




1 of 1

RECEIVED
NOV 19 1993
OSTI

UCRL-ID 112772
Reprinted from *Engineering Research,
Development, and Technology.*
Thrust Area Report FY 92
UCRL 53868-92

Computational Mechanics

Gerald L. Goudreau

March, 1993



Lawrence
Livermore
National
Laboratory

This is an informal report intended primarily for internal or limited external distribution. The opinions and conclusions stated are those of the author and may or may not be those of the Laboratory.

• Work performed under the auspices of the U.S. Department of Energy by the Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

DISSEMINATION OF THIS DOCUMENT IS UNLIMITED

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California and shall not be used for advertising or product endorsement purposes.

This report has been reproduced
directly from the best available copy.

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information
P.O. Box 62, Oak Ridge, TN 37831
Prices available from (615) 576-8401, FTS 626-8401

Available to the public from the
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.,
Springfield, VA 22161

Computational Mechanics

Gerald L. Goudreau

MASTER

THE MASTER OF THIS DOCUMENT IS ON FILE

rb

Computational Mechanics

The Computational Mechanics thrust area sponsors research into the underlying solid, structural, and fluid mechanics and heat transfer necessary for the development of state-of-the-art general purpose computational software. The scale of compu-




tational capability spans office workstations, departmental computer servers, and Cray-class supercomputers. The DYNA, NIKE, and TOPAZ codes have achieved world fame through our broad collaborators program, in addition to their strong support of on-going Lawrence Livermore National Laboratory (LLNL) programs. Several technology transfer initiatives have been based on these established codes, teaming LLNL analysts and researchers with counterparts in industry, extending code capability to specific industrial interests of casting, met-

alforming, and automobile crash dynamics. The next-generation solid/structural mechanics code, ParaDyn, is targeted toward massively parallel computers, which will extend performance from gigaflop to teraflop power.

Our work for FY-92 is described in the following eight articles: (1) Solution Strategies: New Approaches for Strongly Nonlinear Quasistatic Problems Using DYNA3D; (2) Enhanced Enforcement of Mechanical Contact: The Method of Augmented Lagrangians; (3) ParaDyn: New Generation Solid/Structural Mechanics Codes for Massively Parallel Processors; (4) Composite Damage Modeling; (5) HYDRA: A Parallel/Vector Flow Solver for Three-Dimensional, Transient, Incompressible Viscous Flow; (6) Development and Testing of the TRIM3D Radiation Heat Transfer Code; (7) A Methodology for Calculating the Seismic Response of Critical Structures; and (8) Reinforced Concrete Damage Modeling.

Gerald L. Goudreau
Thrust Area Leader

Section 2



2. Computational Mechanics

Overview

Gerald L. Goudreau, Thrust Area Leader

Solution Strategies: New Approaches for Strongly Nonlinear Quasistatic Problems Using DYNA3D

Robert G. Whirley and Bruce E. Engelmann 2-1

Enhanced Enforcement of Mechanical Contact: The Method of Augmented Lagrangians

Bradley N. Maker and Tod A. Laursen 2-7

ParaDyn: New Generation Solid/Structural Mechanics Codes for Massively Parallel Processors

Carol G. Hoover, Anthony J. De Groot, James D. Maltby, and Robert G. Whirley 2-11

Composite Damage Modeling

Edward Zywickz 2-15

HYDRA: A Flow Solver for Three-Dimensional, Transient, Incompressible Viscous Fluid

Mark A. Christon 2-19

Development and Testing of the TRIM3D Radiation Heat Transfer Code

James D. Maltby 2-23

A Methodology for Calculating the Seismic Response of Critical Structures

*David B. McCallen, Francois E. Heuze,
Lawrence J. Hutchings, and Stephen P. Jarpe 2-27*

Reinforced Concrete Damage Modeling

Sanjay Govindjee and Gregory J. Kay 2-35

Solution Strategies: New Approaches for Strongly Nonlinear Quasistatic Problems Using DYNA3D

**Robert G. Whirley and
Bruce E. Engelmann**

*Nuclear Explosives Engineering
Mechanical Engineering*

The analysis of large, three-dimensional, strongly nonlinear structures under quasistatic loading is an important component of many programs at Lawrence Livermore National Laboratory (LLNL). The most widely used formulation for this type of problem is an implicit solution process with a linearization and iteration approach to solving the coupled nonlinear equations that arise. Our research investigates an alternative approach, in which an iterative solution method is applied directly to the nonlinear equations without the use of a linearization. This approach alleviates some of the difficulties encountered when linearizing nonsmooth phenomena such as mechanical contact. The first iterative method explored is the dynamic relaxation method, which has been implemented into the LLNL DYNA3D code, and combined with software architecture and computational mechanics technology developed for explicit transient finite element analysis. Preliminary analysis results are presented here for two strongly nonlinear quasistatic problems to demonstrate the promise of a linearization-free approach.

Introduction

Many programs at Lawrence Livermore National Laboratory (LLNL) use nonlinear finite element structural analysis to guide engineering projects. Applications include the determination of weapon component response to a variety of structural and thermal environments; the study of stresses in nuclear fuel transportation casks; and the simulation of the forming of sheet metal parts to optimize processing parameters and minimize waste. These applications share the common features of being three-dimensional (3-D), quasistatic, and strongly nonlinear, and illustrate wide use of this type of computer analysis.

Nonlinear finite element structural analysis methods may be divided into two categories: implicit methods and explicit methods. Implicit methods are typically used for quasistatic and low-frequency dynamic problems. The LLNL NIKE2D¹ and NIKE3D² codes are based on an implicit formulation. This approach uses a small

number of large increments to step through the simulation, with the increment size chosen by the analyst to satisfy accuracy and convergence requirements. An implicit analysis code must solve a coupled system of nonlinear algebraic equations at each step, usually by a linearization and iteration procedure. This linearization leads to a coupled system of linear algebraic equations that must be solved at each iteration of each step in the analysis. Typically, the iteration process is continued within a step until some convergence measure is satisfied, then the solution is advanced to the next step.

Explicit methods are typically used for high-frequency dynamics, wave propagation, and impact problems. The LLNL DYNA2D³ and DYNA3D⁴ codes are based on an explicit formulation. In contrast to implicit methods, explicit methods use a large number of small increments to step through a problem, with the increment size chosen automatically to satisfy stability requirements. This stability requirement essentially dictates that

the time increment size must be smaller than the time it would take a stress wave to propagate across the smallest dimension of the smallest element in the mesh. An explicit code does not solve coupled equations at each step, and therefore the update from step to step is much faster than in an implicit code.

In practice, implicit methods have worked well for strongly nonlinear quasistatic problems in two dimensions, but have encountered difficulties on 3-D problems. These difficulties can be attributed to three primary factors. First, large 3-D contact problems, especially sheet-forming problems, have large matrix bandwidths due to the large contact area between the sheet and tool surface. This large matrix bandwidth translates into high computer memory requirements and expensive linear solutions at each iteration of the nonlinear solution process. Second, strongly nonlinear problems often contain discontinuous phenomena that are difficult to linearize. For example, in a contact problem, the interface pressure abruptly changes from zero when two bodies are separated to a finite value when the bodies come into contact. Obtaining an accurate linearization of such abrupt changes is difficult, and this is manifest in the code as slow convergence or nonconvergence of the linearization and iteration procedure within a step. Finally, when solution difficulties are encountered in a large 3-D problem, it is much more difficult to troubleshoot the model than it would be in a similar two-dimensional model. Often there are few clues to suggest why the iteration procedure is having difficulty converging to a solution. The development of a more robust solution strategy for strongly nonlinear quasistatic problems is the primary objective of this effort.

One approach to improving the performance of implicit methods for large, 3-D, strongly nonlinear quasistatic problems focuses on the solution of the large linear system that arises from the linearization and iteration approach. An iterative method, such as the use of a preconditioned conjugate gradient, is one approach to solving the linear system. This approach was investigated in the LLNL NIKE3D code,⁵ and culminated in the development of an iterative solver now used in the production code version. Although this approach reduces memory requirements and may reduce CPU costs for the linear equation solution, it does nothing to improve the convergence of the nonlinear iteration.

An alternative approach for difficult quasistatic problems is to use an explicit transient dynamics code, and apply the loads so slowly that the dynamic effects are negligible, and therefore a quasi-

static solution is obtained. Although this approach is often used by engineering analysts, it does have several disadvantages. First, the best rate of load application to minimize dynamic effects while keeping the analysis cost tolerable is not known *a priori*, and often requires some experimentation. Also, it is important to minimize artificial oscillations in the solution when history-dependent material models such as plasticity are included, and this further complicates the choice of analysis parameters. Finally, this approach obtains only an approximate quasistatic solution, and the amount of error due to dynamic effects requires some effort to quantify.

These observations suggest the alternate approach followed in our work. The basic linearization and iteration paradigm is abandoned, and an iterative solution method is applied directly to the nonlinear equations. This method is combined with much of the computational mechanics technology and software architecture developed for explicit transient dynamic analysis to produce a code that solves the nonlinear problem directly by using a large number of rather inexpensive iterations, and without solving a coupled linear system. The essential elements of this approach and its development in the LLNL DYNA3D code are described below.

Progress

In FY-92, we developed an iterative quasistatic solution capability in DYNA3D, based on the dynamic relaxation (DR) method. In addition to the implementation of the basic DR procedure, a load incrementation framework has been incorporated into DYNA3D that allows a true quasistatic solution to be obtained at a load level before the load is increased for the next increment. In addition, a spectrum contraction algorithm has been implemented that greatly improves the efficiency of the method. Also, extensions have been developed for the rigid-body mechanics formulation and the treatment of boundary conditions to accommodate nonlinear quasistatic problems within the DYNA3D framework. The resulting code is now being used as a testbed to evaluate the overall robustness and efficiency of the DR method, and to study improvements in the formulation, contact algorithms, and adaptive damping procedures.

Overview of the DR Approach

In the DR method, the equations governing a quasistatic analysis are first transformed into those governing a dynamic system. The nonlinear cou-

pled equations for a quasistatic problem may be written as

$$\mathbf{p}(\mathbf{x}_0) = \mathbf{f}, \quad (1)$$

where $\mathbf{p}$ is the internal nodal force vector resulting from stress states in the finite elements; $\mathbf{x}_0$ is the nodal displacement solution; and $\mathbf{f}$ is a vector of externally applied loads. An associated dynamic problem may be written as

$$\mathbf{M}\ddot{\mathbf{x}} + \mathbf{C}\dot{\mathbf{x}} + \mathbf{p}(\mathbf{x}) = \mathbf{f}, \quad (2)$$

where dots denote differentiation with respect to time. With the appropriate choice of mass and damping matrices, $\mathbf{M}$ and $\mathbf{C}$, the solution of the dynamic problem as time gets large approaches the solution of the quasistatic problem, i.e.,

$$\lim_{t \rightarrow \infty} \mathbf{x}(t) = \mathbf{x}_0. \quad (3)$$

The iterative scheme is defined by applying the explicit central difference method to integrate the dynamic equations in time.

