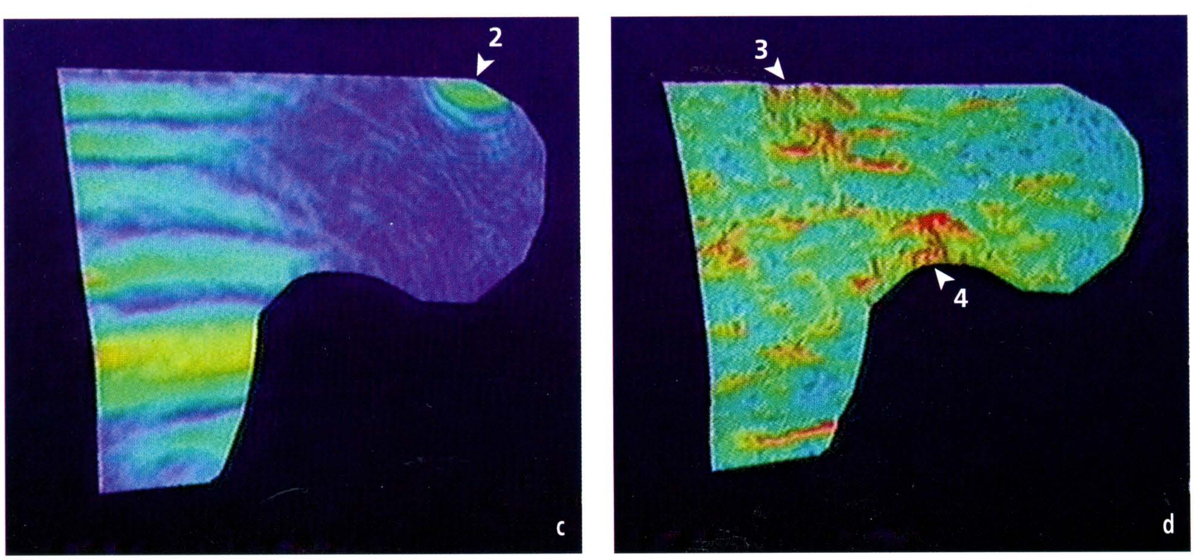

Figure 2. Contours of effective plastic strain during the draw: (a) the fender after binder wrap; (b) post contacts sheet at point [1], spreads to point [2], then up the door line toward the hood line; (c) near end of draw; and (d) fully drawn panel. The model predicts good evolution of contact resulting in acceptable strains within the body of the fender. The higher strains in the addendae surface may cause difficulties later when they are flanged.

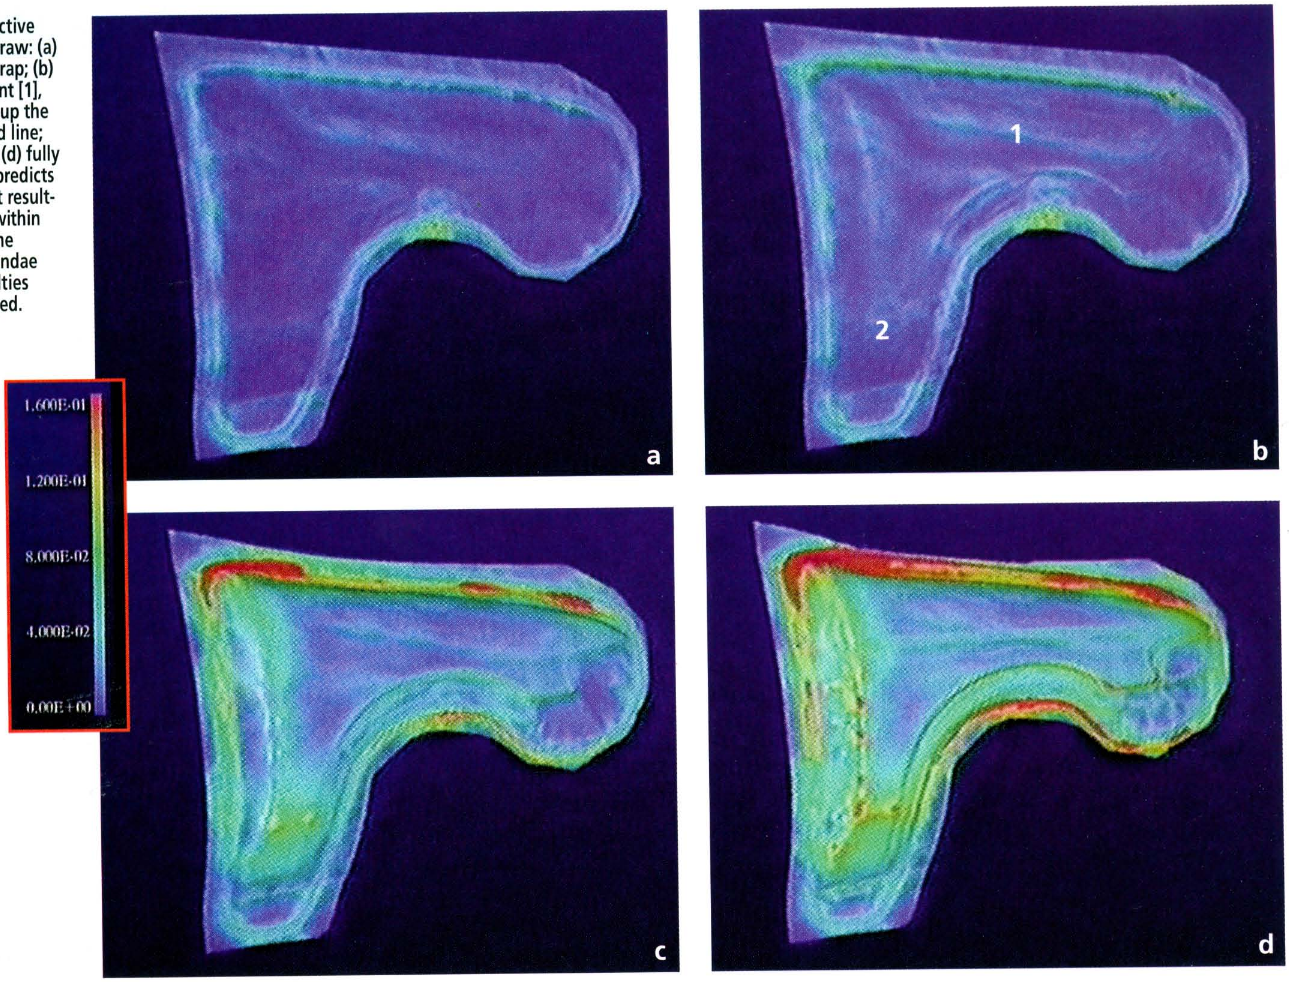

enabling strain values near 50 percent to be obtained. Friction was modeled using the standard Coulomb friction theory, with a friction coefficient of 0.08, which was determined experimentally.

The central difference method used in LS-DYNA3D is not absolutely convergent. Instead, a stable timestep is chosen such that $Dt < Dtcr = 2/w_{max}$, where w_{max} is the highest frequency of the system. LS-DYNA3D calculated the timestep automatically.

The stable timestep can be shown to be roughly equal to the time it takes for a sound wave to travel the length of the smallest deformable element in the model. As such, explicit codes such as LS-DYNA3D are inherently suited for modeling phenomena that occur in the millisecond time domain. Sheet-forming operations typically occur more slowly, taking a number of seconds to complete.

Because the stable timestep for explicit codes is so small, the number of timesteps is quite high. To model the forming of the fender with a natural time scale could take as many as one million timesteps. However, by increasing the velocity of the tools, the model runtime can be reduced accordingly.

LS-DYNA3D is a dynamic FE code; therefore the analyst must be careful about scaling the tooling velocities.[5] If too large a scaling factor is chosen, observable dynamic effects can be introduced. Essentially, if the tools are moving too rapidly, the sheet's mass prevents it from responding quickly enough. This could lead to erroneous predictions of stress and strain. However, carefully scaling the tooling velocities can reduce CPU costs without significantly affecting the accuracy of the solution. For the fender model, the maximum tooling velocity was increased from about 1 to 10 m/s. Experience with other simulations has shown that this was unlikely to affect the accuracy of the solution. The resulting runtime for the forming process was 64 hours on a single processor of a CRAY Y-MP system. On the CRAY C90 system, this time was reduced to 22 hours on one processor. With the parallel version of LS-DYNA3D, the elapsed time could be reduced even further, to approximately 5 hours.

Ideally, it should be possible to incorporate changes in the model to correct forming or springback problems and to have a new result in about two days. This sort of turnaround time is about equal to what is required to make minor changes in the actual prototype tooling—in fact, a quick turnaround would allow for the effects of various changes in the tooling geometry to be explored before actual changes, based on the model predictions, are made to the tooling.

Although the cost of 24-hour runs may seem large, it is comparable to that of crash analyses. In addition, the development and cost during prototype would be greatly reduced, and hopefully a more robust manufacturing window would be developed for the stamping plant. Die design times, which comprise a major portion of car production lead time, could be shortened. The ability to explore and optimize forming conditions prior to finalizing the die design could reduce the risk, both real and perceived, associated with the introduction of aluminum stampings.

Results: model predictions

The FE simulation consists of three stages. In the first stage, the binder wrap, the binders move down and clamp the aluminum sheet with a force of 895 kN. The binders close more quickly than the real event (10 m/s versus 1 m/s) to reduce CPU time. In the second stage, the post moves up into the sheet, forming the panel. Once again the tooling velocity has been scaled by a factor of ten. In the third and final stage, the tooling is removed, and the panel is allowed to elastically spring back to its final shape.

Binder wrap stage

During binder wrap, the binders (the outer tools) close, clamping the sheet. For most automotive panels, the binders are nonplanar surfaces and therefore bend and stretch the sheet. Ideally, the shape of the blank after binder wrap is free of wrinkles.

Figure 1 shows four stages in the binder wrap process of the fender. Contours of Von Mises stress are plotted on the deforming sheet. The yield stress of the aluminum alloy was 165 MPa, and all of the red-colored material has yielded. Figure 1a shows the first contact. Because the first contact between the blank and the upper binder occurs at a single point, the blank twists, resulting in a stress wave that proceeds up the "dogleg" of the fender at an angle (Figure 1b). Secondary contact occurs at the opposite end of the fender (Figure 1b), resulting in a second stress wave. These stress waves meet in the middle of the fender, resulting in a pronounced wrinkle above the wheel well (Figure 1c). When the binders are fully closed, the imprint of the draw bead can be seen (Figure 1d).

The binder design chosen leads to two problems. The first occurs when the upper binder contacts the blank at an angle. The resulting twist is not favorable; it predisposes the dogleg to problems during springback. The second problem is the wrinkle in the wheel well area. This wrinkle results from the double curvature, which causes contact to occur from outside to inside. Fortunately, this wrinkle is stretched out in later stages and does not protrude into the class A surface of the fender.

The first draw stage

After the binders have closed, the punch completes the forming of the fender. During this stage, the blank is stretched and drawn into the shape of the fender. Some of the metal outside the draw bead will be pulled inward and through the bead. The shape of the bead will determine which areas will draw in the most. By altering the profile of the draw bead in certain areas, the die designer can control the strains in the panel and the amount of wrinkling that will occur.

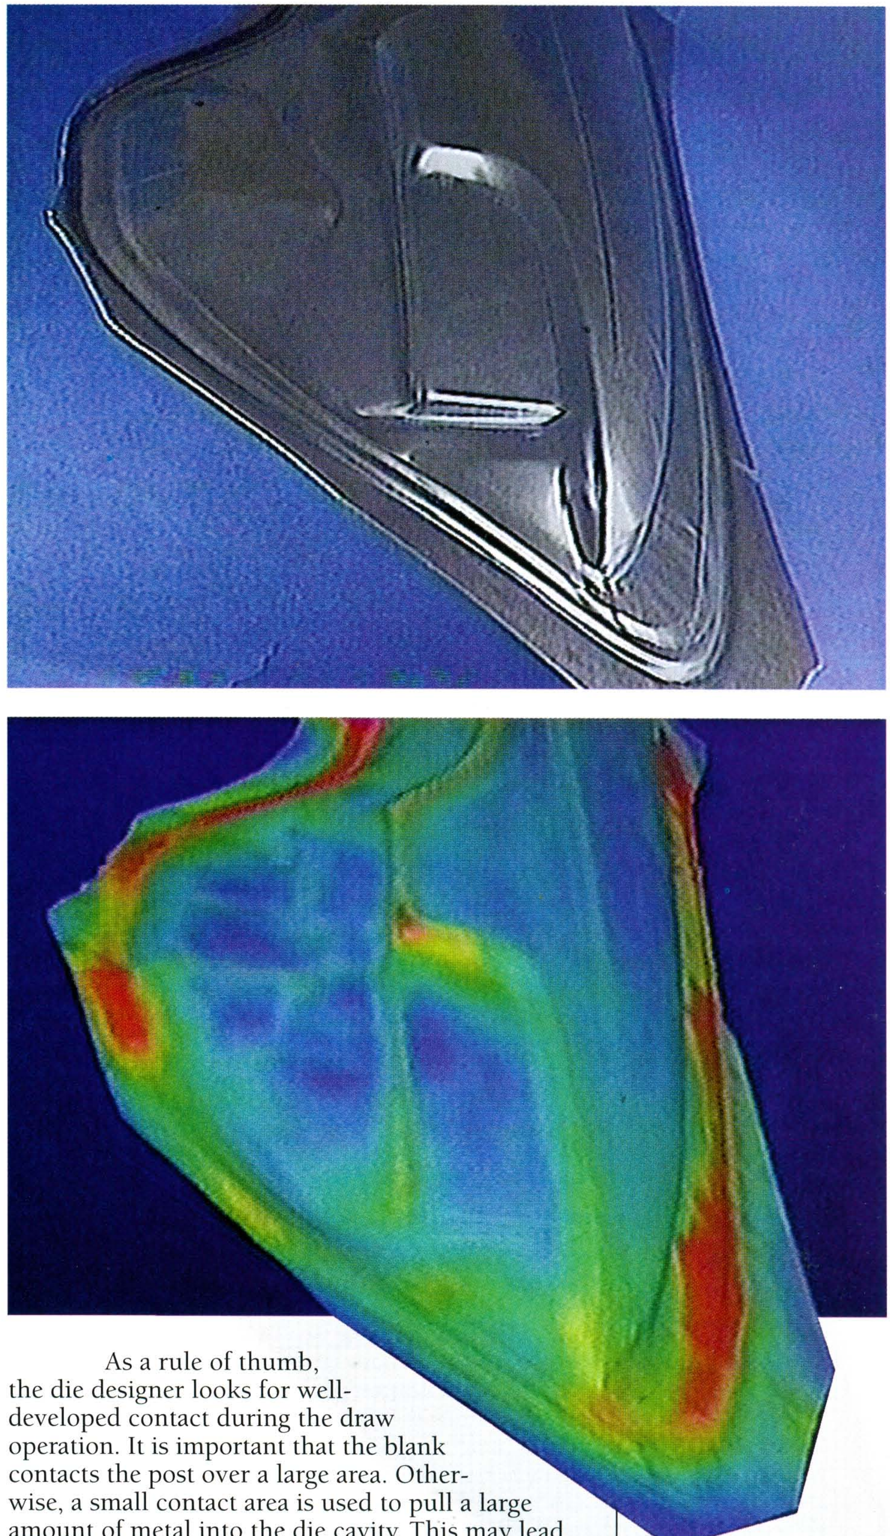

As a rule of thumb, the die designer looks for well-developed contact during the draw operation. It is important that the blank contacts the post over a large area. Otherwise, a small contact area is used to pull a large amount of metal into the die cavity. This may lead to splitting of the blank, as the stresses become too large. Figure 2 shows that the draw process is well-designed for the fender. Contours of effective plastic strain are plotted. Initial contact between the blank and post occurs above the wheel well (Figure 2b) and rapidly spreads toward the door line (Figure 2c). At the conclusion of the draw, the strains are seen to be small enough that the fender will not have any splits, but large enough (Figure 2d) that springback can be controlled. Figure 3 compares the real fender panel and the FE-predicted panel shape.

Figure 3. Comparison of fender (top) and finite element model.

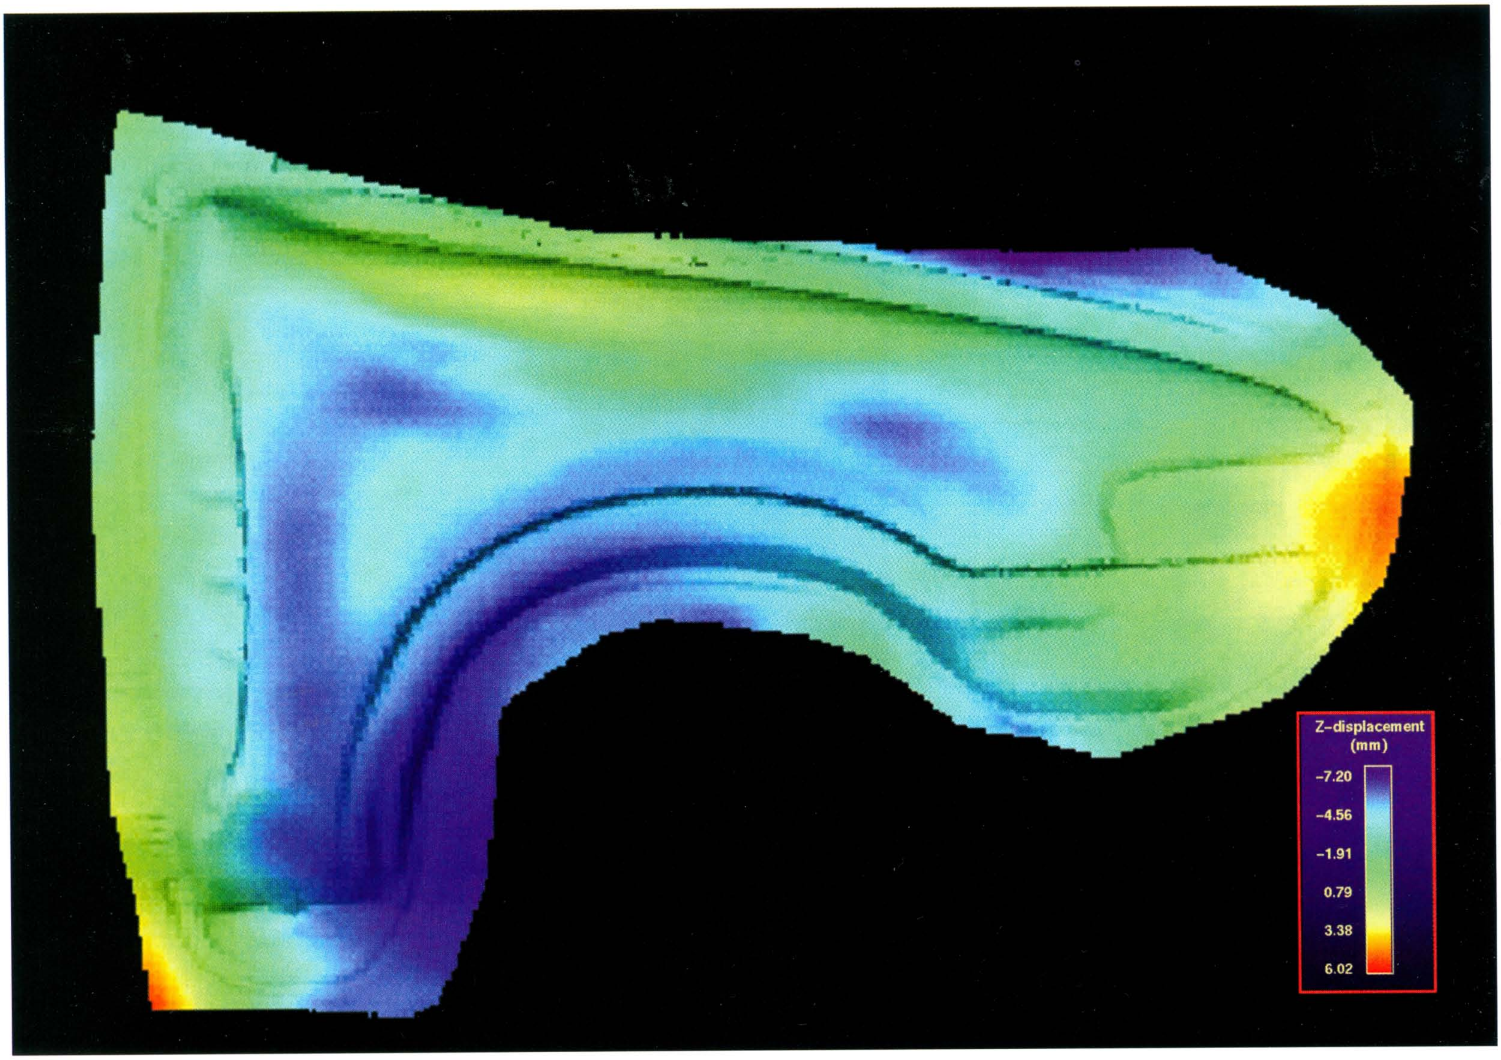

Figure 4. Contours of z-displacement resulting from springback of fender.

Springback stage

When the drawn panel is removed from the tools, all of the forces the tools had been exerting on the panel during the draw are removed. As a result, the panel changes shape so that all of the internal stresses become balanced. This process, called springback, changes the shape of the panel. The amount of change is principally related to the Young's modulus of elasticity, the yield strength, and the plastic strain in the panel. Generally, higher strains reduce the amount of springback, by helping to "set" the panel, as do lower yield strength and a higher Young's modulus. Because the Young's modulus of aluminum is roughly one-third that of steel, springback of aluminum is generally more pronounced than that for steel. For this reason, it is important to predict the springback of aluminum panels so that changes to the tooling geometry may be made to correct the final panel shape.

LS-DYNA3D can be used to predict the panel springback by applying damping to the panel during dynamic relaxation. However, this CPU-intensive process requires hundreds of thousands of timesteps to complete. To address this issue, Livermore Software Technology Corporation modified the LS-DYNA3D code to generate an input deck to the LS-NIKE3D code for the springback analysis.[1]

LS-NIKE3D[6] is an implicit FE code. Implicit codes are typically less robust than explicit codes for handling contact, but they are ideally suited for springback calculations that do not have any contact conditions to satisfy. Springback calculations can be completed in a few timesteps using an implicit code, as compared to the several hundred thousand timesteps for an explicit code. This translates into significant CPU cost savings. For instance, using LS-DYNA3D and dynamic relaxation required 135 hours on the CRAY Y-MP system, whereas LS-NIKE3D solved the same problem in only 48 minutes.

The predicted contours of deflection on removal of the tooling—the parameters of the springback—are plotted on the drawn panel in Figure 4. The greatest deflections occur in the dogleg region where the dogleg twists, causing the wheel well side to move down by 8 mm and the door line side to move up by 6 mm. This twist is judged to be a result of the initial contact condition during binder wrap. To correct this, the tooling could be tipped so that the contact occurs evenly along the bottom of the dogleg.

While implicit methods use less CPU time than explicit codes for predicting springback, they require significantly greater memory. In this instance, the explicit code required 25 Mwords of memory on the CRAY Y-MP system, while the implicit code required 256 Mwords. This memory requirement underscores the advantages of large-memory systems, although in multiuser environments these advantages may be somewhat mitigated.

Visualization and interpretation

The figures in this article were generated using EnSight (formerly MPGS), a visualization package from Computational Engineering International.[7] This package allows the die designer to see a variety of values plotted on the deformed fender geometry. However, sheet-metal forming is a dynamic process in which the geometry changes rapidly. Individual snapshots of the forming process do not necessarily convey enough information about the total forming process. For this reason, EnSight was used to generate several hundred snapshots at equal time intervals. These snapshots were used to make a videotape animation of the forming process for the die designers. In this way, it is possible to determine cause and effect relationships. For example, the development of the wrinkle above the wheel well during binder wrap can be observed and linked back to the initial contact conditions.

Discussion

The effectiveness of using large-scale manufacturing simulation as an engineering tool was clearly demonstrated through this work. Within the rather broad subject of manufacturing and simulation, or more specifically sheet formability modeling, a large number of technical options are available to balance the opposing restraints of time, cost, and modeling accuracy. The guiding principle must be to match the accuracy of the modeling simulation to the expected decision process. The present example has demonstrated that very detailed models are technically feasible using existing hardware and software, and this approach can be used to critique the intended design and manufacturing assumptions prior to developing any physical tooling. Furthermore, simulations of this type can provide valuable insight into the root cause of many stamping defects and can greatly facilitate the concept of designing for manufacture.

The FE model presented has been extensively verified by comparing the model predictions with actual prototype fenders. The model quite successfully predicted the overall strain, wrinkling, and springback behavior. In some locations the thickness strains were only qualitatively correct, and this error was attributed to the assumption of isotropic material behavior ($R=1$), whereas the actual value is closer to an R value of 0.7. It is possible to model the material with anisotropic behavior, but this increases the CPU time. However, when thickness distribution is determined to be an important factor, these refined material models should be used.