The success of the dynamic relaxation iterative method to solve highly nonlinear quasistatic problems depends on many factors including the specification of mass and damping, as well as the development of an incremental loading strategy.

Spectrum Contraction

The efficiency of the DR method may be improved by contracting the spectrum of the global equations. This is easily accomplished by proper choice of the mass matrix $\mathbf{M}$ in Eq. 2. The construction of the mass matrix should not dominate the computation, and thus it should be based on conveniently available quantities. In our algorithm, the mass matrix is constructed from an assemblage of element contributions. The mass matrix of each element is scaled so that all elements have a uniform critical time step, and thus information flows throughout the mesh at an optimal rate during the iteration process. This technique has proven quite useful in accelerating the convergence of the DR method.

Damping

The type and amount of damping can also significantly affect convergence. For linear systems, optimal damping depends on both the highest and lowest eigenvalues of the system. Although bounds on the highest eigenvalue are readily available from the element eigenvalue inequality,




Figure 1. Variation in the number of normalized iterations required for convergence, relative to the magnitude of the damping factor in the dynamic relaxation method.

estimates must be used for the lowest eigenvalue, which can vary greatly throughout a nonlinear simulation. When insufficient damping is used, iterates will oscillate around a solution and reach it very slowly. Too much damping will dramatically retard convergence, especially for problems that include large rigid-body motions such as the motion of the sheet in sheet-forming simulations. Results thus far indicate that adaptive damping approaches, based on the evolving physics of the problem, may prove most effective for highly nonlinear problems. Figure 1 shows the variation in the number of iterations required to converge relative to the magnitude of damping in the DR algorithm. The graph depicts the strong influence of damping value on the number of iterations required by the DR method to converge to a solution. The automated determination of the optimal damping value is a subject of ongoing investigation.

Cantilever Elastic Plate Example

To demonstrate some essential features of the new quasistatic solution capability in DYNA3D, an elastic cantilever plate was subjected to an applied moment on the free end. The problem was solved with two magnitudes of applied load: one




Figure 2. The initial and final deformed shapes of a cantilever plate subjected to an end moment. The upper figure corresponds to a small load, and the lower figure to a larger load. The solution to each of these problems was obtained using only one load increment with the DR method in DYNA3D.




Figure 3. Initial geometry for hydroforming simulation, showing the punch, blankholder, and sheet. The punch and blankholder are geometrically represented by 8-node continuum elements and are treated as rigid bodies. The sheet is represented by 4-node thin shell elements and is modeled as an elastic-plastic material.

that causes a small deformation of the plate, and one that causes an extensive 'roll-up' deformation. The initial geometry and the two final deformed shapes are shown in Fig. 2. An interesting observation is that the DR algorithm required approximately the same number of iterations to converge for both load cases. This is in contrast to conventional implicit solution techniques, where the number of iterations required to converge increases quickly with the degree of nonlinearity. This insensitivity of DR to the degree of nonlinearity is a powerful advantage of the DR method.

Sheet Metal-Forming Example

One major application of the quasistatic solution capability developed in this research is the numerical simulation of sheet metal-forming processes. These problems pose a real challenge since they involve large strains, material nonlinearities such as plasticity, and extensive sliding contact

with friction. In addition, the thin sheets have a wide spectrum due to the large difference between in-plane and bending stiffnesses, thus making them even more difficult for an iterative solver.


Figure 3 shows the finite element model for the numerical simulation of an aluminum hydroforming process. Pressure is applied to the upper surface of the sheet to hold it against the blankholder, and the punch is then advanced to form the sheet into the final shape shown in Fig. 4. The good comparison between the computed results and the shape of the actual part, including the failure locations, is illustrative of the power of a versatile quasistatic analysis tool.

This problem was first solved at LLNL in 1988 by running DYNA3D in an explicit dynamic analysis mode and applying the loads slowly to minimize dynamic effects, an approach requiring approximately two hours of CPU time. More recently, this problem was solved using the LLNL implicit code NIK3D, but it required somewhat more computation time. Using the newly developed iterative methods in DYNA3D, this solution has been obtained in approximately 20 minutes of CPU time. Further improvements in contact algorithms, adaptive damping algorithms, and code optimization should enable solution of problems such as this in even less CPU time and without trial and error. Although much remains to be done, these initial results indicate the promise of the iterative quasistatic solution method in DYNA3D.

Future Work

Our research in FY-92 has led to the development and implementation of a DR iterative strategy for quasistatic problems in the LLNL DYNA3D code. Four general conclusions can be made from our experience thus far: (1) overdamping in the

Figure 4. Comparison of actual deformed shape with that predicted by numerical simulation. The circled areas in the numerical results indicate regions of large strains, and these correspond closely with the tears observed in the real part.



DR method significantly slows the convergence rate, especially for problems with large rigid body motions; (2) the convergence rate of DR appears insensitive to the degree of nonlinearity in many problems; (3) the rate of load application within an increment is important, and a step function is probably not optimal; and (4) adaptive damping algorithms work extremely well for some problems, and are clearly desirable. More study and development will be required, however, before these algorithms can be used for general production analysis.

Our research efforts in FY-93 will explore the promising directions discussed above. We will refine adaptive damping DR algorithms and develop optimal load application schemes for a range of nonlinear quasistatic problems. We will also investigate new contact formulations to eliminate the solution noise introduced by the current penalty-based procedures. In addition, we will evaluate the utility of the nonlinear conjugate gradient algorithm for the problem classes of interest. Finally, the results of this effort will be optimized for vector computers and implemented into a future production version of the LLNL DYNA3D code for general use. In addition, the algorithms developed in this project will be implemented into the ParaDyn project to allow the solution of large quasistatic problems on massively parallel computers.

Acknowledgements

The authors wish to acknowledge Dr. Brad Maker of the LLNL Methods Development Group for sharing his early experiences on sheet forming with DYNA3D and for providing the finite element model and photographs for the sheet forming example.

1. B.E. Engelmann and J.O. Hallquist, *NIKE2D: A Nonlinear, Implicit, Two-Dimensional Finite Element Code for Solid Mechanics—User Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-105413 (1991).
2. B.N. Maker, R.M. Ferencz, and J.O. Hallquist, *NIKE3D: A Nonlinear, Implicit, Three-Dimensional Finite Element Code for Solid and Structural Mechanics—User Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-105268 (1991).
3. R.G. Whirley, B.E. Engelmann, and J.O. Hallquist, *DYNA2D: A Nonlinear, Explicit, Two-Dimensional Finite Element Code for Solid Mechanics—User Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-110630 (1992).
4. R.G. Whirley and J.O. Hallquist, *DYNA3D: A Nonlinear, Explicit, Three-Dimensional Finite Element Code for Solid and Structural Mechanics—User Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-107254 (1991).
5. R.M. Ferencz, *Element-By-Element Preconditioning Techniques for Large-Scale, Vectorized Finite Element Analysis in Nonlinear Solid and Structural Mechanics*, Ph.D. Thesis, Stanford University, Palo Alto, California (1989). 

Enhanced Enforcement of Mechanical Contact: The Method of Augmented Lagrangians

Bradley N. Maker

*Nuclear Explosives Engineering
Mechanical Engineering*

Tod A. Laursen

*Duke University
North Carolina*

We have introduced the method of augmented Lagrangians into our stress analysis codes, NIKE2D and NIKE3D. This approach provides a simple and effective enhancement to the penalty method for enforcing contact constraints. Also, by using augmented Lagrangians, accuracy is determined by physically motivated convergence criteria, independent of the penalty parameter.

Introduction

Contact between deformable bodies occurs commonly in mechanical systems. Stress analysis codes that are applied to multi-body systems and assemblies must accommodate this contact to be useful to design engineers. Our NIKE and DYNA finite element codes have a widely recognized capability to capture the mechanics of contact in complex systems, as the models in **Fig. 1** demonstrate. The results of this research effort have further enhanced our contact algorithms by introducing the method of augmented Lagrangians into NIKE2D and NIKE3D.

In the finite element method, bodies are discretized into assemblies of elements whose boundaries are described by a set of node points. In this context, mechanical contact conditions act to constrain the node points of one body from penetrating the boundary surface of another. **Figure 2** represents the discrete contact problem in two dimensions. Driven by the action of externally applied loads, a single node point from the 'slave' body penetrates the boundary of the 'master' body. This penetration is identified by a search algorithm, and activates the constraint enforcement algorithm. As the contact constraint is enforced, penetration is minimized, and stress and deformation are induced in each body. This deformation may cause other slave nodes to penetrate the master body, which in turn activates additional constraints. As this iterative process reaches equilibrium, the proper contact area and pressure

distributions that balance the applied loads are obtained.

This simple example highlights the nonlinear nature of the contact problem. Indeed, the deformation of each body may be large, generating both geometric and material nonlinearities. But the more fundamental nonlinearity in the contact problem arises from the discontinuous manner in which the contact area evolves. Since the surfaces are faceted, the contact area grows or shrinks in discrete increments. These abrupt changes in contact area are sharp nonlinearities, which complicate the equilibrium search process.

Progress

The constraint algorithm used to minimize penetration in most finite element codes, including our own, is the penalty method. This simple but effective approach introduces penalty springs between the two bodies wherever penetration occurs. As the penetration increases, the springs are stretched, generating forces that oppose further penetration. The springs act unilaterally, i.e., when the bodies separate, the penalty springs are removed, allowing gaps to open.

One obvious drawback of the penalty method is that penetration must occur before any constraint forces are generated. Thus, in the equilibrium state, where each penalty spring is properly stretched to balance the applied loads, the two bodies are interpenetrated, and the exact contact condition is violated.




Figure 1. Examples of NIKE3D contact algorithms applied to engineering problems: (a) the bolted superflange and (b) the Kestrel bulkhead.

To minimize this penetration, the penalty spring stiffness may be increased, generating a large contact force through a very small penetration. This approach works well in theory, but in practice introduces poor numerical conditioning, and inevitably numerical errors. But a more fundamental deficiency of the penalty method is that the

results are dependent upon the value chosen for the penalty stiffness. This effect is demonstrated in Fig. 3. Clearly, as a larger stiffness is chosen, the bodies are driven further apart, and the contact area and/or pressure changes. This arbitrariness motivated our work toward an enhanced constraint algorithm.

The augmented Lagrangian method is an effective and intuitively obvious enhancement to the penalty method, and proceeds as follows. Using the penalty method as a kernel, equilibrium is obtained in the usual manner. With known penetration depth and penalty stiffness, the contact force may be computed. This force is taken as the initial value for the Lagrange multiplier. The Lagrange multiplier defines a static load that is applied to the slave node, and the equilibrium search is then repeated. In the presence of the Lagrange multiplier load, penetration is reduced. The new penetration distance is then used to compute a new increment in contact force, the Lagrange multiplier is augmented by this increment, and the iteration process is repeated.

This equilibrium search and Lagrange multiplier augmentation loop proceeds until convergence is obtained. But now convergence may be defined in physically meaningful terms. For example, the augmentation loop can proceed until the contact force (Lagrange multiplier) stabilizes to within 1%, or until the largest penetration is less than a user-specified distance.




Figure 2. Enhanced enforcement of mechanical contact. (a) A search algorithm detects penetration of the master body by the slave node point. The penalty method introduces a spring of stiffness k between nodes S and M . When stretched, the spring generates interface force $F = kd$. (b) The augmented Lagrangian method applies F as a static force on nodes S and M , and iteratively augments this force, i.e., $F_{n+1} = F_n + kd_n$ until a convergence criterion is satisfied.




Figure 3. Using the penalty method, results vary dramatically with penalty stiffness. In this example, a contact interface is defined between two flat plates (arrow). The lower plate is fixed at its lower edge. A downward motion is prescribed to the upper edge of the upper plate. Erratic stress distributions result using NIKE3D's default penalty stiffness (a). Increasing the penalty stiffness by 10^4 produces a more uniform stress distribution (b). The augmented Lagrangian method gives the most accurate solution (c) using a convergence tolerance of 0.1% on the interface force. This same answer was obtained for penalty stiffnesses ranging from 10^2 to 10^4 . The new method therefore provides an insurance policy against errors from a poorly chosen penalty parameter.

The new method has several advantages. In the limit of a large number of augmentations, equilibrium contact force is obtained without penetration. Further, the solution is independent of penalty parameter, since augmentations proceed until the (physically based) convergence criterion are satisfied. The exception to this independence is the case where the penalty stiffness is chosen so large that the original penalty method (the kernel of the new method) will not converge due to numerical conditioning. This case is obviously moot, since both methods fail.

The obvious drawback to the augmented Lagrangian method is that an additional iterative loop is introduced into the solution process. For a very soft choice of penalty parameter, this iteration loop can be slow to converge. However, our implementation allows for immediate convergence with no iteration if the penalty stiffness is cleverly (or luckily) chosen to satisfy convergence criteria in the first step. The method is therefore an insurance policy against a poor choice of penalty parameter, which before would have yielded an inaccurate result.

The final and perhaps most dramatic advantage to the new method is that the Lagrange multipliers are preserved for use in the next loading step. Thus, for a problem in which load is applied

in several steps, the initial guess at contact pressure is the converged value from the previous step. This history information often speeds convergence of the equilibrium search in the second and later steps in the problem, and can result in an overall reduction in CPU run time for a complex problem.

The augmented Lagrangian method provides a simple and effective enhancement to the penalty method for enforcing contact constraints in NIKE 2D and NIKE 3D. Accuracy is determined by physically motivated convergence criteria, and is independent of the penalty parameter.

Future Work

The method of augmented Lagrangian also offers a new mathematical framework for considering the frictional contact problem, which will be pursued in future work.

Acknowledgement

We gratefully acknowledge the extensive collaboration of Dr. Bruce Engelmann in the algorithm development and NIKE 2D implementation, and Messrs. M. A. Gerhard, D. L. Trummer, and L. A. Platt for the superlunge and keshet examples. ■

ParaDyn: New Generation Solid/Structural Mechanics Codes for Massively Parallel Processors

Carol G. Hoover

*National Energy Research
Supercomputer Center
Computation Directorate*

James D. Maltby

*Nuclear Test Engineering
Mechanical Engineering*

Anthony J. De Groot

*Engineering Research Division
Electronics Engineering*

Robert G. Whirley

*Nuclear Explosives Engineering
Mechanical Engineering*

The objective of this work is to develop DYNA3D for massively parallel computers. In this last year, we have worked with the DYNA2D program on a Thinking Machines CM-5 computer to develop strategies for distributing the data and parallelizing the finite element algorithms. We are using the experiences gained with DYNA2D to guide the parallelization of the algorithms for the much larger and more complex DYNA3D. We have measured performances comparable to Cray Y-MP speeds for a DYNA2D test problem on systems with as many as 512 processors. The performance results show moderately large communication times relative to computing times, particularly for the global force assembly (scatter). We attribute this performance to the early developmental releases of the CM-5 software.

Introduction

Recent advances in microprocessor chip technology and parallel computer architectures are revolutionizing the concept of supercomputing. Vector supercomputer architectures have reached technology limits that preclude the orders-of-magnitude performance improvements expected for the massively parallel architectures.¹ A massively parallel computer is an arrangement of hundreds to thousands of microprocessors interconnected with a high-speed internal network (currently up to 250 megabytes/s). Typical microprocessor peak speeds range from a low of 10 MFLOPS per processor to a high of 100 MFLOPS per processor for pipelined (vector-like) processors. Performances between 10 and 100 GFLOPS are possible today on systems of 1000 processors. By comparison, the latest vector super-

computer is a 16-processor system with a peak performance of 1 GFLOP per processor. The motivation for developing a parallelized version of the solid mechanics programs (DYNA and NIKE) is the potential in the next three to five years for running applications that are larger by two or three orders of magnitude than are possible on vector supercomputers. This would allow simulations of hundreds of millions of elements rather than a few hundred thousand elements with DYNA3D.

Figure 1 illustrates the speed and storage requirements for typical advanced applications in metal forming, materials science, earthquake simulations, and crash dynamics. Notice in **Fig. 1** the increased complexity of the models for points in the upper right portion. These applications are of high value in government research and for their impact on industrial competitiveness.




Figure 1. Advanced applications for massively parallel processors. The data represents systems as follows:
E = earthquake simulations: (a) bridges, buildings and other structures, and (b) full Bay Area earthquake simulations.
C = crash dynamics simulations: (a) automobile component simulations, (b) automobile/barrier simulations, (c) multiple automobile crash simulations, and (d) aircraft crash simulations.
M = metal forming applications: (a) two-dimensional simulations, and (b) three-dimensional simulations.
T = tribology and nanometer-scale machining simulations: (a) large-scale (10 to 100 million atoms) molecular dynamics simulations with no electronic structure (ab initio) calculations, and (b) hybrid molecular dynamics and continuum mechanics models with billions of particles/zones and electronic structure calculations with million particle molecular dynamics. Projected times for the next generations of systems are given along the top of the plot.

Progress

The DYNA3D program is nearly twice the size of DYNA2D, and the three-dimensional algorithms (e.g., for contact between slide surfaces) are more elaborate. Our strategy is to experiment with conversion techniques, parallel language paradigms, and algorithm parallelization with DYNA2D rather than with DYNA3D.

The development of a parallelized version of a large vectorized program necessarily proceeds in steps. The first and most tedious step is the conversion of array storage. The storage allocation for a distributed-memory massively parallel computer is dramatically different than for a common-memory serial computer. Careful analysis of reused storage, detailed conversion of array layout, and parameterization of the element vector block length absorbed well over one third of our effort. A benefit of this work and of the following timing analysis has been the insight we have gained into techniques for greatly reducing this same effort for the DYNA3D conversion.

The computationally expensive step in the DYNA algorithm is the element-by-element force evaluation. The vectorized version of this element processing translates readily into the data parallel paradigm on the CM-5. We have completed a data parallel version of the force update and time integration for

Table 1. Timing for the 7 cycles of an elastic/plastic bar impacting a rigid wall. There are 32,768 elements in the 64x512 mesh. Results are for a 512-processor CM-5.[†] The gather time is associated with the block processing for multiple material and element formulations. The scatter time is associated with the global force assembly step. The parallel reduction time accrues for calls to an intrinsic parallel library routine.

	June 1992	November 1992
Processor CPU time	29.3%	20.8%
Gather/scatter time	33.8%	24.5%
Parallel Reductions	0.8%	0.7%
Front-end to processor time	36.1%	54.0%
Totals	183.7 s	0.962 s

CM-5 elapsed time for 7 cycles:

2.03 s

Time per element-cycle (November 1992):

9 μ s for the 512 processor CM-5
 6 μ s for the Cray Y-MP

The above results do not include the use of vector software. At the time of this printing our timing results for the CM-5 have improved to .732 μ s per element time step.

[†] Disclaimer by Thinking Machines Corporation. These results are based upon a test version of the software where the emphasis was on providing functionality and the tools necessary to begin testing the CM-5 with vector units. This software release has not had the benefit of optimization or performance tuning and, consequently, is not necessarily representative of the performance of the full version of this software.




Figure 2. Partitioning of a standard three-dimensional mesh for an automobile piston: (a) the unstructured, three-dimensional finite element mesh, (b) the partitioning of the piston mesh for two processors, (c) the partitioning of the piston mesh for four processors, (d) the partitioning of the piston mesh for eight processors.

both elastic and elastic-plastic material models. We chose a standard test case, a bar impacting a rigid wall, for timing and performance analysis.

Balancing the parallel and scalar calculation time with network time is essential for efficient use of a massively parallel computer. An unbalanced problem with communication times exceeding calculation times prevents the desirable linear speed-ups predicted by Amdahl's law. On the CM-5, the performance analysis tool, PRISM, has been effective in providing the breakdown of hardware times. The most valuable feature of PRISM is the availability of timing data throughout the program, from upper level subroutines down to individual FORTRAN statements. Two timing analyses are shown in Table I. The performance difference in the two runs is a combined effect of hardware, software changes at the Army High Performance Computing Research Center, and several programming changes inspired by the PRISM statistics. The speed achieved for our most recent run is 9 microseconds per element cycle, which is comparable to the one processor V-MIP performance of 6 microseconds per element cycle.

The communication times we have measured are still excessively high for a balanced calculation. We have collaborated with computational analysts at Thinking Machines to analyze the imbalance in the timings for this test problem. The development of system software such as compilers and communication libraries for massively parallel systems is in its infancy. We find that the newer alpha-test versions of the software, used now by the company analysts, will change these results up to an order of magnitude. This software may be available to us within the next three to six months. With the new versions of the software, we expect to exceed single processor Cray C-90 performances for a single material problem with a regular topology. At the time of this printing our timing results for the CM-5 have improved to 732 μ s per element time step.

Several techniques have been developed for balancing the computational work among processors while minimizing the communication time. We are testing a recursive spectral bisection technique with three dimensional meshes. We have developed a method for visualizing the results, as shown in Fig. 2.

Future Work

We will continue to use DYNA2D to experiment with algorithm parallelization. In this next year, we plan to investigate: (1) message-passing and data-parallel versions of selected contact algorithms, (2) a data-parallel and message-passing hybrid system software model available in the next year from Thinking Machines, and (3) parallel table lookup and sort algorithms, which are appropriate for contact algorithms. We will begin the conversion and parallelization of DYNA3D and develop kernel algorithms for DYNA3D for further evaluation of parallel programming paradigms and architectures.

DYNA3D is an 80,000 line analysis program including ten finite element formulations (solid elements, shells, and beams), 35 material models, equations of state for hydrodynamic models, several algorithms for contact at arbitrary interfaces, and a list of additional boundary conditions and mechanics algorithms, all of which make the program one of the most widely used tools for nonlinear structural response simulations. Our plan over the next three years for demonstrating a prototype massively parallel version of DYNA3D includes implementing an eight-node solid (continuum)

element, at least one shell element, one contact algorithm, and solid/structural mechanics capabilities needed for three large-scale demonstration problems. The demonstration problems include the simulation of a nanoindentation problem, an automobile/barrier simulation, and a weapons penetration application.