Conclusions

The FE method can be used to simulate the forming of large and complex sheet-metal stampings for automotive applications. To predict strains accurately and the final shape after springback, large models that reflect the true shape of the forming tools, including draw beads, must be developed. This approach has been adopted here using the explicit code LS-DYNA3D running in a supercomputing environment. The FE simulation accurately predicts the strains within the formed fender panel as well as the size and location of wrinkles which were formed in the prototype fender. LS-NIKE3D successfully predicted the final shape of the simulated fender in under one hour, using the input deck generated at the completion of the LS-DYNA3D analysis.

With this technology, complemented by the use of computer graphics as an aid to technology transfer, sheet stamping operations can be evaluated before the tooling has even been cast. As a result, the manufacturing simulation of automotive sheet stampings may one day be a common means of developing tooling for the forming of automotive components.

Acknowledgments

The authors are grateful to the Ford Motor Company for supplying the CAD data for the fender tooling and agreeing to the publication of that data. Control Data Corporation provided training and a great deal of assistance in the mesh generation using ICEM CFD/CAE. Cray Research was an invaluable partner and supplied CPU time and assistance in postprocessing using EnSight. John O. Hallquist provided timely enhancements to both the LS-DYNA3D and LS-NIKE3D codes.

About the authors

Chris Galbraith joined Alcan's Kingston Research and Development Centre (KRDC) in 1989 after obtaining a Ph.D. degree in mechanical engineering from Queen's University at Kingston. Since then he has worked on modeling sheet-metal formability for automotive applications.

Mark Finn has been employed at KRDC as a researcher for 13 years in a variety of positions, including material characterization, computer network administration, and statistical analysis. He is currently developing FE modeling techniques for the forming of sheet materials for automotive applications. He graduated as a mechanical engineering technologist from St. Lawrence College in 1981.

References

1. Finn, M. J. et al., "Use of a Coupled Explicit-Implicit Solver for Calculating Spring Back in Automotive Body Panels," Numisheet '93, Tokyo, 1993.
2. Hallquist, J. O., "LS-DYNA3D User's Manual," January 1, 1993, Livermore Software Technology Corporation.
3. PDGS Design Guide, Computervision Corporation, UPD30300-007, January 1993.
4. ICEM CFD/CAE Version 3.1.2, Publication Number 60000202, Control Data Corporation.
5. Galbraith, P. C. et al., *Cray Channels*, Vol. 13, No. 4, pp. 10-14.
6. Hallquist, J. O., "LS-NIKE3D User's Manual," January 1993, Livermore Software Technology Corporation.
7. Computational Engineering International, Inc., "EnSight," 1994.

Structure-occupant integrated analysis of an automobile equipped with an airbag

Yang Wook Choi,
Hyundai Motor Company, YongIn Gun, Korea

Executive summary

Automobile companies routinely use crash simulation to analyze the structural integrity of their vehicles. Hyundai has used computerized crash analysis since 1989, most recently incorporating it into the design phase of vehicle development. Hyundai engineers have developed a structure-occupant integrated analysis on their CRAY Y-MP4E supercomputer; this simulation model can incorporate more than 50,000 finite elements and includes the automobile structure, an occupant, and an airbag. This extremely complex simulation of objects with diverse characteristics reliably matches experimental results.

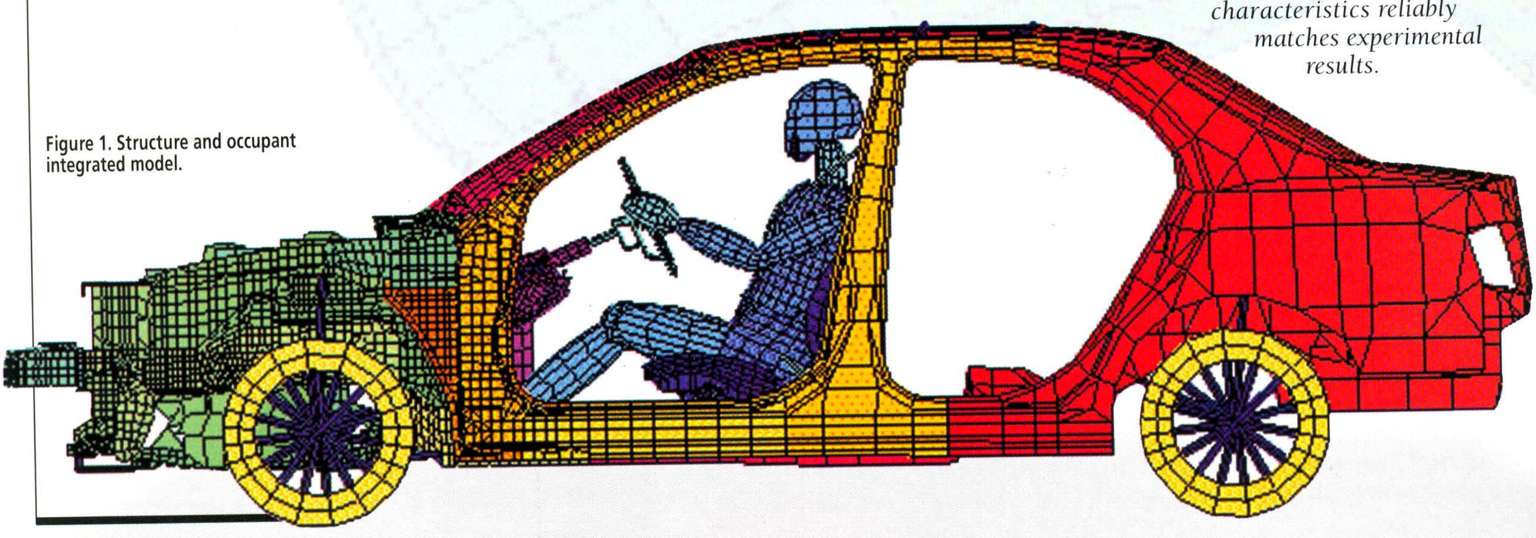

Figure 1. Structure and occupant integrated model.

The ultimate purpose of computerized automobile crash simulation and analysis is to estimate and decrease occupant injuries. Hyundai Motor Company engineers have used the PAM-CRASH structural analysis code since 1989 to analyze crash simulations of several existing car models and have seen a strong correlation between simulated and experimental results. Hyundai engineers now use structural crash analysis as a convenient tool for estimating crashworthiness when developing a new vehicle. They also have done occupant simulation with the MADYMO code developed by TNO. Combining these two simulations into a structure-occupant integrated analysis is the best way to determine a new vehicle's crashworthiness in the early stage of its design.

In the past, engineers have used a structure-occupant related analysis, in which the structure is analyzed before the occupant analysis takes place. The structure-occupant related analysis method is useful for simulating a frontal crash in which a passenger (or dummy) is restrained by a seat belt. Forecasting results is easier than with other types of crash models. However, crash models are becoming more complex, as is the simulated motion of the dummy. At this point, the related method is unable to estimate the various interactions found in different crash models. Furthermore, the occupant restraint system has been developed to include an airbag, which passively protects an occupant in a crash. With the airbag, the motion of a dummy differs from that of a frontal impact simulation that includes just a seat belt. The airbag system is so sophisticated that it must be tuned car by car. The related method cannot account for the airbag.

Each car company has developed its own method of integrating the structure and occupant analysis into a finite element model that includes all of the interior parts and the dummy. One Hyundai integrated model, shown in Figure 1, contains more than 50,000 elements and includes various kinds of materials, so performing analyses is time-consuming and can be expensive. Advances in computer hardware and software have made such analyses possible. Using Hyundai's CRAY Y-MP4E supercomputer, Hyundai engineers worked with Cray Research analysts on a joint integrated analysis project.

To obtain a reliable model, the integrated model was merged part by part after each subpart was analyzed and passed the experiment correlation test. The procedure included four phases and each phase had several tasks. These phases were:

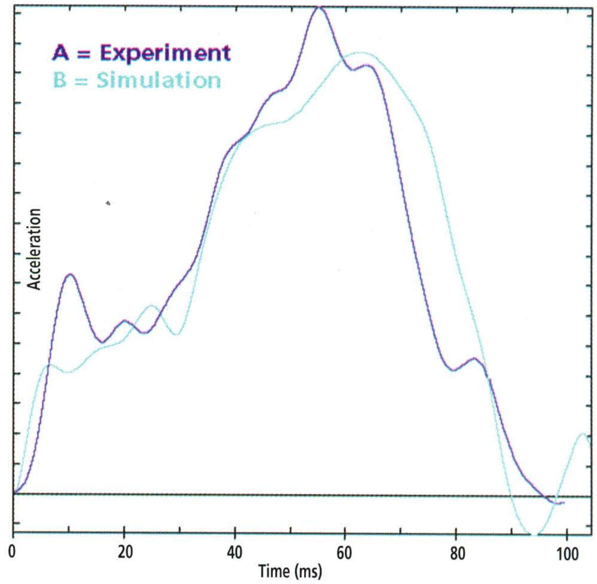

Figure 2. Comparison of acceleration results of structural model and results from experimentation.

- Structural model analysis
- Understanding the dummy model and analysis of interior parts
- Airbag analysis
- Model integration and analysis

Available experimental data and/or simulation results from MADYMO were used to compare results. The first phase was done at Hyundai, and the other phases were done at Cray Research's Eagan, Minnesota, computing center. To solve the model, Hyundai engineers and Cray Research analysts used both CRAY Y-MP and CRAY C90 systems.

Structural model analysis

Hyundai selected one model from several that were built before the project started. The model was validated with experiments, and only some interior parts were modified. Figure 2 shows the correlation of numerical results with experimental results. The structure was modeled in a typical way, with about 25,000 shells, beams, and rigid bodies, as shown in Figure 3. The model contained every important part in the engine compartment, including detail up to the dash panel (firewall). The element size behind the dash increases to the rear part of the vehicle.

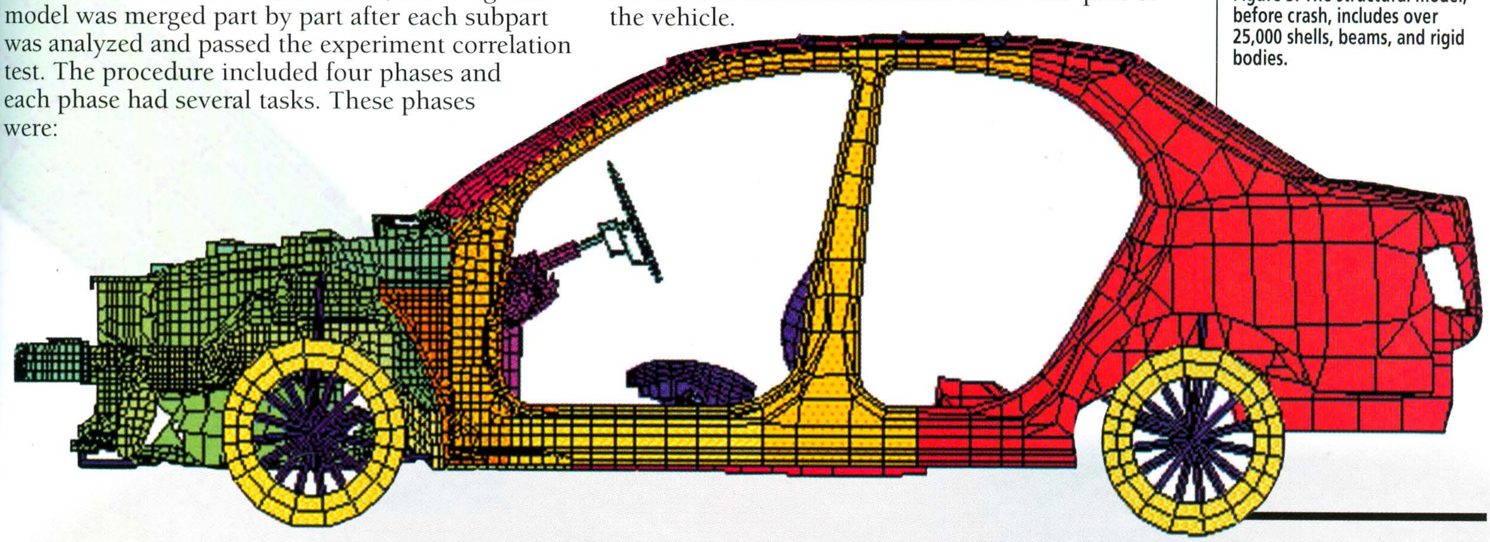

Figure 3. The structural model, before crash, includes over 25,000 shells, beams, and rigid bodies.