Acknowledgement

We gratefully acknowledge the Army High Performance Computing Research Center for providing CM-5 computer time for this work as part of their General Plan for Developing Structural Analysis Programs for Advanced Massively Parallel Computers. Funding for computer time was supported by, or in part by the Army Research Office contract number DAAL03-89-C-0038 with the University of Minnesota Army High Performance Computing Research Center. We thank Earl Renaud from Thinking Machines Corporation for his advice and cooperation.

1. B. Bhoghosian, *Comput. Phys.* **4**, 1 (1990).
2. H.D. Simon, *Computing Systems in Engineering* **2** (2/3), 135 (1991). 

Composite Damage Modeling

Edward Zywicz

*Nuclear Explosives Engineering
Mechanical Engineering*

A progress damage model for continuously reinforced, polymeric-matrix composites is being developed and implemented in the implicit finite element code NIKE3D. The constitutive model replicates the discrete laminae with an equivalent homogenized material prior to the onset of damage. Failure criteria eventually trigger damage evolution laws that track individual failure mechanisms within each lamina and degrade the stiffness and strength of the laminated composite. Failure criteria and damage evolution laws are currently being developed, as well as numerical procedures, to efficiently address the multilayer nature of laminates. This work will allow analysts to simulate the redistribution of load as the composite materials degrade and, therefore, to design minimal mass composite structures.

Introduction

Continuously reinforced, polymeric-matrix composites offer substantial weight savings over conventional materials, such as steels and aluminums, and at the same time provide equal or superior mechanical properties. For example, at Lawrence Livermore National Laboratory, continuously reinforced graphite/epoxy (Gr/Ep) composites are used in lightweight earth-penetrator weapons, advanced conventional munitions, and enhanced nuclear safety systems. Commercial applications of Gr/Ep composites include high-speed aircraft, automobile drive shafts, bicycles, and tennis rackets. Currently, components manufactured with continuously reinforced, polymer-based composites are designed very conservatively or must be tested extensively, because the failure response of the material is not fully understood. To overcome this barrier, a composite damage model is currently being developed and implemented in the implicit finite element code NIKE3D.¹

A progress composite damage model permits analysts to simulate the complex three-dimensional (3-D) response of composite components in both subcritical (e.g., dings in an aircraft wing) and catastrophic (e.g., car crashes) loading environments. To be useful, the damage model must accommodate a wide range of fiber layups, track damage evolution based upon individual failure mechanisms, and predict resid-

ual life and strength. At the same time, the material model must be numerically efficient and resolve the complex lamina behavior within each laminate region modeled.

Progress

A continuum-based framework has been assembled to represent composite behavior. The approach uses conventional 8-node, solid isoparametric, 3-D elements with conventional $2 \times 2 \times 2$ Gaussian quadrature. Element stresses and stiffness are calculated in the usual way at each Gaussian point; however, the constitutive evaluations use homogenized material properties that are calculated uniquely for each element. During initialization, the 'virgin' elastic properties of all laminae present within each element, which can vary between one and two hundred, are homogenized and stored along with element-level, strain-based, failure criterion coefficients. Throughout the analysis, the small-strain, finite-deformation (total-Lagrangian)-based constitutive relation continually updates the element stresses, using effective stiffnesses and monitors for failure initiation, prior to the onset of damage.

Element-level failure triggers an in-depth lamina level or microanalysis. The microanalysis checks for failure, and tracks and evolves individual damage mechanisms for each lamina present in the element. Furthermore, it degrades

the individual lamina stiffnesses and calculates a material tangent matrix. The updated stiffnesses and tangent matrices are then homogenized for use at the element level.

The two-tier homogenized approach provides a rational and precise mechanism for tracking and integrating the complex response of damaged laminated composites. For undamaged material points, the homogenization technique, which incorporates bending and coupling effects, yields accurate solutions at a substantial computational savings, since laminate integration is performed only in initial-ization. Traditional methods use single elements or integration points for each lamina or homogenized material property and neglect bending and coupling effects. Efficiency in the undamaged region is very important since, in general, only small regions of typical composite components reach critical load levels. To date, the element homogenization technique, a conservative, element-level, strain-based failure criteria for fiber-direction strain to failure, and a micro-lamina-level subintegration scheme have been developed, implemented, and verified.

Element Homogenization and Representation

Homogenized stiffness functions² are used in the element to represent the total sub-lamina response. Within an element, the homogenized local stiffness $C^{II}(z)$ is given by

$$C^{II}(z) = C_0^{II} + C_1^{II}z + C_2^{II}\cos(\alpha z), \quad (1)$$

where C_0^{II} , C_1^{II} , and C_2^{II} are element-based stiffness matrices, z is the normalized distance from the element's central plane, and α is a constant used to minimize element integration error. To determine C_0^{II} , C_1^{II} , and C_2^{II} , all laminae present in an element are identified. Next, using the closed-form long wavelength solution of Pagano,³ the current lamina stiffnesses are 'integrated' through the element thickness, yielding the effective extensional (A), coupling (B), and bending (D) matrices of the element. This approach treats each element as a unique sublaminate. Using the same long wavelength procedure, Eq. 1 is integrated. The resulting extensional, coupling, and bending matrices are equated with the previous ones and manipulated to yield C_0^{II} , C_1^{II} , and C_2^{II} directly in terms of the actual lamina properties and local geometry.

This approach, as we have noted,² ensures identical net mid-surface forces and moments between

any two systems for a specified displacement field. The kinematics assumed in the effective long wavelength solution are, with the exception of the through thickness shear strains, identical to an 8-node, rectangular, isoparametric brick element for small strains. Therefore, the finite element solution reflects the same behavior assumed in the homogenization. Thus, the effective properties represent precisely the varying lamina orientations, and the stacking sequence relates behavior, i.e., the bending and coupling responses.

Equation 1 allows approximate, but very accurate, element-level integration with conventional Gaussian quadrature. With $\alpha = 0.25$, the maximum normalized error in any single stiffness or force term is less than 8.4×10^{-3} . Although smaller values of α reduce the integration error, they introduce other undesirable numerical problems. By restricting individual laminae to be orthotropic, only 19 coefficients per pair of Gaussian points are necessary to describe $C^{II}(z)$, independent of the number of laminae present.

Failure Criteria

Accurate, strain-based, element-level failure criteria minimize computational costs by postponing, as long as possible, the use of expensive microlevel analysis. Criteria must be conservative to ensure that failure initiation is not missed within any of the sublaminae, and thus, a criterion is needed for each failure mechanism.

Laminate strengths are inherently limited by the extreme stresses and strains that the individual fibers and matrix can sustain. Since fibers are typically brittle, one convenient and commonly used criterion bases damage initiation or failure upon the minimum and maximum fiber direction strains.⁴ A conservative, element-level failure criterion based upon fiber direction strains was formulated. Tensile and compressive failure initiates when

$$\begin{aligned} & a \left| \max \left(0, \frac{\epsilon_{11}}{\epsilon_f^t} \right) \right|^2 + b \left| \max \left(0, \frac{\epsilon_{22}}{\epsilon_f^t} \right) \right|^2 \\ & + \left| \max \left(c_1 \frac{\epsilon_{12}}{\epsilon_f^t}, c_2 \frac{\epsilon_{12}}{\epsilon_f^t} \right) \right|^2 = \frac{4\beta - 1}{4\beta^2} \end{aligned} \quad (2)$$

and

$$\begin{aligned} & a \left| \min \left(0, \frac{\epsilon_{11}}{\epsilon_f^c} \right) \right|^2 + b \left| \min \left(0, \frac{\epsilon_{22}}{\epsilon_f^c} \right) \right|^2 \\ & + \left| \min \left(c_1 \frac{\epsilon_{12}}{\epsilon_f^c}, c_2 \frac{\epsilon_{12}}{\epsilon_f^c} \right) \right|^2 = \frac{4\beta - 1}{4\beta^2}, \end{aligned} \quad (3)$$

respectively, where

$$a = \max\{\cos^2 \theta_i\} \quad (4)$$

$$b = \max\{\sin^2 \theta_i\} \quad (5)$$

$$c_1 = \max\left\{\left[\max\left(0, \frac{2 \cos \theta_i \sin \theta_i}{\alpha}\right)\right]\right\}_i \quad (6)$$

$$c_2 = \min\left\{\left[\min\left(0, \frac{2 \cos \theta_i \sin \theta_i}{\alpha}\right)\right]\right\}_i \quad (7)$$

In these expressions, ϵ_{11} , ϵ_{22} , and ϵ_{12} are the in-plane strains, expressed in the element's natural coordinate system; ϵ_f^t and ϵ_f^c represent the composite's tensile and compressive fiber-direction strains to failure, respectively; and θ_i is the angle between the fiber direction (in the i -th lamina) and the 1-axis of the elements' natural coordinate. The value of β positions the failure surface and is bounded by $1/2 \leq \beta \leq 3/2$. For arbitrary layups, the optimal value is $\beta = 0.785$. Equations 4 through 7 are evaluated using all lamina present within the element. In 8-node brick elements, the failure criteria, Eqs. 2 and 3, need be checked on only the 'upper' and 'lower' element surfaces.

Internally Pressured Thick-Walled Cylinder

To demonstrate the new model's ability to predict the elastic response of a laminated composite material, a thick-walled composite cylinder subjected to an internal pressure was analyzed. The cylinder was axially constrained and had an inside-to-outside ratio of $r_i/r_o = 3/4$. There were 72 Gr/Ep plies randomly oriented with their fibers in either the axial (0°) or hoop (90°) direction. The 3-D quarter model used, shown in Fig. 1, contained only 60 elements, i.e., 12 circumferentially and 5 radially oriented. Radial displacements on the inner $[\delta r(r = r_i)]$ and outer $[\delta r(r = r_o)]$ surfaces were compared to a baseline solution and are listed in Table 1. The baseline solution was obtained with an axisymmetric, two-dimensional model that contained 1152 elements in the radial direction. It modeled each lamina with 16 elements. Calculations were performed with and without incompatible modes. Overall, there is excellent agreement between the axisymmetric baseline and the 3-D homogenized solutions.




Figure 1. 3-D finite element discretization of the internally pressurized thick-walled cylinder. A quarter section is represented here.

	$\delta r(r = r_i)$	$\delta r(r = r_o)$	
Baseline	1.112×10^{-4}	8.294×10^{-4}	Table 1. Baseline and 3-D calculated radial displacements. *With 'incompatible modes' turned on.
3-D	1.112×10^{-4}	8.297×10^{-4}	
3-D*	1.110×10^{-4}	8.260×10^{-4}	

Future Work

Additional failure mechanisms are necessary before component responses can be realistically tracked beyond initial failure. This requires that an element-level failure criterion as well as damage evolution relationships be developed for each mechanism. The lamina-level constitutive laws must ensure solution convergence with mesh refinement, include all mechanisms, and permit interaction between the various modes. In the immediate future, development efforts will focus on the predominant failure modes, namely, tensile and compressive failure, delamination, and in-plane shear failure. Development of the evolution laws will use both micromechanical models and non-traditional experimental results.

1. B.N. Maker, *NIKE3D User's Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-105268 (1991).
2. E. Zywicz, *Int. J. Num. Meth. Eng.* **35**, 1031 (1992).
3. N.J. Pagano, "Exact Moduli of Anisotropic Laminates," *Mechanics of Composite Materials* **2**, Academic Press (New York, New York), 23 (1974).
4. R.M. Christensen, *J. Compos. Mater.* **22**, 948 (1990)

HYDRA: A Flow Solver for Three-Dimensional, Transient, Incompressible Viscous Flow

Mark A. Christon

*Nuclear Explosives Engineering
Mechanical Engineering*

This article describes the current effort to develop a high-performance flow solver for addressing the incompressible class of complex-geometry transient flow problems that require very-high-resolution meshes. The code development effort is described in terms of the algorithm-to-architecture mapping issues involved in both vector and parallel supercomputers. An example problem showing the application of the current code to a streamline submarine hull is presented to demonstrate the class of problems being considered.

Introduction

This work is part of a collaborative effort involving the Mechanical Engineering and Physics Departments, and Military Applications at Lawrence Livermore National Laboratory (LLNL). The development of a high-performance, three-dimensional, transient, incompressible, viscous flow code is being undertaken primarily to study submarine performance in a fluid dynamics sense. The effects of flow separation and vorticity upon vehicle lift, drag, and ultimately steering, are of primary interest. The final goal is to provide a design simulation tool that will help to reduce the costly submarine design cycle.

While this effort addresses one of the National Grand Challenges of Computing, simulating flow fields about vehicles and in turbomachinery, this computational fluid dynamics (CFD) capability is unique because it also finds application within multiple divisions at LLNL, the Department of Energy, and in U.S. industry. Applications include the study of casting processes, heavy gas dispersion, and flow in the planetary boundary layer. There is also immediate application in industries critical to U.S. competitiveness in the world economy, such as the automotive industry where CFD is being used to augment engineering design in the areas of vehicle aerodynamics, heating and air conditioning, and engine and underhood cooling.

For the full-body, transient flow simulation of a submarine, it is anticipated that upwards of one million elements will be required to resolve im-

portant flow-field features such as vortices and regions of separated flow, which directly influence vehicle lift and drag. In addition to the high degree of spatial discretization, the temporal resolution is also demanding, requiring the efficient mapping of the flow-solution algorithm to current vector and parallel supercomputer architectures to make such simulations practicable.

For the solution of the time-dependent Navier-Stokes equations with complex geometry, it is estimated that computers with memory sizes of 1000 to 10,000 million words and performance rates of 10 to 1000 GFLOP's (1 GFLOP = 1 billion floating point operations per second) will be required. Today, a fully configured CRAY C-90 vector supercomputer provides a peak performance rate of 16 GFLOP's with a memory size of 256 million words. In contrast, a fully configured parallel computer such as the Thinking Machines CM-5 provides a peak performance rate of 120 GFLOP's with 4096 million 64-bit words of memory. By focusing on the rapidly evolving parallel computing platforms and making use of advanced numerical algorithms, the goal of rapid simulation of complex geometry flow simulations for design may be achievable in the near future.

Progress

The current finite element code for solving the Navier-Stokes equations is based primarily upon the work of Gresho *et al.*,¹⁻⁴ making use of ad-

vanced solution algorithms for both implicit and explicit time integration. In the case of the second-order fractional step algorithm^{3,4} (implicit), a consistent-mass predictor in conjunction with a lumped mass corrector legitimately decouples the velocity and pressure fields, thereby reducing both memory and CPU requirements relative to traditional, fully coupled solution strategies. The consistent-mass predictor retains phase speed accuracy, while the lumped mass corrector (projection) maintains a divergence-free velocity field. Both the predictor and the corrector steps are amenable to solution via direct or preconditioned iterative techniques, making it possible to tune the algorithm to the computing platform, i.e., parallel, vector, or super-scalar. The second-order projection algorithm can accurately track shed vortices, and is amenable to the incorporation of either simple or complex (multi-equation) turbulence submodels appropriate for the driving applications.

The explicit solution algorithm^{1,2} sacrifices some of the phase-speed accuracy of the fractional-step algorithm for the sake of minimizing memory and CPU requirements. However, the momentum equations are still decoupled in the explicit algorithm. While both the diffusive and Courant stability limits must be respected in the explicit algorithm, balancing tensor diffusivity^{1,2} is used to lessen the restrictive diffusive stability limit in the explicit algorithm. This, in combination with single-point integration and hour-glass stabilization, makes the explicit algorithm very efficient computationally, and because of this, the explicit algorithm was chosen as the initial focus of the parallelization effort.

The fractional step and the explicit algorithms both rely upon the implicit solution of a linear system arising from an elliptic operator. In the case of the fractional step algorithm, this solution is used to project an intermediate velocity field to a divergence-free space. In the explicit solution strategy, the elliptic operator appears in the pressure Poisson equation, which is used to advance the pressure field in time. Because the linear system solver is a key component of the algorithm-to-architecture mapping for both algorithms, it has been necessary to develop modified, conjugate-gradient iterative solvers^{5,6} that minimize the impact on memory requirements and allow the natural data parallelism of element-level processing to be exploited. For both the explicit and the projection algorithms, no additional storage is required for the elliptic operator itself, making the current conjugate gradient solver essentially matrix-free.

During the past year, our efforts have been directed primarily towards the vector and data-parallel or SIMD (Single Instruction Multiple Data) implemen-

tation of the code suitable for the laminar flow regime. In the case of the vectorized version of the code, element operations are blocked into groups of contiguous, data-independent operations by using a simplified domain-decomposition strategy to group the elements. This approach results in a code that is completely vectorized, yielding performance comparable to DYNA3D for the time integration of the momentum equations. However, the solution to the pressure Poisson operator limits the overall performance of the code, taking up to 95% of the CPU time per simulation time-step in problems with strong pressure-velocity coupling.

Because the element data structures for the vectorized version of HYDRA are adjustable, they are also used for the SIMD (CM-2/CM-5) implementation, where element-level operations are performed in a lock-step parallel fashion. For the CRAY architecture, the vector block size is configured as 128 (twice the length of the vector pipeline). In the case of the CM-2/CM-200, the element block size is configured as a multiple of the minimum virtual processor ratio (4) and the number of available processing elements. For the CM-5, the block size is configured as an integral multiple of the processor vector pipeline length and the number of available processors enabling processor pipelined operations in conjunction with SIMD parallelism.

In the SIMD (CM-2/CM-5) version of the code, data dependence in the element blocks may be resolved using hardware-specific communication/combining operations for the parallelized assembly of element data to nodal data. Instead of data dependency, the constraint on domain decomposition in the SIMD implementation requires that the elements be grouped in a spatially contiguous manner to minimize the deleterious effects on performance of off-processor communication. However, because the same data structures are used for the vector and SIMD versions of HYDRA, it is possible to appropriately reconfigure the block size for each architecture for the sake of performance. In effect, the element grouping strategy provides a mechanism to account for variations in granularity across supercomputer architectures ranging from vector to SIMD to Multiple Instruction Multiple Data (MIMD).

Many alternative domain-decomposition algorithms are available, including methods that consider the graph of the finite element mesh⁷ when subdividing the physical domain, and are not restricted to logically regular meshes. By matching the domain decomposition strategy to the supercomputer architecture, it will be possible to maintain optimal performance on vector, SIMD, and MIMD machines.

HYDRA has been written using standard, UNIX

software-development tools, enabling the code to be simultaneously developed in FORTRAN 77 and FORTRAN 90 by making use of compile time configuration of the software. This approach has made it possible to provide HYDRA in a form suitable for computing platforms including workstations, and CRAY vector and Thinking Machines SIMD supercomputers. The top-down design and bottom-up implementation have required the design of a memory management package that makes it possible to perform dynamic memory allocation on a single processor workstation, multi-processor CRAY, or on the processors of the CM-2/CM-5 with a single interface definition.

Application

Currently, initial calculations are being performed on a range of simplified submarine-hull configurations. The top frame in Fig. 1 illustrates the mesh used in the computation of the flow field around a streamline submarine hull at a Reynolds number of 830, based on the hull diameter. A 1/4 symmetry model has been used, resulting in a mesh with 18,000 nodes (16,000 elements) or 72,000 degrees of freedom. Tow-tank conditions were imposed on the computational domain to simulate the case when the vehicle is moving straight ahead.

Isosurfaces of pressure are shown in the middle frame of Fig. 1 for the initially divergence-free and potential flow field. At the leading edge of the vehicle, a stagnation point is apparent, with the pressure decreasing in the streamwise direction along the hull of the vehicle. Near the trailing edge of the submarine, a low pressure 'bubble' is present due to the acceleration of the fluid as it tries to turn and follow the streamline surface of the hull. In the bottom frame of Fig. 1, isosurfaces of the x-velocity are shown. At the inlet to the domain and the far-field boundary, a uniform x-velocity has been imposed to simulate tow-tank conditions. The bubbles at the front and back of the straight section of the hull correspond to locations where the fluid has been accelerated to track the contours of the vehicle hull. At the surface of the hull, no-slip boundary conditions have been imposed.

Future Work

Future efforts will address two key issues: the acceleration of the solution to the linear system, arising from the pressure Poisson operator; and the inclusion of the recursive, spectral domain-decomposition strategy for SIMD architectures. The pressure computation currently relies upon a data-parallel, element-by-element diagonally scaled, matrix-free conjugate gradient solver. While this approach offers




Figure 1. Results of calculations performed on simplified submarine-hull configurations, showing (a) the mesh used in the computation of the flow field; (b) isosurfaces of pressure; and (c) isosurfaces of the x-velocity.

scalability in terms of the mesh resolution, multigrid acceleration can provide enhanced convergence rates by effectively damping the low-mode error components via coarse grid corrections. It also fits well in the current parallel-code architecture.

While the current, vector-blocking, domain-decomposition algorithm is adequate for vector supercomputers, robust decomposition techniques yielding element-to-processor assignments that minimize communications overhead are necessary to achieve peak performance rates on both SIMD and MIMD parallel architectures. The recursive spectral bisection⁷ algorithm, which uses the second eigenvector of the mesh connectivity graph, is currently being investigated as a candidate for performing domain decomposition.

1. P.M. Gresho, S.T. Chan, R.L. Lee, and C.D. Upson, *Int. J. Numer. Methods Fluids* **4**, 557 (1984).
2. P.M. Gresho, S.T. Chan, R.L. Lee, and C.D. Upson, *Int. J. Numer. Methods Fluids* **4**, 619 (1984).
3. P.M. Gresho, *Int. J. Numer. Methods Fluids* **11**, 587 (1990).
4. P.M. Gresho and S.T. Chan, *Int. J. Numer. Methods Fluids* **11**, 587 (1990).
5. E.S. Beckman, "The Solution of Linear Equations by the Conjugate Gradient Method," *Mathematical Methods for Digital Computers* **1**, 62 (1965).
6. R.M. Firenzi, *Element-by-Element Preconditioning Techniques for Large-Scale, Vectorized Finite Element Analysis in Nonlinear Solid and Structural Mechanics*, Ph.D. Thesis, Stanford University, Palo Alto, California (1989).
7. H.D. Simon, "Partitioning of Unstructured Problems for Parallel Processing," *Computing Systems in Engineering* **2** (2), 135 (1991).