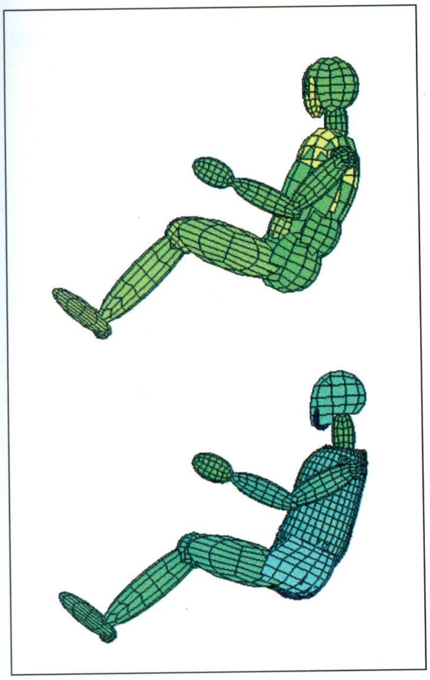

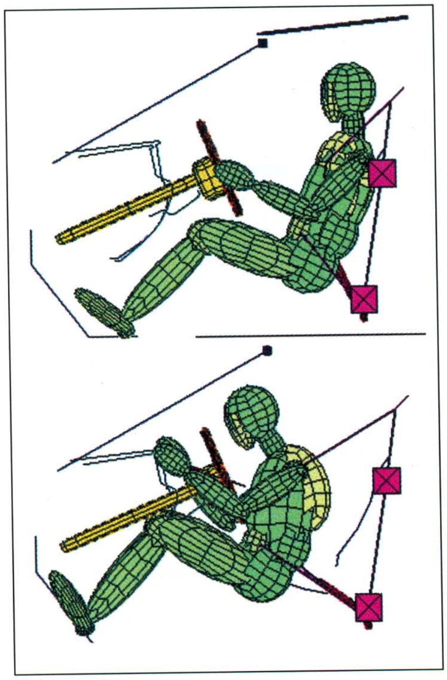

Figure 4 (above). Rigid (top) and finite element (bottom) dummy models.

Figure 5 (above right). Rigid sled model before crash (top) and after crash (bottom).

Understanding dummy model and analysis of interior parts

To help understand dummy motion quickly, MADYMO dummies input data was used, and David Lasry from Engineering Systems International (ESI) helped set up the dummy models. Two types of dummies were used—the rigid body dummy and finite element dummy. The schematic view of the two models is shown in Figure 4. The rigid body dummy model, the same as that of MADYMO, consists of lumped masses and contact ellipsoids and joints. The MADYMO input data was easily converted into PAM-SAFE format with special codes developed by ESI and Hyundai.

The rigid sled model, shown in Figure 5, was used to test the rigid body dummy model. The rigid sled model was composed of the rigid body dummy, plates for describing the contact of interior parts, and structural parts. The restraint system consisted of a seat belt modeled with special bar elements created for a seat belt. The model was moved by the acceleration field that resulted from structural crash simulation. The results from the rigid sled model using PAM-CRASH were almost the same as those from MADYMO, so the model was considered acceptable. However, the concept was the same as for the related analysis. To replace the contact plate with a 3-D finite element model, validation of the characteristics of each part had to be performed first.

Several impact tests of dummy parts crashing into interior parts of the automobile were simulated to validate the interior finite element model. The knee bolster and steering wheel are the most important parts in a frontal crash. The knee bolster is composed of a cowl cross member, a contact plate (a real, not arbitrary, plate to describe contact during a crash), and brackets. The knee bolster model before and after analysis is shown in Figure 6. Creating a model of a plastic cover that contains steel parts of the contact plate and brackets is not easy. Plastic has different material characteristics; therefore it is hard to define contact between those different materials.

The simulation was done without the plastic cover, and the results revealed lower resisting force than that shown by experiment. The two curves are shown in one graph in Figure 7. The tendency of the resisting force was the same, so the model was thought to be good enough to use in the integrated model.

The steering wheel system is more complex than that of the knee bolster. The steering system is designed to absorb crush from the engine compart-

Figure 6. Knee bolster model before crash (top) and after crash (bottom).

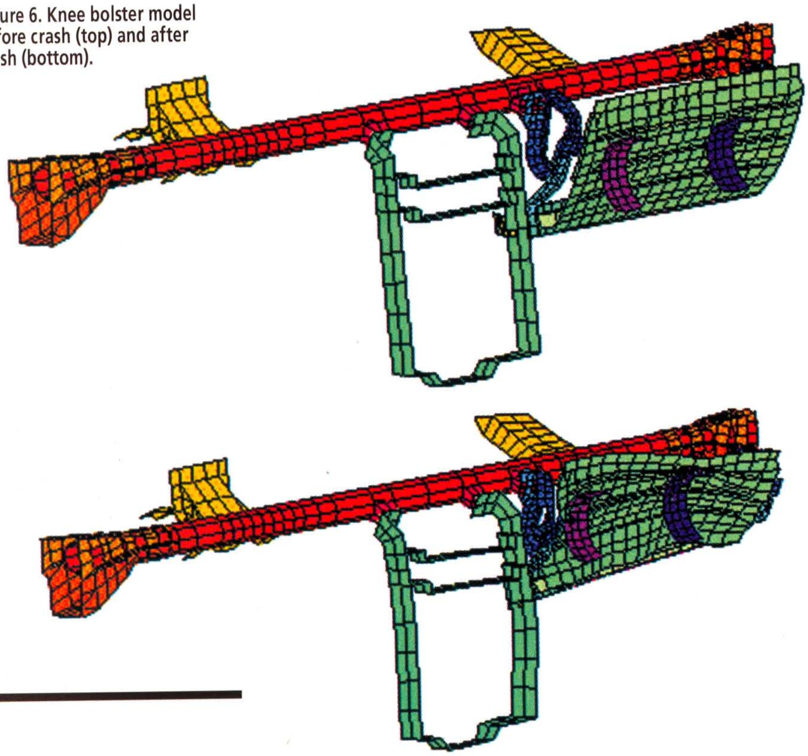

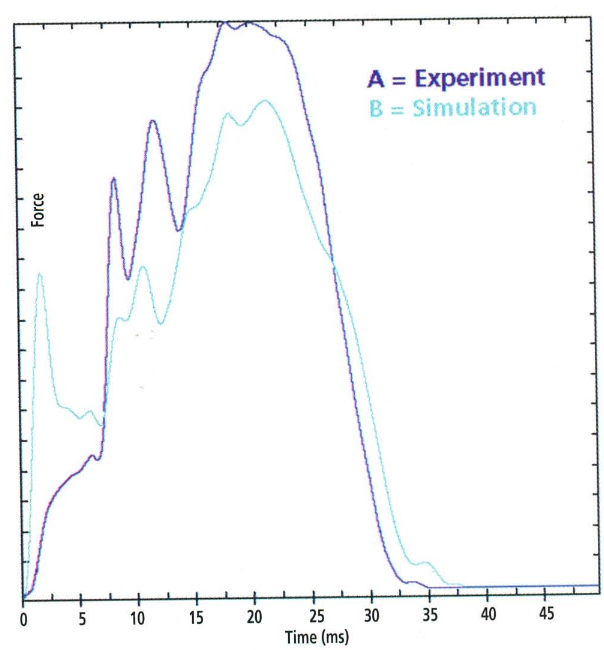

Figure 7. Comparison of simulation and experimental knee bolster crash results.

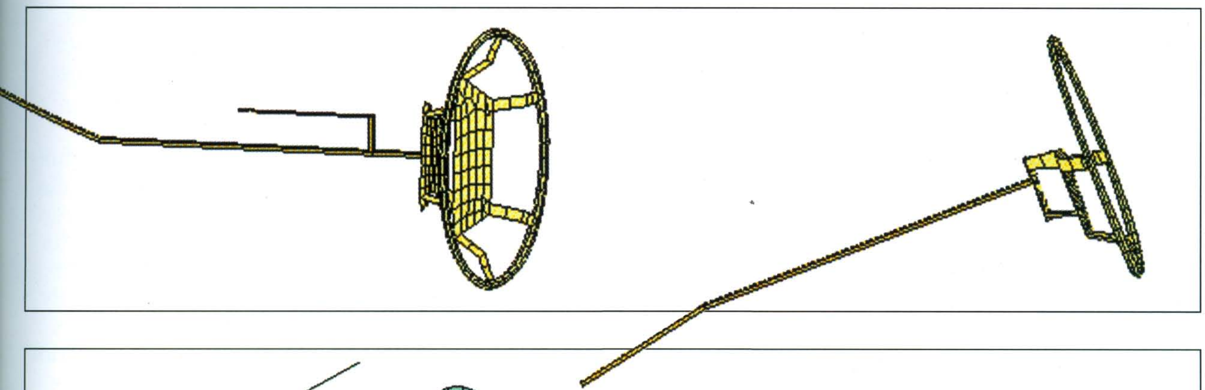

Figure 8. Steering system model, with the energy absorbing characteristics modeled by nonlinear spring elements.

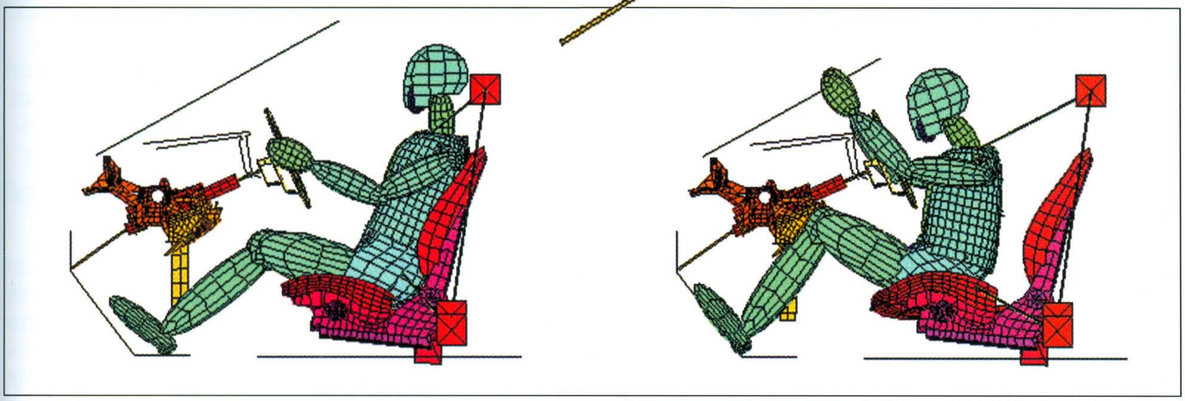

Figure 9. Finite element sled model before crash (far left) and after crash (left).

ment and the crash energy of a driver hitting the steering system. The crush-absorbing system was modeled using joints, and the energy-absorbing system was modeled using nonlinear spring elements that represent the experimental characteristics. The model is shown in Figure 8.

The seat was modeled with solid elements, and the key was to avoid the hourglass mode. The belt system was modeled using PAM-CRASH's belt, slipring, and retractor elements. The retractor caused some problems, and because the slipring element in PAM-CRASH transfers belts element by element, a little oscillation occurred in the sliprings.

After the parts tests were finished, all subparts were merged into a sled model. Output from the finite element sled model approached test results, and it was thought that the contact between the dummy and interior parts was more realistic than with a rigid element dummy. The finite element model could describe deformation in any direction, so the contact force was different from the parts test. The parts test was set in only one direction and did not permit sliding in any other direction, but the finite element model was allowed to move in any direction, and the contact force was generated even though the direction was different from the desired direction. The behavior of dummy and interior parts is shown in Figure 9.

Airbag analysis

The airbag restraint system uses crash-detecting sensors, usually accelerometers, mounted in several locations. The sensor signals are measured and analyzed to define whether a car's motion is that of a crash. Once the controller recognizes the signals as crash signals, the inflator is activated. After activation, the inflator explosive pours gas into the airbag, and the airbag is inflated. When an occupant hits the airbag, the gas leaks out through an orifice and small holes in the fabric to absorb the occupant's kinetic energy. This whole process is completed in only 0.08 seconds; the inflating process of the airbag should be completed in 0.04 seconds. Because each vehicle has its own crash characteristics, tuning an airbag is very important. Several basic parameters such as airbag volume, outlet hole size, inflator temperature, and pressure change must be adjusted. To find the optimum parameters efficiently, the airbag must be analyzed.

PAM-CRASH has its own airbag model—a membrane element that can control gas thermodynamics. The airbag model is separated into two models, the flat airbag and the folded airbag. The flat airbag is like a flat plate that has no gas in it—no folding has been done in this model. The flat airbag is calculated first to tune the airbag. The shape of the flat airbag model before and after inflation is shown in Figure 10. After completing the flat airbag model, it is folded into several layers. Folding the airbag is so complex that the folding line must have nodes on it, and the gap between the layers must have some distance to avoid initial penetration. ESI provided a special program, PAM-FOLDER, to fold an airbag interactively. The gap is usually smaller than the element size, so a penetration problem always exists. Therefore, the following limits must be used in an airbag analysis:

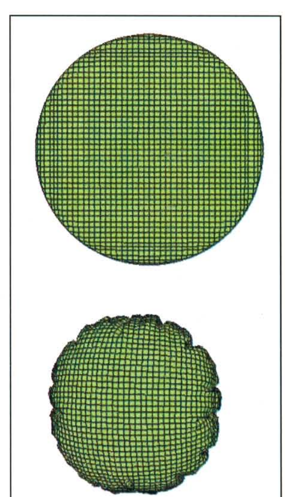

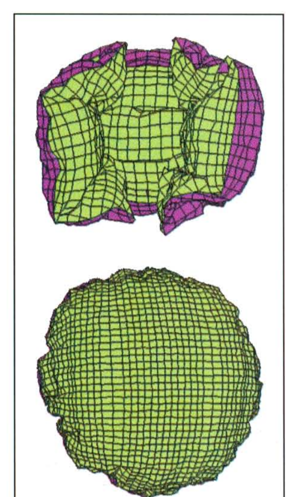

Figure 10 (below left). Flat airbag model at 0 ms (top) and 50 ms (bottom).

Figure 11 (below). Folded airbag model at 10 ms (top) and 50 ms (bottom).

- Set contact search time interval to 1 cycle
- Set friction coefficient to 0.0 (no friction)
- Set contact thickness under 0.5 mm

The flat and folded airbag simulations are similar except for mass flow. The flat airbag can be used in an integrated model; the crash mode is restricted to frontal crash. The folded airbag model is useful to estimate crashworthiness of an unrestrained crash, a crash involving children, and several out-of-position crashes. The airbag model has more than 3,000 membrane elements and no strap inside. Self contact is important. Figure 11 shows the folded airbag model.

The airbag has top and bottom sheets. Each sheet has top, bottom, and parent layers. The fabrics of top and bottom sheets cross 45° and 90°. Generally, the driver side airbag is coated to prevent gas leakage through the fabrics.

Model integration and analysis

After the airbag analysis was successfully completed with the finite element method—made possible with new features in PAM-CRASH v12.1—the integrated analysis for crashworthiness became possible. Before incorporating the airbag into a full vehicle model, the sled model was used first to test the airbag with other interior parts and then with an occupant. Several modifications were made to the sled model equipped with an airbag. For example, penetration between the airbag and steering wheel caused some problems, which were solved by increasing the element size of the steering wheel. Making the steering wheel rim a little bigger equalized the element size with that of the airbag. Severe problems also occurred with the dummy's impact on the airbag. The airbag material was so soft that the hard dummy materials could push back the airbag when contact was detected. This was due to a penalty factor in the contact algorithm. Several trial and error tests were conducted to find the proper penalty factor value for this kind of contact. The dummy's head and chest are the points of contact, and those two parts have different material properties. The head is stiffer than the chest, and even the chest has a sternum which makes itself soft. Therefore, trial and error should be done for both cases.

After all of the contact parameters were tuned, the airbag worked well in the sled model. Encouraged by our series of successes, we built a fully integrated model by merging the structure and occupant models and adding the tuned airbag. The analysis was easier than the several phases that came before. Figure 12 shows the fully integrated model and the results.

The only problem was CPU time. The fully integrated model contained more than 50,000 elements, and consumed over 30 wall clock hours, even on a CRAY C90 system. However, the results were very realistic. As usual, numeric simulation and visualization show useful data that can hardly be seen in real tests. Once the integrated analysis was done, interactive behavior between the interior and the dummy, and structural parts such as the dash, toe pan, floor, and dummy, and restraint system and dummy can be optimized. The most important behavior is between the restraint system and structure. Interference of structure and occupant was demonstrated first in this fully integrated model.

This project and the results helped us overcome many new problems revealed by initial results. To adjust the reliability of the model, we modified the crash characteristics of plastic materials of interior parts. In addition to necessary algorithm improvement for compressive foam solid, the algorithm also has to account for belt material (especially in the retractor); the contact algorithm or contact force monitoring features of an airbag must be taken into account, too.

Conclusion

Hyundai now has a new method of structure-occupant integrated analysis, developed with the help of Cray Research analysts and automotive industry consultant Paul Du Bois. While the model is quite good, many things can be improved. After we used the model for a series of analyses on non-steel materials, we know that the model reveals a realistic correlation and can now be used on other types of crash models.

About the author

Yang Wook Choi is a section chief engineer in the research and development department of the Hyundai Motor Company Research Center, where he has worked on crashworthiness since 1985. He received a B.S. degree in mechanical engineering from Korea University in 1985 and an M.S. degree in engineering mechanics from The Ohio State University in 1991.

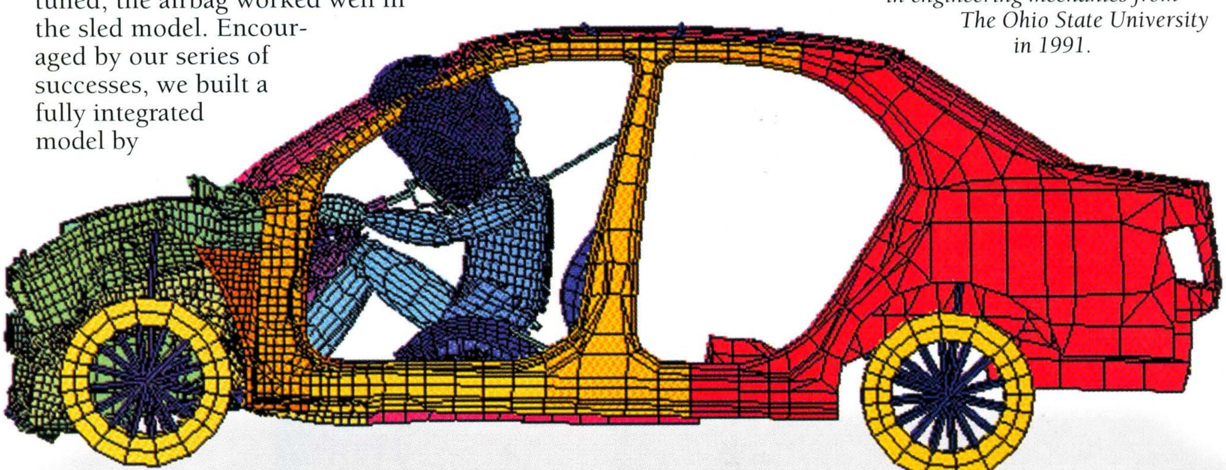

Figure 12. Fully integrated model after crash.

OptiStruct

A topology optimization tool for conceptual design

Jeff Brennan,
Altair Computing, Inc.,
Troy, Michigan

Executive summary

Classical design optimization methodology has begun its migration from the classroom to use by industrial designers. OptiStruct is the first commercial software package to use a new topology optimization technique that creates arbitrarily shaped designs with minimal input. Two examples from the automotive industry (decklid design and control arm design) are presented to demonstrate the viability of this new simulation technology. The examples also demonstrate speedups and increased problem resolution when using OptiStruct on a CRAY EL system.

The design of strong, lightweight structures is a common and important engineering goal in the automotive industry. Increased global competition has put an additional burden on engineers to not only produce efficient structures, but to do so quickly, prompting engineering groups to rely more on computer-aided engineering (CAE) tools for analyzing potential designs. The success of finite-element analysis (FEA) and the increased use of optimization tools (coupled with FEA) are evidence that computational methods can significantly affect structural design, yet a clear need exists for more powerful computational analysis and optimization tools. Often thought useful in academia only, classical optimization methods are beginning to proliferate in industrial design. Several of the major FEA code vendors are developing powerful optimization modules, and some new codes, such as GENESIS by VMA Engineering, have been introduced specifically to apply classical optimization methods to modern engineering problems.

The flexibility of optimizing designs with many load cases and constraints for diverse objectives, and the ability to solve large-scale models, have opened the door for optimization methods in all types of engineering design. One difficulty with traditional size/shape optimization algorithms is that to pose the problem, one must begin with an initial design layout and define, in advance, potential changes to the size/shape of this design. In other words, the topology (or layout) of the design is fixed by the user, and major structural features such as the connectivity, number of members, and the placement of "holes" in the structure cannot be changed. However, a much more efficient design may exist—one with a completely different topology. Only by first solving the layout optimization problem, and then obtaining a conceptual design with an optimal topology, can the true potential of size/shape optimization be realized during design refinement.

The difficulty of solving the layout optimization problem has long been prohibitive due to the ambiguity in posing the problem mathematically, as well as the computational expense. Fortunately, modern supercomputers now can solve realistic engineering optimization problems with many design

Figure 1 (below). Design of a reinforcement for a decklid outer panel.

Figure 2 (bottom). OptiStruct-predicted topology for decklid reinforcement placement.

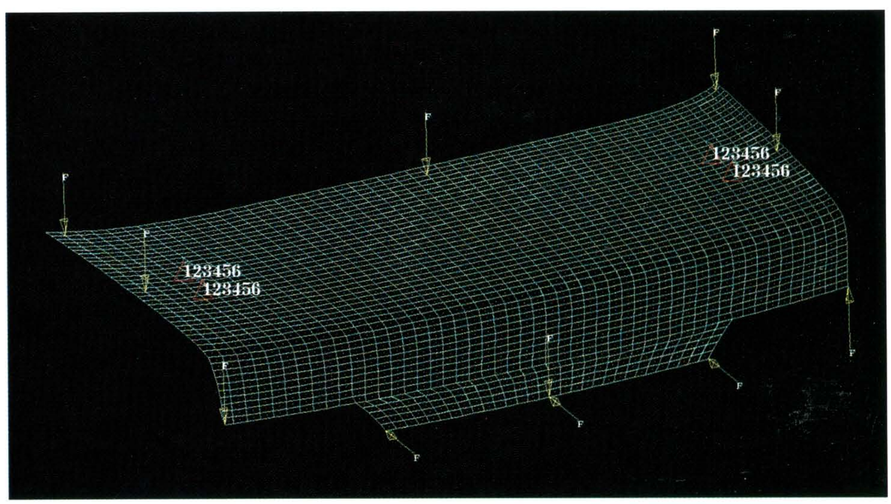

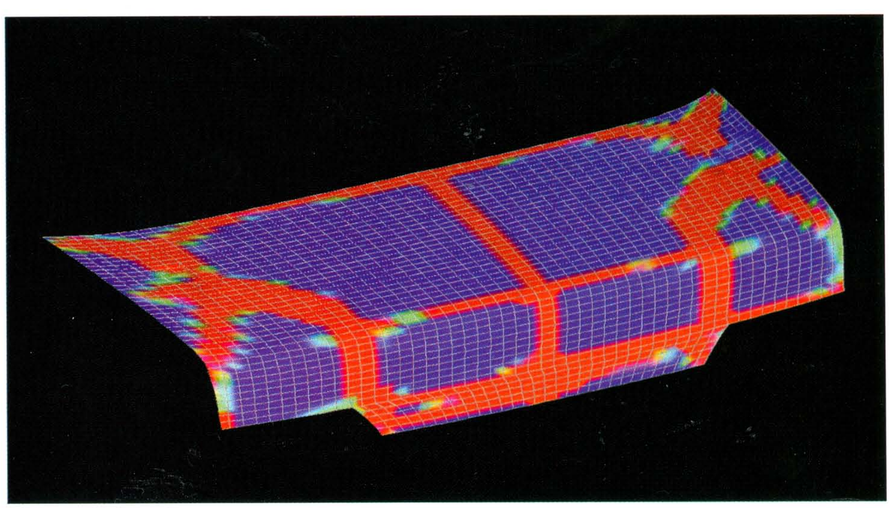

variables and involving multiple solutions of large finite element models. What is needed is a robust methodology to solve the generalized layout problem. Recently, a topology optimization technique has been developed that can create optimal designs of arbitrary shape with minimal input. In 1993, Altair Computing, Inc. introduced OptiStruct, the first commercial software package using this technique.