Development and Testing of the TRIM3D Radiation Heat Transfer Code

James D. Maltby

*Nuclear Test Engineering
Mechanical Engineering*

We have developed a new code, TRIM3D, to solve radiative heat transfer problems involving a participating medium. The code uses a Monte Carlo formulation to solve problems with absorption, anisotropic scattering, and specular boundaries. It is designed to work with other codes to solve coupled radiation/conduction/thermal stress problems, and has been verified against known analytic solutions.

Introduction

Radiation heat transfer problems involving a semi-transparent medium that participates in the radiative exchange occur often in areas such as high-power optics, crystal growth and glass manufacture, coal furnaces, annealing ovens, and analysis of fuel fires. Unfortunately, these problems are often difficult to solve due to the complex nature of the radiative transfer equation. A Monte Carlo approach to radiation heat transfer problems without a participating medium has proved very successful, resulting in the computer code MONT3D.^{1,2} The objective of this research was to develop a Monte Carlo code to analyze radiation heat transfer in the presence of a participating medium.

The resulting code, TRIM3D, represents the state of the art for radiation heat transfer analysis, and is also the first production code with detailed participating medium capability. The addition of TRIM3D to our code suite allows the solution of coupled radiation/conduction/thermal stress problems with a level of detail not previously attainable.

Progress

During FY-92, a working version of the computer code TRIM3D was developed and given preliminary testing and verification. The TRIM3D code is formulated in a similar manner to the successful MONT3D non-participating medium heat transfer code used by programs at

Lawrence Livermore National Laboratory (LLNL). The current working version of TRIM3D solves three-dimensional (3-D) radiation heat transfer problems in absorbing, emitting, and anisotropically scattering media. Problems may be solved that are non-homogenous and non-isothermal, and material properties may vary with wavelength. Boundaries may be diffuse, specular, or mixed, with directional reflectivity and transmissivity.

The code has been verified against a series of analytic problems with absorbing or scattering media and specular boundaries, with agreement within the statistical accuracy of the simulation. Currently, no other code exists that can handle participating media problems of this complexity.

Theoretical Formulation

TRIM3D generates a matrix of direct exchange areas (DEA's) that describe the radiative interaction among all surfaces and volumes in an en-
clo-




Figure 1.
TRIM3D code flow. Temperature output from TOPAZ3D is passed through REMAP to NIKE3D for solution of radiation/conduction/thermal stress problems.

sure. The net exchange between any two surfaces or volumes may be described:

$$\begin{aligned} Q_{ij} &= \sigma s_i s_j (T_i^4 - T_j^4) && \text{(surface to surface)} \\ &= \sigma s_i g_j (T_i^4 - T_j^4) && \text{(surface to volume)} \\ &= \sigma g_i g_j (T_i^4 - T_j^4) && \text{(volume to volume).} \end{aligned}$$

This matrix is then passed to TOPAZ3D for solution of the coupled radiation/conduction problem. Since the matrix of DEA's is temperature-independent, boundary conditions and temperatures in an analysis may be changed without re-running TRIM3D. This approach has been very successful with MONT3D and TOPAZ3D, resulting in a large savings in computer time. The temperature output from TOPAZ3D may then be passed through REMAP to NIKE3D for solution of radiation/conduction/thermal stress problems. The code flow during the solution of such a problem is shown in Fig. 1. The mesh generator INGRID and post-processor TAURUS are also used.

TRIM3D simulates thermal radiation by emitting a large number of monoenergetic photons from each surface and volume. These photons are traced through multiple reflection and/or scattering events until they are absorbed in another surface or volume. The DEA's are then calculated from these photon tallies. For a given row of the DEA matrix,

$$\begin{aligned} s_i s_j &= A_i F_{ij} N_{ij} / N_i \\ g_i s_j &= 4 V_i a_i N_{ij} / N_i \end{aligned}$$

where N_{ij} is the number of photons emitted by element i and absorbed by element j , and N_i is the number emitted by element i .

Figure 2. Analytic verification of TRIM3D for pure absorption.





Figure 3. Analytic verification of TRIM3D for isotropic scattering.

If the material properties change significantly with wavelength, as is typical with gas-radiation problems, a band-wavelength model is available. This model splits the wavelength range into separate gray bands, with a separate simulation per band.

Surfaces consist of 4-node shell elements, degenerating to triangles. Volumes are represented as 8-node bricks, with triangular prisms and tetrahedra as subsets. Both surfaces and volumes are designed for mesh compatibility with INGRID and TOPAZ3D. Material properties are assumed constant within a single element, but any number of materials may be defined. In this manner, non-homogenous problems may be solved.

Analytic Verification

TRIM3D has been verified against a series of participating medium heat transfer problems with known analytic solutions. Though the analytic solutions are one-dimensional, they were simulated with a 3-D geometry with specular mirrors on four sides. Optical depths from 0.1 to 10 were tested, with good agreement throughout the entire range. Some of the results of this verification are shown in Figs. 2 and 3 for pure absorption and isotropic scattering, respectively. Agreement with the analytic solution in both cases is very good.

An additional result from this verification was that the speed of the code appears to decrease only linearly with optical depth, and that even at an 'optically thick' depth of 10, the speed is practical on a SUN workstation.

Future Work

One of the difficulties in verifying a participating medium code is the small number of problems with analytic solutions that exercise all the code features. To address this problem, a symposium was held at the 1992 American Society of Mechanical Engineers Heat Transfer Conference to assess the current capability for solving non-gray, anisotropically scattering, multidimensional radiation problems. Thirty-four benchmark problems ranging from one to three dimensions at optical depths from 0.1 to 10 were specified.³ These problems will be solved using TRIM3D and should provide a good platform for verification of the code features.

Additional features are planned for the production version of the code to simplify the solution of large problems and make the code more 'user friendly.' A complete user's manual, including test problems, will be produced for TRIM3D. In addition, all the solved analytic and benchmark problems will be organized into a quality assurance manual for code validation purposes.

Once the code is released, we intend to collaborate with groups inside or outside LLNL to test the utility and accuracy of TRIM3D on experimental problems. This will provide valuable feedback on code robustness and performance on large problems, as well as on which features are most useful to the analysis community.

Because TRIM3D uses a Monte Carlo formulation, it is very well suited to the new class of massively parallel computers. A test version of TRIM3D will be developed for whichever parallel computer becomes available at LLNL, and its performance will be assessed. If successful, it should provide a good example of a production-parallel application.

1. J.D. Maltby and P.J. Burns, *Numer. Heat Transfer* **9** (2), (1991).
2. R. Siegel and J.R. Howell, *Thermal Radiation Heat Transfer*, 4th ed., Hemisphere Publishing Corporation (Bristol, Pennsylvania), 1992.
3. T.W. Tong and R.D. Skyocypec, "Summary on Comparison of Radiative Heat Transfer Solutions for a Specified Problem," *Developments in Radiative Heat Transfer*, HTD **203** (New York), 1992. 

A Methodology for Calculating the Seismic Response of Critical Structures

David B. McCallen

*Nuclear Test Engineering
Mechanical Engineering*

**Francois E. Heuze,
Lawrence J. Hutchings, and
Stephen P. Jarpe**

Earth Sciences Department

We are developing a methodology chain that will allow estimation of the seismic response of critical structures to large earthquakes. The methodology consists of three distinct steps: generation of synthetic bedrock motion at the structure site due to a postulated large earthquake; nonlinear finite element analysis of the soil profile at the site to transform the bedrock motion to surface motion; and linear/nonlinear finite element analysis of the structure based on the predicted surface motions. Progress in all steps is reported here. Our ultimate goal is to allow accurate, site-specific estimates of structural response for a specified earthquake on a specified fault.


Introduction

Our computational simulation of the seismic response of a critical structure is illustrated in Fig. 1. To envelope the motions that might be observed at the structure site, the seismological portion of the methodology develops a suite of possible earthquake rupture scenarios for each fault that can contribute significant ground motion at the site. Field instrumentation is placed on bedrock at the structure site, and over a period of time, bedrock motions due to micro-earthquakes emanating from the causative fault(s) are recorded. These recordings serve as empirical Green's functions, which characterize the motion at the structure site location due to slip of an elemental segment of the fault. By appropriate summation of the responses due to each element of the fault rupture zone for a given rupture scenario, the bedrock motion due to slip over a large area of the fault (corresponding to a large magnitude earthquake) can be estimated. By considering a standard suite of 25 possible fault rupture models, which characterize the different manners in which the fault rupture can propagate across the total fault rupture zone, a suite of 25 acceleration time histories are generated. The suite of time histories is representative of the maximum ground accelerations that could be expected at the site for a given size earthquake. Hutchings^{1,2,3} and

his coworkers have led the development of the empirical Green's function technique and demonstrated the utility of this method using Loma Prieta earthquake data.²

The transmission of earthquake motion from bedrock through the soil to the soil surface can result in significant modification of the bedrock motion. Traditionally, the nonlinear behavior of the soil under strong seismic motion has been modeled with 'equivalent linear' methods, which iterate with a linear model to approximate the nonlinear response of the soil deposits. The classical computer program SHAKE⁴ has typically been used to perform site-response analysis. SHAKE is operational at Lawrence Livermore National Laboratory (LLNL), but such equivalent linear models cannot describe the evolution of pore pressure and predict liquefaction; i.e. they cannot perform "effective-stress" analysis which we deemed essential for this project. So, the effective stress nonlinear finite element program DYNAFLOW⁵ has been obtained from Princeton University. As part of the methodology development and validation, the DYNAFLOW and SHAKE programs will be applied to the Loma Prieta earthquake data obtained at Treasure Island, California. The Treasure Island site consists of saturated soils that exhibited liquefaction during the Loma Prieta earthquake. Site-response calculations are being performed by

Figure 1. Computational simulation of the seismic response of a critical structure showing (a) the physical system and (b) the three-step computational model.



a number of researchers, and a portion of our model validation efforts will consist of a comparison of DYNATLOW and SHAKE results with measured Treasure Island response data for the Loma Prieta earthquake.


Nonlinear structural-response computations are being performed with nonlinear finite element software developed at LLNL. The implicit, nonlinear, finite deformation program NIK3D is being used to model structures and the nonlinear near-field soil. NIK3D has a number of nonlinear constitutive models and advanced contact-surface capabilities for modeling gap opening and closing.

The seismic analysis procedures and capabilities under development are being applied to two

transportation structures in California. The first structure is the Dumbarton Bridge, which is the southern-most crossing of the San Francisco Bay. The Dumbarton Bridge Study was initiated by LLNL at the request of the California Department of Transportation (CDOT). The second study is concerned with the seismic analysis of the Painter Street Bridge in Rio Dell, California. The Painter Street Bridge study is very important from the standpoint of validation of our methodology and procedures. This study is the focus of this report.