How OptiStruct works

An elegant mathematical algorithm that requires relatively simple input data resides within OptiStruct. An engineer needs only to define a finite element model of the potential package space, the necessary constraints and load cases, and an estimate of the available material. OptiStruct creates a robust structure optimized for multiple loading conditions when an accurate description of the service load conditions is used. OptiStruct also contains its own finite element solver, and the user is free to define a complex finite element model or assembly and to optimize only a portion of the model or assembly layout. OptiStruct contains three solution schemes: a 2-D plane stress elasticity solution, a 3-D shell solution, and a 3-D solid elasticity solution. Currently, the user is free to mix elements of these three types outside the design domain, but all elements within the design space must be of a single type.

The OptiStruct algorithm assumes the design space to be a porous medium consisting of a number of small voids. Taking the void sizes and orientations as design variables in the optimization routine, OptiStruct alters the material distribution throughout the design space to minimize the structural compliance. In other words, the void size will shrink where local compliance is the greatest and expand in areas of little or no compliance. Although the void size is adjusted with respect to local compliance, the void orientation angle is adjusted with respect to the principal stress directions, an important feature in the optimization process.

Homogenization allows OptiStruct to relate the inhomogenous problem of void/solid mixtures to the practical problem of a homogenous material distribution. Homogenization is a mathematical technique that converts the microscopic material behavior of a periodic, cellular composite material into an equivalent set of macroscopic properties for an anisotropic, homogenous material. This technique lets OptiStruct search in the anisotropic space of solutions for the optimal layout, while the code actually attempts to converge to an isotropic solution (that is, the algorithm favors densities close to 0 or 1). By expanding the space of solutions, OptiStruct tends not to converge to local minima and provides a robust optimal topology. OptiStruct uses homogenization on an element basis to calculate the effective properties of each element after each iterative update of void size and orientation. The finite element equations then are used to solve for the displacements of this new material distribution, and the process iterates until the solution converges. The final material distribution, or topology, can then be postprocessed to visualize the shape of the structure. This shape becomes the optimal starting point for design.

Applications in the automotive industry

While OptiStruct is a robust tool for design in any industry, the majority of applications to date have been automotive related. Successful uses of OptiStruct in the automotive industry include design of sheet metal reinforcements, rib placement on castings, suspension member design, and bracket design. The following examples show how the Cray Research supercomputer version of OptiStruct is used in automotive design.

Shell solution example: decklid design

The design of reinforcement patterns for sheet-metal parts, such as automobile body panels, is an excellent example of the use of topology optimization. Traditionally, the pattern for reinforcements such as a decklid, or trunk, inner panel is determined by experience, design intuition, or both. This pattern, perhaps in the shape of the letter "W", is then refined into open channels for manufacturing (stamping), and the width and depth of these channels and the gauge of the material are then the variables during the design optimization process. The problem is that a better configuration, say a "double-X" pattern, might lead to a more efficient design. Consider the design of a reinforcement for the decklid outer panel shape shown in Figure 1. In this example, each of the ten loads shown is considered a separate load case (except for the torsion couple load at the corners), and

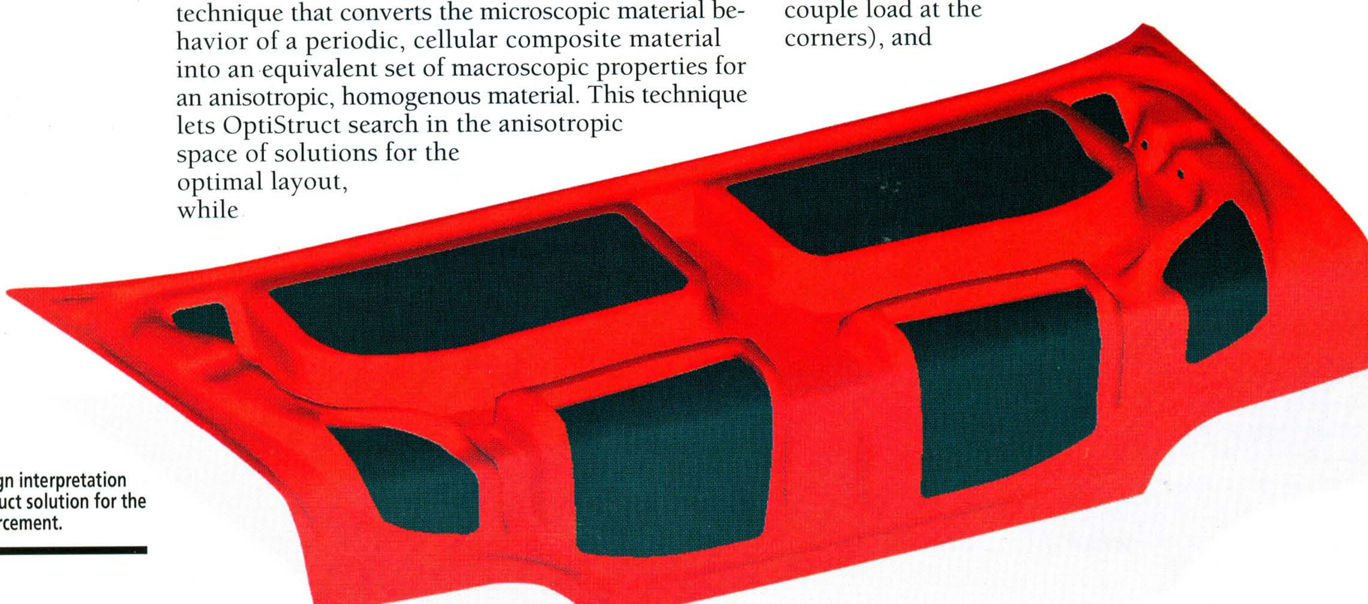

Figure 3. Design interpretation of the OptiStruct solution for the decklid reinforcement.

OptiStruct is then asked to produce an optimal shape of a reinforcement pattern of maximum stiffness for a specified target mass. The shell solution within OptiStruct allows the definition of a "core" thickness and a separate "reinforcement" thickness for each element. In this example, the 1.0 mm outer panel contributes to the stiffness of each element and is not changed, and the design (package) space is then the user-defined depth of the inner panel across the surface area of the decklid. Figure 2 shows the OptiStruct-predicted optimal topology. Of course, this is simply a design concept that shows where reinforcements need to be placed, and the user must convert this idea into a workable design model using "U"-type channels where the ribs were placed and introducing a seal around the perimeter.

Figure 3 shows the design interpretation of the OptiStruct solution. This design then serves as the optimal starting point for the design process. Further analysis and size/shape optimization should follow. However, because this design began with the correct topology, these further improvements to the base design should yield the most efficient design possible for the requested specifications. OptiStruct is a compute-intensive code; in this model, 30 iterations of the finite element solution were required to achieve a refined optimal shape. On a high-end workstation, this would amount to nearly 15 hours of CPU time. However, using a single processor on the four-processor CRAY EL system at Altair Engineering, this model took less than 3 hours of CPU time. This five-fold increase in speed will further increase when Altair upgrades its CRAY EL system to a CRAY J916 system.

Solid solution example: lower control arm design

This example is typical of layout design problems in the automotive industry. The goal is to produce a lightweight, efficient topology for the design of a cast aluminum lower control arm. The finite element model of the package space is shown in Figure 4. In this example, the frame connections, the shock mount position, and the spindle attachment (shown in red) are fixed and not part of the optimization process. The potential package space is shown in blue. The three principal load cases applied to the model are from braking, cornering, and vertical load service conditions experienced by a typical automobile suspension. The OptiStruct-predicted topology is shown in Figure 5. The distinct topological features are the two "V"-shaped members that connect the shock mount to the frame and the thin cross-member connecting the bushings at the frame (longitudinal travel was allowed between these bushings). This view does not show the sculpting of material under the shock mount. Also, because of the significant bending in the shock-to-spindle connection area for all three load cases, nearly all elements in this area converged to full density, signifying the importance of material placement there.

The design interpretation of this OptiStruct solution can be seen in Figure 6. This concept model can then be used to assess the baseline design and provide an optimal topology to classical shape/size optimization tools. This particular model has already been analyzed using CSA NASTRAN, and with one modification of member thickness, it produced an extremely lightweight design that met all stress targets. Of course, this is a relatively simple solution to a simple, symmetric design problem, and more complex models tend to yield less intuitive solutions. In fact, much larger and more complex models probably will be analyzed using OptiStruct; the Cray Research version of OptiStruct provides a significant advantage. Increased mesh density within OptiStruct means an improved image of the predicted topology, and perhaps even preliminary size and thickness information can be obtained from fine models. The 1 Gbyte of memory on Altair's CRAY EL system allows the solution of this control arm problem with 8 times the resolution, providing much more clarity and resolution in design interpretation. In another example, an OptiStruct customer in the automotive industry performed an analysis that took 2 minutes per iteration on a CRAY C90 system. This same 27,000 degree-of-freedom analysis took about 16 minutes per iteration on the CRAY EL system and 90 minutes per iteration on a high-end workstation. As the desire for larger models within OptiStruct and other analysis codes grows, the Cray Research supercomputer will continue to be the workhorse at Altair.

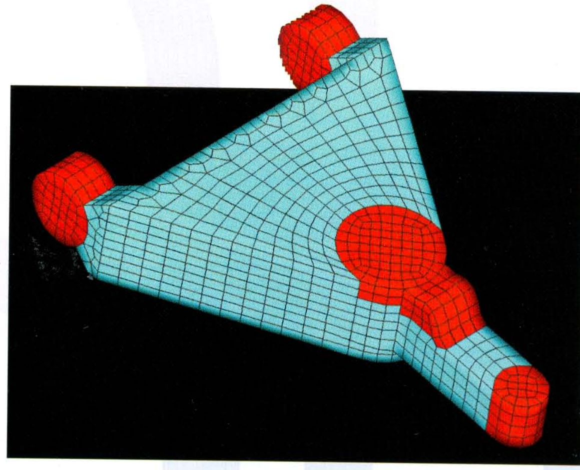

Figure 4 (left). Finite element model of a cast aluminum lower control arm. Fixed design elements are shown in red; blue designates the package space that could be optimized.

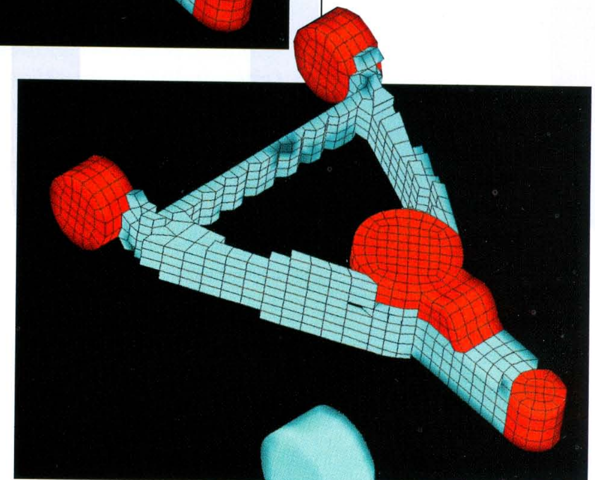

Figure 5 (below). OptiStruct-predicted topology for a control arm.

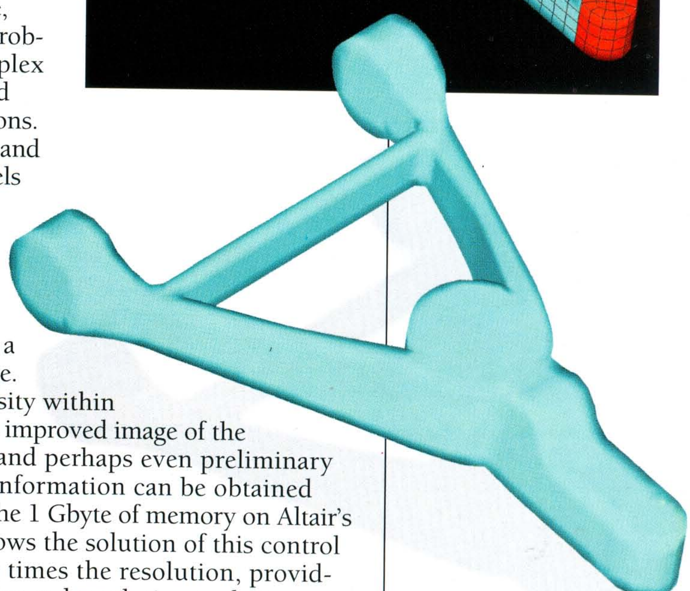

Figure 6 (bottom). Design interpretation of the OptiStruct solution for a control arm.

About the author

Jeff Brennan is the OptiStruct product manager for Altair Computing, Inc. He holds a B.S. degree in mechanical engineering from the University of Notre Dame and an M.S. degree in mechanical engineering from the University of Michigan.

HEXAR 1.0

Breaking the mesh generation barrier

Computer-aided engineering (CAE) simulation applications are used throughout automotive manufacturing companies. These applications result in better designs, lower manufacturing costs, and safer products. Frequently, however, a hidden barrier exists to using computational fluid dynamics (CFD) and finite element analysis (FEA) simulations of complex models. Handcrafting the meshes that describe model geometries is painfully slow. Thus, the time required to generate these meshes has made many simulations impractical. Until now.

Cray Research has recently introduced HEXAR 1.0, a software package that automatically generates complex 3-D unstructured hexahedral meshes on any Cray Research parallel-vector computer system.

"HEXAR software is a genuine breakthrough," said Reza Taghavi, head of the Cray Research HEXAR product development team. "Until now, engineers typically have had to take initial CAD data and painstakingly handcraft 3-D meshes using workstation-based software before beginning computer analysis of new product or component designs. It could easily take six months and $150,000 or more to develop a million-element, hexahedral (six-sided "brick") mesh. Now this typically takes less than 30 minutes using HEXAR software on a Cray Research supercomputer system."

Volume mesh for the fluid and thermal analysis of the engine compartment in a Ford Taurus.

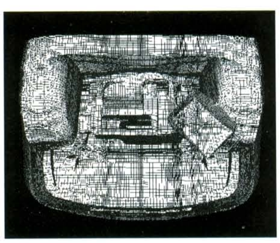

Automatic mesh generation is here

HEXAR, designed, developed, and supported by the Cray Research Engineering Applications Group, is an automatic volume grid generation software package that works directly on raw computer-aided design (CAD) surface data. A solid modeling step is unnecessary. HEXAR produces unstructured, boundary-fitted meshes containing only hexahedral elements. This element type is an optimal choice for numerically accurate simulations. HEXAR produces meshes that are suitable for most real-world CFD and FEA applications.

HEXAR turns months into minutes

Common CAE projects generally involve three basic steps: computer-aided design, hand generation of a 3-D mesh, and computational analysis. A nine-month, large-scale CAE project might require six months for mesh generation alone. HEXAR can reduce those six months to minutes. A low-cost Cray Research computer system and HEXAR can provide mesh generation services typically requiring a large staff of consulting engineers—for a lot less money.

HEXAR accepts CAD surface definitions, even relatively poor quality ones with small gaps and slight overlaps. It uses surface feature recognition, local mesh adaptation, and internal volume scanning techniques to create unstructured volume meshes that describe your model with the accuracy required for CFD and FEA simulations. HEXAR takes your CAD files and quickly returns a mesh that you can immediately submit to an analysis application.

HEXAR enables engineers to explore many design options instead of only one or two. Performing more iterations on a design can greatly improve the quality of the final product or component. Now you have time to consider factors that can reduce manufacturing costs. You can ask more "What if...?" questions than ever before.

Make changes on CAD drawings, not in the mesh

Simulation results often dictate design changes that must be reflected on the original CAD documentation. The usual practice has been to alter the computational mesh during the analysis phase as necessary, then update the CAD drawings later. HEXAR allows you to specify design changes on the CAD data and quickly create a new mesh. When you generate meshes directly from the CAD design files, your CAD documents can be up to date when the final analysis is complete.

HEXAR takes mesh generation one step closer to push-button, black-box, automatic operation. It reads CAD surfaces from all major CAD applications that produce the industry standard IGES and SLA (stereo lithography) formats, or our generic triangulated surface patch format. HEXAR produces mesh files in PATRAN-neutral, I-DEAS universal, NASTRAN bulk, or Cray Research generic engineering format. For more information, contact Lee Roll at 612/683-3552 or lee.roll@cray.com.

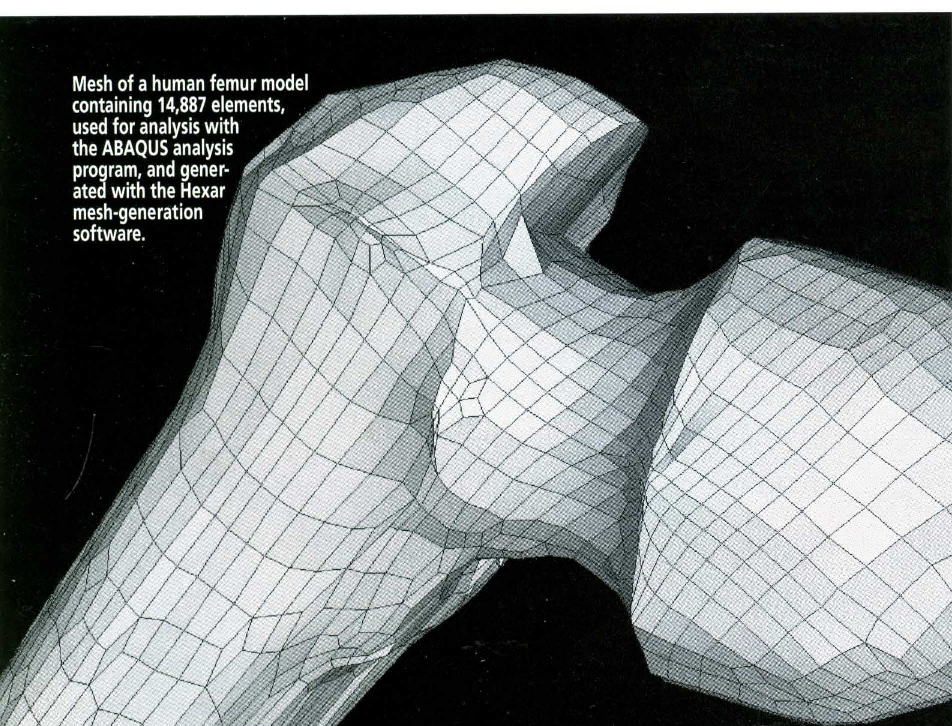

Mesh of a human femur model containing 14,887 elements, used for analysis with the ABAQUS analysis program, and generated with the Hexar mesh-generation software.

CORPORATE REGISTER

Hawtal-Whiting Group installs CRAY EL98 system, will add CRAY J916 system

Hawtal-Whiting Group, an engineering design consulting firm based in the United Kingdom, has installed a CRAY EL98 supercomputer at its Basildon, Essex, headquarters and has also ordered a CRAY J916 system. Hawtal-Whiting's clients include most of the world's automobile manufacturers. In addition to offering styling, design, and manufacturing services, the company also specializes in structural, dynamic, acoustic, and crashworthiness analysis of vehicle bodies.

Hawtal-Whiting will use the CRAY EL98 system primarily to run structural analysis applications. The CRAY EL98 system will enable the company to analyze substantially larger models in significantly less time. Colin Cox, senior principal engineer, said Hawtal-Whiting selected the CRAY EL98 system after evaluating it against other supercomputers as well as networked workstation offerings. He also said that the system strengthens their credibility with customers and prospects in the automotive industry who are often already Cray Research customers.

Chiba University installs nine Cray Research systems for general education

Chiba University of Tokyo has installed eight CRAY EL92 deskside supercomputing systems and one CRAY SUPERSERVER 6400 (CS6400)—all to be used for general education. This is the first time anywhere in the world that a university has dedicated a Cray Research system to this activity.

The installed Cray Research systems will be used for computational science education as part of Chiba University's required courses. The systems will be used not only by engineering students but by law, economics, medical, humanities, education, physical sciences, nursing, and horticulture students as well. The Cray Research systems will be connected to more than 200 computer terminals located throughout the university campus.

The National Center for Atmospheric Research (NCAR) has installed four Cray Research systems, including a 64-PE CRAY T3D system, a CRAY Y-MP8I system, and two CRAY EL92 deskside systems.

The CRAY T3D and CRAY Y-MP8I systems will provide the foundation for the establishment of the Climate Simulation Laboratory (CSL) at NCAR. The aim of the CSL will be to partner with other organizations in research initiatives of mutual interest in the areas of global and regional climate change.

The **Manchester Computing Centre,** University of Manchester, added to its history of leading-edge computing by installing a 12-processor CRAY SUPERSERVER 6400 (CS6400) system as the new computing platform for the United Kingdom (UK) census data and other National Dataset (database) Services the Centre provides under government contract to academic researchers throughout the UK. The new system will support more than 90 higher education institutions and 2000 academic researchers registered for the National Dataset Services provided through the Centre.

The **University of Tokyo** has ordered and installed a CRAY SUPERSERVER 6400 (CS6400) system with 16 processors and one Gbyte of central memory to be used for research in machine design and control, computational fluid dynamics, and molecular dynamics. Bid criteria included the ability to run unmodified CAD (computer-aided design) software developed on Sun systems, the ability to meet future parallel processing needs, and high-speed scalar processing ability.

SERC obtains fastest supercomputer in Europe

The Science and Engineering Research Council (SERC) has chosen a 256-PE CRAY T3D system for use in Grand Challenge scientific applications. The system offers over five times the power of any existing UK academic facility. It will provide world-class high-performance computing facilities to UK researchers for use in oceanographic and atmospheric science, studies of the structures of new materials, realistic simulation of complex structural and fluid flow problems in engineering applications, modeling of large molecules, and simulation of the fundamental constituents of matter. Edinburgh Parallel Computing Centre (EPCC) at the University of Edinburgh has been chosen to provide a national service based on the new system. Users will interact with the system through a single-processor CRAY Y-MP4E system with 212 Gbytes of disk storage. Researchers throughout the UK will be able to access the system via the high-speed SuperJANET network.

The **United Kingdom Meteorological Office** has obtained a CRAY C916 system for its Bracknell, Berkshire location. With 140 Gbytes of disk storage and 256 Mwords of central memory, the fully loaded, 16-CPU system will be six times faster than each of the two CRAY Y-MP systems it replaces. The new supercomputer will be used for climate prediction research and to improve operational weather forecasting.

Germany's primary climate research support facility, the **Deutsches Klimarechenzentrum GmbH (DKRZ),** has installed a 12-processor CRAY C916 system at its Hamburg facilities. With 256

Mwords of memory, the new system will be one of the most powerful supercomputer systems in the world used exclusively for climate research.

The CRAY C916 system replaces a CRAY-2 system and a CRAY Y-MP4E system currently in use at DKRZ. Funded by the German Federal Ministry for Research and Technology, the new Cray Research system will be used for climatological research, primarily global modeling of environmental phenomena such as the predicted global warming due to greenhouse gas emissions.

The **Defense Nuclear Agency of the Department of Defense** has installed a 2-Gword CRAY M98 system at Los Alamos National Laboratory, supplementing their CRAY Y-MP and CRAY X-MP systems. The CRAY M98 system provides the primary computing resource for the more than 500 Defense Nuclear Agency users. The CRAY M98 system will permit the agency to consolidate workloads, address larger and more complex problems, and speed up throughput.

UNICOS 8.0 earns Trusted Network security rating

Cray Research is now shipping the UNICOS 8.0 operating system, which includes enhanced scalable parallel processing, security and availability enhancements, and, for the first time, centralized resource management facilities.

With its security enhancements, Cray Research has become the sole supercomputer vendor and only the second major computer vendor to receive a U.S. Department of Defense B1 "Trusted Network" security rating, which means a Cray Research supercomputer can operate in a heterogeneous networked environment and still maintain its evaluated rating.

The particular configuration of UNICOS 8.0 software called Trusted UNICOS contains security enhancements that go far beyond standard UNIX features to effectively protect mission-critical data against accidental or malicious access or corruption. Trusted UNICOS was designed to meet the B1 MDIA requirements of the Trusted Network Interpretations ("red book"), provides a fully functional, practical, and usable trusted computing environment, and is now included in the Evaluated Products List (EPL).

UNICOS 8.0, the eighth release of the product first made available in 1984, is bundled with Cray Research systems. (Trusted UNICOS is available on CRAY Y-MP with an IOS model E, CRAY C90, CRAY M90, and CRAY EL systems.) In addition to enhanced features for batch, tape, and volume tape management, as well as scheduling and accounting, which have long been available in UNICOS, the new version has kernel multithreading capabilities for improved scalable parallel processing.

CraySoft's NQE software upgrades to version 1.1

The Network Queuing Environment (NQE) software 1.1, a set of sophisticated batch processing and automatic load balancing software for UNIX computer networks, is available for several platforms, including the CRAY SUPERSERVER 6400 (CS6400) and Sun Microsystems family of products. NQE 1.1 software will also support other SPARC/SunOS systems, IBM RS6000 systems, SGI systems, HP PA-RISC systems, and DEC Alpha systems running OSF/1.