The Painter Street Bridge, which has been heavily instrumented by the California Department of Mines and Geology (CDMG), provides an excellent case study. The high rate of occurrence of

Figure 2. Location of April 1992 Petrolia earthquake epicenters and Painter Street Bridge site.






Figure 3. Photographs showing (a, b, and c) experimentation and field work for the Painter Street Bridge site; and (d) illustration of finite element model.

earthquakes in the seismically active region of northern California has allowed the measurement of response of this bridge to a number of significant earthquakes. In April 1992, three large earthquakes occurred within close proximity to Rio Dell and the Painter Street Bridge location (see Fig. 2). During the largest of these shocks, the Painter Street Bridge structure was shaken quite violently, with lateral deck accelerations on the order of 1.23 times the acceleration due to gravity. These measured accelerations represent the largest accelerations ever measured in any structure during an earthquake. Prior to the April earthquakes, McCullen, Romstad, and Goudreau had constructed a detailed finite element model of the Painter Street bridge-abutment system (see Fig. 3) and had performed detailed parameter studies on the dynamic response of the system. Since an extensive modeling effort had already been initiated on this bridge, the latest set of quakes was a fortuitous event for our project.

Progress

As a result of the April quakes, the LBNL efforts at the Painter Street site have been scaled up signif-




Figure 4. Region of aftershock locations used for measurement of empirical Green's functions.




Figure 5. Five sample fault rupture scenarios with resulting Painter Street time histories.

icantly. At the request of Heuze,⁸ the CDOE recently drilled four bore holes at the Painter Street site (see Fig. 3) and performed down-hole, shear-wave-velocity measurements. Soil samples were retrieved from the boreholes, and Heuze has contracted with the Department of Civil Engineering at (UCB) to perform laboratory tests on the soil samples. The field shear-wave-velocity measurements and the laboratory soil testing will provide quantitative soil properties for use in the site soil response calculations and the structural model calculations. Two of the boreholes were drilled to bedrock (a depth of approximately 80 ft), and two seismic instrumentation packages were placed, one at the surface and the second at the bedrock depth of 80 ft (Fig. 3). The package at bedrock depth is currently being used by seismologists to measure empirical Green's functions for micro-earthquakes emanating from nearby faults.

To date, Painter Street site bedrock responses have

been measured for eight micro-earthquakes emanating from the fault locations indicated in Fig. 4. Based on the empirical Green's functions obtained from these measurements, synthetic bedrock-ground-motion time histories have recently been generated by Hutchings and Jarpe for a number of earthquakes. Samples of Painter Street site synthetic time histories, each based on a different fault rupture propagation model, are shown in Fig. 5.

In parallel to the seismological work, finite element modeling of the Painter Street bridge/abutment system has progressed into the nonlinear regime. For nonlinear time history analyses, the superstructure, pile foundation, and approach embankment soil masses have been modeled as shown in Fig. 3. In this type of bridge structure, nonlinear hysteretic behavior of the soil embankments has been experimentally identified as a very important factor in the dynamic response of the bridge system.^{9,10} The primary objective of constructing a detailed, three-dimensional model of the bridge/soil system was to allow incorporation of the effects of nonlinear soil stiffness and soil mass. Traditional finite element models for this type of bridge, which are used in bridge design and analysis calculations, neglect the soil mass, and the soil stiffness is represented by linear elastic, amplitude-independent springs. We decided to truncate the detailed finite element model at approximately the original ground surface elevation, and apply the surface free field motion directly to the base of the model at this elevation (see Fig. 3). This approach neglects potential soil-structure interaction effects between the piles and soil below this level, and prevents radiation of energy vertically back into the soil. However, interaction between the soil and piles typically occurs in the top portion of the piles, and energy loss through radiation will be small relative to the energy dissipated by the nonlinear hysteretic behavior of the soil embankments.

Until the experimental tests are completed at UCB in January 1993, there is only sparse quantitative data on the soil properties for the Painter Street site. The small-amplitude shear moduli for the approach embankments and original grade soils have been estimated¹¹ based on P and S wave surface refraction measurements which were performed. To represent the nonlinearity of the soil in the bridge/abutment finite element model, the small-strain shear moduli obtained from these measurements were used with standard soil modulus degradation and damping curves.¹² To represent the standardized modulus degradation and damping curves in the NIKI3D finite element program,

a simple Ramberg–Osgood constitutive model was used to model the soil. The material parameters were set such that the Ramberg–Osgood hysteresis loop would yield modulus degradation and damping curves very similar to Seed's¹² standardized curves. The procedure for determining the Ramberg–Osgood parameters to approximate given modulus degradation and damping curves was developed by Ueng and Chen.¹³ The modulus and damping curves obtained from the Ramberg–Osgood constitutive model fit with Ueng and Chen's technique are shown in Fig. 6 along with the original curves of Seed. The shear stress-strain behavior generated with the fitted Ramberg–Osgood model in the NIKE3D finite element program is also shown in Fig. 6.

A number of time history analyses have been carried out with the detailed bridge/abutment model shown in Fig. 3, as well as with simple reduced-order stick models of the bridge.⁴ The bridge instrumentation records for the April 1992 Petrolia earthquakes have not yet been completely processed by the CDMG; thus, the measured free field motions were not available to apply to our model prior to this report. However, free field and bridge-response data for a magnitude 5.5 earthquake of November 1986 were available and were used to examine the accuracy of the finite element models of the bridge system.

The 1986 free-field acceleration time histories were used as input motion to the base of the bridge system models. The model response predictions were compared to the actual bridge response data measured by the CDMG bridge instrumentation array. Since the details of all of the response predictions are given elsewhere,¹⁴ only an illustrative example of the response predictions is provided here. The detailed model response predictions for the absolute displacement at channel 7 (transverse motion at mid-span) are shown in Fig. 7. Figure 7a shows the response of the detailed model when a linear elastic soil model is used, and mass- and stiffness-proportional Rayleigh damping is used to provide approximately 5% damping in the first transverse and longitudinal modes of the bridge system. For the linear analysis, soil properties were set equal to the small-strain soil properties estimated by Heuze and Swift from field measurements. Two observations can be made: (1) the frequency content of the bridge model is significantly too high when the small-strain soil properties are used; and (2) the amplitude of the response prediction is too large relative to the measured response. The bridge response prediction using the detailed model with the nonlinear Ramberg–Osgood soil model





Figure 6. Simple nonlinear soil characterization for finite element model.

Figure 7. Response predictions with
(a) linear and
(b) nonlinear models.



is shown in Fig. 7b. This model also used mass-proportional Rayleigh damping, in which the damping in the first transverse mode was set to 10%. The nonlinear model exhibits significant improvement over the linear model. The nonlinear model displays appropriate softening and energy dissipation in the system, such that the frequency content and amplitude are more representative of the actual structural response.

Future Work

Significant progress has been made in the study of the Painter Street overcrossing site. Construction of the seismological model and the structural model have been completed, and calculations have been generated with both models. Additional field measurements of Green's functions from future micro-earthquakes will continue to enhance the site seismological model, and laboratory experimental data will improve the soil characterization in the finite element model of the bridge/abutments. The site-soil characterization will also allow site-response

analysis to transform bedrock motion to soil surface motion. Specific tasks that we intend to perform during the next year include:

- (1) Use of the nonlinear model of the bridge/abutment system to predict the response of the bridge to the April 1992 Petrolia earthquake. The predicted response will be compared to the actual bridge response measured by CDMG. This earthquake should have resulted in significant nonlinear behavior of the bridge/abutment system, and this analysis will allow us to further verify the ability of the nonlinear model to accurately predict bridge/abutment response.
- (2) Based on measured Green's functions, the seismological model will generate a final suite of 25 time histories for the April 1992 Petrolia magnitude-7 earthquake.
- (3) The bedrock-motion time histories will be transformed to surface motion with a site-response analysis, and the suite of surface time histories will be compared to the actual free field motion measured at the site by CDMG.

- (4) The suite of predicted free field responses will be run through the structural model, and response statistics will be compared to the actual response from the April 1992 magnitude-7 earthquake.

The ultimate goal of our project is to allow accurate site-specific estimates of structural response for a specified earthquake on a specified fault. For practical applications of this methodology, it will be essential to decide how the structural engineer may best use the information provided by the suite of time histories developed by the seismological portion of the study. It will generally be impractical to perform 25 time history analyses (or more if multiple faults/multiple rupture zones are considered) for a large structural model. It is necessary to consolidate the information obtained from the time histories into a simplified form (e.g., a representative response spectrum and corresponding single time history) to achieve practical application.

The Painter Street site study will allow a critical evaluation of the accuracy of the method that is being developed, and a demonstration of our technology in all segments of the methodology chain. It will also provide an opportunity for interaction between structural analysts and seismologists, so that appropriate procedures for using the earthquake ground motion in structural response calculations can be developed.

1. L. Hutchings, *Modeling Earthquake Ground Motion with an Earthquake Simulation Program (EMPSYN) That Utilizes Empirical Green's Functions*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-ID-105890 (1992).
2. L. Hutchings, *Bull. Seism. Soc. Am.* **81** (5), (1991).
3. L. Hutchings, *J. Geophys. Res.* **95** (B2), (1990).
4. P.B. Schnabel, H.B. Seed, and J. Lysmer, *SHAKE - A Computer Program for Earthquake Response Analysis of Horizontally Layered Sites*, Earthquake Engineering Research Center, University of California, Berkeley, California, Report EERC-72-12 (1972).
5. J.H. Prevost, *DYNALOW, User's Manual*, Department of Civil Engineering, Princeton University (1992).
6. B.N. Maker, R.M. Ferencz, and J.O. Hallquist, *NIKE3D: A Nonlinear, Implicit, Three-Dimensional Finite Element Code for Solid and Structural Mechanics, User Manual*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-MA-105268 (1991).
7. D.B. McCallen, K.M. Romstad, and G.L. Goudreau, "Dynamic Response of a Reinforced Concrete Box-Girder Bridge," *Engineering Research, Development, and Technology*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-53868-91, 2-12 (1992).
8. Personal communication between E. Heuze (LLNL) and Kenneth Cole (CTXT) (1992).
9. J.C. Wilson, and B.S. Tan, *ASCE J. Eng. Mech. Div.* **226** (8), (1990).
10. S.D. Werner, J.L. Beck, and M.B. Levine, *Earthquake Eng. and Struct. Dyn.* **15**, (1987).
11. E.E. Heuze and R.P. Swift, *Seismic Refraction Studies at the Painter Street Bridge Site, Rio Dell, California*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-ID-108595 (1992).
12. H.B. Seed, R.T. Wong, I.M. Idriss, and K. Tokimatsu, *Moduli and Damping Factors for Dynamic Analysis of Cohesionless Soils*, Earthquake Engineering Research Center, University of California, Berkeley, California, Report EERC-84/14 (1984).
13. T.S. Ueng, and J.C. Chen, *Computational Procedure for Determining Parameters in Ramberg-Osgood Elastoplastic Model Based on Modulus and Damping Versus Strain*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-ID-111487 (1992).
14. D.B. McCallen and K.M. Romstad, *Dynamic Response of a Reinforced Concrete Box-Girder Bridge*, Lawrence Livermore National Laboratory, Livermore, California, UCRL-ID-110640 (1992). □

Reinforced Concrete Damage Modeling

**Sanjay Govindjee and
Gregory J. Kay**

*Nuclear Explosives Engineering
Mechanical Engineering*

The modeling of reinforced concrete structures is currently performed by empirical codified formulae and linear elastic calculations. This state of the practice, however, can lead to both non-conservative designs on the one hand and to over-designed and costly structures on the other. This wide range of outcomes arises from the lack of an adequate constitutive model to describe the behavior of concrete as it cracks under applied loads. This report briefly describes work at Lawrence Livermore National Laboratory in the development of an appropriate constitutive model for concrete damage.