NQE software automatically distributes a job to the most appropriate network resource and provides reliable data transfer, so users can share enterprise-wide network resources more effectively.

A customer does not need a Cray Research system to use CraySoft software. CraySoft's NQE software is a client/server product that provides reliable server batch processing, and it is based on the industry-proven Cray Research NQS-based software developed and continually enhanced for the company's supercomputing systems.

CraySoft, formed in October 1993, brings Cray Research software—networking and application software, compilers, tools, and libraries—to more users on a variety of computer platforms, including workstations, servers, and PCs.

Cray Research spins off Computational Engineering International

Cray Research, Inc. has spun off its popular MPGS engineering postprocessing software package. A new, independent company, Computational Engineering International, Inc. (CEI), has been formed by former Cray Research employees to develop, support, and expand the market for this technology.

The software, marketed under the new name EnSight, is a distributed engineering postprocessing package used around the world by engineers and scientists to display and manipulate the results of large computational analyses.

The latest version of the software does not require users to have access to a supercomputer—therefore, the potential market for EnSight now includes nearly all engineering workstation users and users of Cray Research supercomputers.

Current customers have used the software as a visualization tool for all types of finite element modeling, including structural analysis, fluid dynamics, electromagnetics, injection molding and thermodynamics. For more information about EnSight on Cray Research systems, contact your local Cray Sales office, or email ensight@cray.com. CEI may be contacted by phone at 919/481-4301 or by email at ensight@ceintl.com.

Cray Research debuts low-priced DD-301 disk drive

Cray Research has introduced a new disk drive product priced at one-third the unit cost of the company's previous lowest-priced drives. The DD-301 drives, priced at $6000 each (in the U.S.), are the first intelligent peripheral interface (IPI), 3.5-inch disk drive product for Cray Research supercomputers. Each DD-301 drive delivers a sustained data transfer rate of 8.2 Mbytes/s and has a storage capacity of 1.4 Gbytes. The new disk drive is also available in a RAID level-3 disk array package that includes five DD-301 disk drives.

Cray Research: on the web

Cray Research is now serving up information about its products, services, company background, financial results, and more via the World-Wide Web (WWW). The WWW, which grew out of an initiative at CERN (the European Laboratory for Particle Physics) into one of the Internet's most popular information systems, provides a quick and easy means of accessing a wide variety of information including text, graphics, sound, and video.

The Cray Research Internet Information Server currently offers information on all of our major product families, including product descriptions and specifications, pictures of the products, configuration diagrams, and in some cases, online manuals.

Other highlights include a hypertext version of the 1994 Directory of Applications Software for Cray Research systems and a section on CraySoft. Use the URL http://www.cray.com/ and your favorite World-Wide Web browser to access the server.

APPLICATIONS UPDATE

CRI/*Turbo*Kiva 2.0

To help significantly reduce the average two-year period needed to test efficiency and emissions of new engines, Cray Research has introduced CRI/*Turbo*Kiva 2.0, an upgrade to the powerful software environment. Developed in cooperation with the T-3 Group at Los Alamos National Laboratory (LANL), Los Alamos, New Mexico, CRI/*Turbo*Kiva 2.0 is based on the KIVA-3 computational fluid dynamics (CFD) code from LANL.

CRI/*Turbo*Kiva 2.0 is the first engine combustion software to support the simulation of two-stroke engines, an important software feature as the automotive and engine industries look to two-stroke engine technology to design low-vibration, compact engines that are less expensive to manufacture and less polluting than today's engines.

Two-stroke engines have unique flow characteristics that make it challenging to predict emissions. CRI/*Turbo*Kiva software models fluid flows—air intake, fuel injection, combustion, and exhaust characteristics—of reciprocating internal combustion engines. CRI/*Turbo*Kiva 2.0 can be used to simulate all these engine types and will be an attractive computational tool for the wider engine market with this new capability.

"CRI/*Turbo*Kiva actually helps engineers see the combustion process in more detail. It's almost like watching the engine breathe on the computer workstation screen," said Cray Research's Reza Taghavi. "In order to design the best engine, good analysis is required of the flow and combustion—particularly in two-stroke engines, where efficient scavaging and pumping processes are key in designing the most efficient, least polluting engine possible."

CRI/*Turbo*Kiva is designed for use on all Cray Research supercomputing systems and is currently being used by 13 automotive, bus, truck, and engine manufacturers and universities throughout the world, including Nissan Motors, Hyundai, and the Korean Institute of Science and Technology. For more information on CRI/*Turbo*Kiva 2.0, call Lee Roll at 612/683-3552, or email lee.roll@cray.com.

Renault eliminates die tryout with simulation

In an effort to eliminate the time- and labor-consuming die tryout phase of a typical auto stamping part, the Production Technology Division of Renault has used PAM-STAMP software for sheet-metal forming simulation. For the production of deep drawn exterior body panels, soft tools are built first to determine the manufacturing feasibility of certain parts. Once a part geometry is defined in the design process, decisions are made on how the part can be manufactured successfully. From the part geometry, die engineers design the tooling surfaces, including the punch, die cavity, blank holder and the die addendum. Draw beads are added to the die addendum if deemed necessary to ensure a well-shaped part.

Common difficulties encountered in the forming process include the formation of wrinkles, bulges, and the occurrence of material failures such as splitting. In a traditional manufacturing process, a proposed tooling design is first tested with soft tools, which are less expensive to produce and allow for easier minor modifications. During this die tryout phase, it becomes clear whether the proposed die design is capable of producing a part that will satisfy rigorous requirements for allowable thickness reduction, maximum elongations, and surface quality. After this die tryout phase is successfully completed, the hard tools used for the actual production stampings are manufactured and put into service.

Responding to the tremendous pressures of today's shortened design cycles in the automotive industry, Engineering Systems International, the supplier and manufacturer of the PAM-CRASH crashworthiness code, designed PAM-STAMP, a numerical simulation tool to address the needs of the sheet forming community.

At Renault, some recent work done with PAM-STAMP by Bernard Thomas, Mostapha El Mouatassin, and Benoit Abadie has clearly demonstrated the benefit of using numerical simulation techniques to eliminate the die tryout phase for deep drawn parts, which produced substantial savings in time and materials. This novel approach of virtual prototyping successfully answered questions about the manufacturability of a suspension dome for the Renault Twingo model. This example shows that numerical simulation of sheet-forming processes can be applied in today's design environment in a very cost-effective manner.

For the past several years, auto companies such as Renault have been working closely with ESI developers in designing a new software tool that can realistically simulate the sheet-forming process. To date, the software PAM-STAMP is routinely used in Renault as a result of this joint development.

Further efficiencies in integrating virtual testing into the design process are clearly achievable through the use of Cray Research supercomputers. Because the numerical simulation of the sheet-forming process is an extremely challenging applied mechanics problem of a very compute-intensive nature, supercomputers will allow for exploration of many design alternatives, without incurring undue delays.

Big memory speeds up data processing at ARSC

Race car drivers and supercomputer users share the quest for speed. "Give me fast engines" or "Give me fast CPUs."

But researchers are discovering that raw CPU power coupled with in-memory, or in-core, programming techniques can achieve even faster computing times.

Researchers at the Arctic Region Supercomputing Center (ARSC) are optimizing programs to get speed advantages from the center's four-processor, CRAY M98 system with 1 Gword of main memory. This system, named Denali, is one of the largest-memory supercomputers in the world. (Denali is the Alaska Native name for the tallest mountain in North America.)

A large-memory supercomputer can have a shorter time-to-solution than a faster-CPU, small-memory computer. This is because in-memory problem solving is faster than problem solving that requires input/output (I/O) between main memory and rotating storage disks.

The speed advantage of ARSC's large-memory supercomputer is clear to meteorologist John W. Glendening. He writes, ". . . the large memory size of the CRAY Y-MP computer at ARSC allows me—for the first time—to run my boundary-layer turbulence model in core. Because of the large matrices involved, requiring 270 Mbytes of in-core memory, I have previously been forced to utilize a matrix-swapping technique, which demands large amounts of I/O and is therefore inefficient."

A simple matrix multiply shows the speed advantage of in-memory versus out-of-memory problem solving.

David Slowinski, a Cray Research computer scientist said, "A high-performance matrix multiply can be programmed with just four lines of Fortran code when the data fits in main memory. But if the main memory is not big enough to hold all of the data, a more complicated out-of-memory method is necessary. A crude out-of-memory matrix multiply can be written with about 60 lines of Fortran, but the performance would be poor."

Optimizing this out-of-memory code can greatly improve performance, but the code gets bigger and looks less and less like a matrix multiply.

For example, an out-of-memory matrix multiply with multiple streams of asynchronous I/O requires perhaps 1000 lines of code. On Denali, in-memory code would take four lines and much less I/O.

All researchers—smaller memory and large memory users—can obtain speedups from other code optimization techniques. One of the easiest methods is to replace user-written code with one of the high-performance routines in the Cray LibSci optimized numerical library or the Cray LIBM optimized math library. For example, rather than write a matrix multiply program, a researcher could call a matrix multiply subroutine from the Basic Linear Algebra Subprograms (BLAS) included in LibSci.

Another technique is to inline repeatedly called subroutines to reduce the overhead associated with the multiple subroutine calls.

Optimization does not require a complete rewrite of a researcher's program. In a typical code, a 90-10 rule applies: 90 percent of the time is spent in only 10 percent of the code. Strategically placed optimizations can bring good speedups.

Software tools on Denali such as hpm, perfview, and profview help researchers identify bottlenecks in their code so optimizations can be concentrated in the few areas where they will result in the most speedup.

The Arctic Region Supercomputing Center's big-memory CRAY M98 system went online in January 1993. From January through December 1993, main memory usage increased 2600 percent. This increased memory usage is, in part, a result of researchers learning how to optimize code to gain processing speed.

Reprinted with permission from the Arctic Region Supercomputing Center's newsletter Challenges, *Vol. 2, No. 1., written by Eric Muehling. ARSC is the world's northernmost supercomputing center and operates a CRAY M98 system with 8 Gbytes memory and a CRAY T3D massively parallel processing system with 128 processing elements.*

CUG NOTES

Gary A. Jensen, CUG President

Why do Cray Research customers join the Cray User Group?

Cray User Group (CUG) meetings provide a forum for over 200 Cray Research customers to discuss issues that are defining "supercomputing" today. These gatherings provide an opportunity for Cray Research customers and users to voice their needs, concerns, and expectations. CUG meetings also enable Cray Research technical staff to provide guidance and direction to users to help solve common and individual challenges. In short, CUG members will help guide Cray Research into the future.

The CUG Board of Directors, comprising representatives from other companies, meet three times a year with Cray Research executives to discuss plans for hardware and software and goals in areas such as reliability, resiliency, and repairability. CUG Special Interest Committees (UNICOS and Compilers, for example) have influenced important improvements in Cray Research products.

CUG meetings take place twice a year; meetings have been held within the United States, Europe, and Asia. The next CUG meeting will be held on March 13-17, 1995, in Denver, Colorado—the mile-high city. The meeting will feature U.S. Government room rates for all attendees, and the rate can be extended a few days before and after the meeting. The conference hotel is the Denver Radisson hotel on the 16th Street Mall in downtown Denver; the hotel phone number is 303/893-3333.

The theme for the Denver CUG meeting will be "Mile-High Performance." By conference time, the debate on faster shared-memory, vector parallel processing supercomputers versus massively parallel processing is sure to be a lively one. Another exciting development is the inclusion of an applications track into the CUG program. There will also be one full day of CUG parallel sessions dedicated to the CRAY T3D system. A Newcomers Reception is planned for Tuesday, March 14, for anyone who has not previously attended a CUG meeting to provide an opportunity to meet the CUG Committee Chairs and Board members.

Cray Research is a company that is vital to the computing needs of governments and corporations worldwide. Their customers' needs and ideas are integral to supercomputing today and tomorrow. It is in all of our best interests that Cray Research be in a position to deliver the products we need, at the time we need them. We can do this by working together.

We encourage you to join the Cray User Group, and participate in these worthy efforts to help ensure both your success and that of Cray Research. So dust off your cowboy hat and mark March 13-17, 1995 on your calendar!

For information on becoming a CUG member, please contact Joan Palm at cug@cug.org

GALLERY

"The Wading Pool." A 1500-line by 2000-pixel image created by Melvin Prueitt of Los Alamos National Laboratory. Prueitt used his SCOPE ray-tracing program on a CRAY Y-MP system to create the image, which incorporates reflection, texture mapping, bump mapping, and two light sources.