Introduction

In the modeling of reinforced concrete structures, the current state of the practice involves the use of codified empirical formulae and linear calculations. While these methods are very useful, they can also produce unwanted results. When using empirical formulae, there is risk involved in applying them to a situation that is not absolutely identical to the tests from which they were deduced. In particular, formulae for limit loads scale in a rather non-linear fashion and must be applied with care and experience to avoid a non-conservative design. On the other hand, one does not want to have to over-build a structure and hence make it overly costly because of uncertainties in modeling.

A vast improvement to the design cycle is obtained if some of the empirical formulae currently used are replaced by analytical models. The main unknown that most of the empirical formulae try to address involves the behavior of the concrete itself as it cracks under various loading conditions with different reinforcement patterns. Thus, the thrust of our work has been to develop a constitutive model that describes the behavior of damaging concrete. Because this work is being performed for the Computational Earthquake Initiative at Lawrence Livermore National Laboratory (LLNL), the level of complexity of the model has been chosen to be commensurate with that needed to model critical sections of large reinforced concrete structures under seismic loading conditions. This requires the constitutive model to be able to track

the progression of damage induced by arbitrary three-dimensional (3-D) loading histories in complex 3-D geometries. Because of these requirements, the model has been developed as a 3-D damage theory that is suitable for large-scale finite element calculations.

Such thinking is not new to the modeling of reinforced concrete structures.¹ This original work, and almost all that has followed since, has been confined to two-dimensional (2-D) problems. Under seismic excitations, however, one must look at the more general situation that includes 3-D effects, because of the high likelihood of complex loading paths. There does exist a handful of 3-D models.^{2,3,4} However, none of these models is suitable for the present problem. The first two models and others like them are only suitable for isotropic compressive type behavior, and the third, while promising, still requires some development. The present model takes advantage of the insights and developments of this previous work and extends them to a new framework for damage modeling. The framework we have developed most closely resembles the framework proposed by Ortiz.⁴

Progress

Progress for FY-92 has been made on many different aspects of the problem: choosing an appropriate class within which to develop the model; developing the features to incorporate into the model; developing appropriate numerical algorithms to efficiently perform finite element calcu-

lations; and determining how reinforcing bars should be modeled in conjunction with the cracking concrete.

Model Class and Features

Model class refers to the basic style of the model: plasticity-like or damage-like. In plasticity-like models, material unloading takes place elastically with a stiffness equal to the initial elastic stiffness of the material (**Fig. 1a**). In a damage-like model, material unloading takes place elastically with a degraded stiffness (**Fig. 1b**). The plasticity-like models have strong appeal for a number of reasons, but mainly because their algorithmic properties are reasonably well understood and are known to be suitable for finite element calculations. The true behavior of cracking concrete, however, resembles more closely damage-like model behavior.

Nevertheless, at the beginning of this project, we used a plasticity-like model to examine some of the numerical and theoretical issues that are unique to materials displaying softening behavior like that shown in **Fig. 1**. The main use of this model class was to examine the issue of ill-posed boundary-value problems. Materials displaying

stress-strain behavior like that in **Fig. 1** often generate ill-posed boundary-value problems.⁵ While there are several ways around this issue, for concrete the most physically realistic one is the notion of constraining the amount of energy dissipated in the system on a per-unit-volume basis to equal that dissipated on a per-unit-area basis when opening new crack faces. This type of constraint results in the appearance of a characteristic length in the model formulation. For the development of the present model, the continuum formulation⁶ was used to render the present formulation well-posed for both the plasticity and damage model classes.


In the domain of damage models, there is a wide variety of model choices. To choose the appropriate one usually requires a fair amount of insight into the micromechanical mechanisms of the observed damage and their relationship to the free energy density of the material. In the case of concrete with Mode I, II-, and III-type cracks, such information is not available. Therefore, several general hypotheses of continuum mechanics have been used instead to generate a complete model.

The basic premise of the model is that the damage state of the material will be represented by the rank 4 stiffness tensor of the material. Hence, as is known to occur in other damaging systems,⁷ the 'elastic stiffness' of the material is allowed to evolve with the loading history. To determine the evolution law for this degraded stiffness, the notion of maximum dissipation is used. To use this idea, one first postulates restrictions on the admissible stress or strain states of the material. For the concrete, two restrictions are postulated. The first restriction states that the normal tractions across cracks in the system must be below a given critical value and that the critical value evolves as the damage increases. The second restriction states that the shear tractions across cracks in the system must be below a given critical value, which also evolves with progressing damage. Cracks are assumed to nucleate in the material when the maximum principal stress at a point exceeds a given value. Using these two restrictions and the concept of maximum dissipation, an evolution law can be derived for the rank 4 stiffness tensor of the material that gives the correct anisotropic structure to the damaged stiffness tensor.

The other dominant phenomenological features of cracking concrete that have been incorporated into the model are:

- (I) The choice of restrictions on the admissible stress states in the material provides

Figure 1. Material unloading in (a) plasticity-like and (b) damage-like model classes.



for Mode I-, II-, and III-type crack growth (damage evolution).

- (2) The notion of crack closure has been included by monitoring the tractions across crack faces. When the traction across a crack face becomes compressive (negative) and the shear tractions are below their critical value, the material behaves as though it is undamaged (up to the compressive yield limit of the concrete).
- (3) The notion of shear retention is built into the model by limiting the amount of shear degradation allowed in the system.
- (4) The softening evolves with an exponential character.
- (5) The damage evolution is anisotropic.

Algorithms

The algorithmic implementation of the proposed model in a finite element setting has involved the development of several novel algorithms. Of foremost importance for softening models has been the development of a characteristic-length interpolation scheme for 3-D problems. While an interpolation scheme for 2-D problems has been presented,⁶ a straightforward extension of this method to 3-D leads to singular characteristic lengths for certain crack orientations. In our work, a new interpolation method has been developed that does not have these singularities and yet remains faithful to the original characteristic-length idea.

The other algorithmic issues that have been addressed deal with local and global integration algorithms. On the local level, a concave (as opposed to convex, as in metal plasticity) optimization problem governs the stress point calculation. Because of the concave nature of the problem, a unique answer to the stress point calculation does not exist; there are two answers, with one being inadmissible. However, by picking a suitable starting value, the stress point algorithm can be made to always produce the admissible answer. On the global level, the non-linear balance equations of the boundary-value problem have multiple bifurcation paths that lie extremely close to each other and cause global convergence difficulties. To circumvent these well-known convergence difficulties, an aggressive, automatic time-stepping scheme has been developed. The scheme uses logarithmic-based time step control in conjunction with a special oscillating norm check. The combination of these two ideas greatly enhances the ability of the global solvers to achieve equilibrium.

Reinforcing Bars

Since using fixed rebar bars (i.e., compatible displacements between concrete and rebar) gives reasonable results, to date only a small effort has been devoted to rebar issues. Our results are, however, slightly non-conservative. To address this, some preliminary work has been done on rebar release methods. Force- and damage-based slide-line release methods have been used, as have bond-link elements. The damage-based slideline release has been found to be superior to the force-based model and the bond-link element for accuracy against experimental data. However, the best overall robustness for these methods (after the fixed rebar model) is given by the bond-link element, which is a node-on-node contact element with a displacement-based release law.

Examples

Two examples are shown to partially demonstrate the proposed model. The first example involves the 3-point bending of a lightly reinforced beam; the second example involves the 3-point bending of a heavily reinforced beam.

In the first example, the beam is 12 feet long with a 8 × 20 in. cross section that contains two #8 rebars in the lower fibers. The load deflection curve at mid-span is shown in Fig. 2. Overall agreement is seen to be quite good. At point (A), the concrete starts to crack, and load is transferred into the rebars. Cracking progresses up through the cross section with more load being transferred into the rebars until at point (B) the rebars yield. These observations from the simulation are consistent with experimental observations.⁸





Figure 2. Load deflection curves at mid-span for beam with two #8 rebars in the lower fibers. The damage initiation point (A) and the point of yield (B) are marked.

Figure 3. Load deflection curves at mid-span for beam with four #9 rebars in the lower fibers and two #4 rebars in the upper fibers. The damage initiation point (A) and the rebar-concrete interface failure initiation point (B) are marked.



In the second example, the beam is 12 feet long with a 12×21.75 in. cross section that contains four #9 rebars in the lower fibers of the beam and two #4 bars in the upper fibers of the beam. Additionally, there are #2 stirrups every 8.25 in. along the length of the beam. **Figure 3** shows the load deflection curve at mid-span for the experiment⁹ and the calculation. At point (A), the concrete starts to crack, and there is a large load transfer to the #9 rebars. The #4 rebars do not carry much of the load. Vertical cracks develop along the span and grow upwards and towards the centerline of the beam. At point (B), the calculation deviates from the data because rebar release was not included in the simulation.

Future Work

Future work will focus on making the local stress point algorithm more robust and efficient. In addition, a few new features will be added, such as compressive flow of the concrete and crossing cracks.

Acknowledgements

The authors wish to acknowledge Dr. B. Maker of LLNL, Prof. R.L. Taylor of the University of California, and Prof. J.C. Simo of Stanford University for their help and interest in carrying out this work.

1. Y.R. Rashid, *Nuc. Eng. Des.* **7**, 334 (1968).
2. D.C. Drucker and W. Prager, *Q. Appl. Math.* **10**, 157 (1952).
3. E.L. DiMaggio and I.S. Sandler, *J. Eng. Mech.* **97**, 935 (1971).
4. M. Ortiz, *Mech. Mat.* **4**, 67 (1985).
5. L.J. Sluys, *Wave Propagation, Localization and Dispersion in Softening Solids*, Ph.D. Dissertation, Delft University of Technology (1992).
6. J. Oliver, *Int. J. Numer. Meth. Eng.* **28**, 461 (1988).
7. S. Govindjee and J.C. Simo, *J. Mech. Phys. Solids* **39**, 87 (1991).
8. N.H. Burns and C.P. Seiss, University of Illinois Civil Eng. Studies SRS No. 234, (1962).
9. B. Bresler and A.C. Scordelis, *J. Am. Concr. Inst.* **60**, 51 (1963).

**DATE
FILMED**

1 / 3 / 94

END

