


LA-UR-92-4306 C.1

HIGH PERFORMANCE COMPUTING

1993

GRAND CHALLENGES IN COMPUTER SIMULATION

Edited by
Adrian Tentner



SCANNED MAR 11 1998



LOS ALAMOS NATIONAL LABORATORY
3 9338 00374 8232



PAGOSA: A MASSIVELY-PARALLEL, MULTI-MATERIAL HYDRODYNAMICS MODEL FOR
THREE-DIMENSIONAL HIGH-SPEED FLOW AND HIGH-RATE MATERIAL DEFORMATION*

D. B. Kothe, J. R. Baumgardner, J. H. Cerutti, B. J. Daly, K. S. Holian,
E. M. Kober, S. J. Mosso, J. W. Painter, R. D. Smith, and M. D. Torrey
Los Alamos National Laboratory
Los Alamos, NM 87545

ABSTRACT

A new data parallel model for three-dimensional (3-D) high-speed fluid flow and high-rate material deformation is described. The model, called PAGOSA, has been developed recently by a team of computational physicists, numerical analysts, and computer scientists at the Los Alamos National Laboratory (LANL) on the Connection Machine (CM) series of massively-parallel supercomputers. With its efficient material interface reconstruction algorithm and finite-difference approximations on an Eulerian mesh, PAGOSA is well-suited for modeling transient flows involving multiple immiscible fluids and/or distinct materials experiencing large distortion. The evolving suite of physical models in PAGOSA currently includes models for compressible hydrodynamics, realistic equations of state, elastic-plastic material deformation, reactive burn of energetic materials, neutron transport, and turbulence effects. The PAGOSA algorithms are clear, concise, and portable due to implementation in data parallel fashion with the Fortran 90 programming language. A brief overview of the physical models, numerical methods, and parallel algorithms embodied in PAGOSA is given. Its capabilities are then illustrated with a simulation of an oil well perforation process.

INTRODUCTION

Massively-parallel computing architectures represent a new era in the field of supercomputing, making possible the solution of problems previously impractical with conventional vector supercomputers. Greater memory capacity and faster CPU speeds are afforded by parallel supercomputers through the cooperative effort of hundreds or thousands of individual processors, each working on its own portion of the problem in parallel with the others, and communicating its results with other processors when necessary. When used efficiently, parallel architectures offer large benefits, including the ability to solve larger problems in less time.

* This work was supported by the U. S. Department of Energy under contract W-7405-ENG-36.

Two principal strategies for parallel computing are the single-instruction, multiple-data (SIMD) paradigm and the multiple-instruction, multiple-data (MIMD) paradigm. In the SIMD paradigm, processors perform their operations in lockstep under the control of a master processor, proceeding more or less simultaneously with all other processors. The SIMD paradigm is often referred to as a data parallel model because it focuses on exploiting parallelism in the data (Thinking Machines Corporation 1991). The Connection Machines at LANL that have been used in formulating and developing the PAGOSA high-speed flow model are examples of data parallel computers. In the MIMD paradigm, processors operate independently on their data using potentially different instructions, with their work coordinated via the explicit passing of messages.

While PAGOSA originated at LANL and is still undergoing active data parallel development, it is also being rewritten in message-passing fashion at Sandia National Laboratories for performance analysis on MIMD parallel architectures (Gardner, Cline, and Vaughan 1992a; Gardner and Fang 1992b; Gardner *et al.* 1992c; Gardner, Cline and Vaughan 1992d). The ultimate goal for models like PAGOSA is to combine the attractive features of both paradigms in a hybrid, portable model that optimizes performance and ease of implementation.

At the present, the basic physical models and underlying numerical algorithms in PAGOSA are similar to the MESA (Holian *et al.* 1990) high-speed flow model. Recent comprehensive reviews of the physics and numerics existing in related modern high-speed flow models (often called "hydrocodes") can be found in the literature (Benson 1992). Most of these hydrocodes are currently designed for high performance on conventional vector supercomputers.

PHYSICAL MODEL

The time evolution of each material (subscript m) in the PAGOSA computational domain is governed by the continuum mechanical conservation relations. A "one-fluid" approximation is used in which all materials move with the same velocity V_i . Materials are initialized with volume fractions α_m , density ρ_m , and specific internal energy ϵ_m . Given a velocity field V_i ,



the total strain rate $\dot{\epsilon}_{ij}$ and deviatoric strain rate $\dot{\epsilon}_{ij}^p$ can be computed,

$$\dot{\epsilon}_{ij} = \frac{1}{2} \left(\frac{\partial V_i}{\partial x_j} + \frac{\partial V_j}{\partial x_i} \right); \quad \dot{\epsilon}_{ij}^p = \dot{\epsilon}_{ij} - \frac{1}{3} \delta_{ij} \dot{\epsilon}_{kk}, \quad (1)$$

enabling the response of each material to be time-integrated via solutions of

$$\frac{d\rho_m}{dt} = -\rho_m \dot{\epsilon}_{kk}, \quad (2)$$

for the conservation of mass,

$$\rho \frac{dV_i}{dt} = \frac{\partial s_{ji}}{\partial x_j} - \frac{\partial P}{\partial x_i}, \quad (3)$$

for the conservation of momentum, and

$$\rho_m \frac{d\epsilon_m}{dt} = s_{mij} \dot{\epsilon}_{mij}^p - P_m \dot{\epsilon}_{kk}, \quad (4)$$

and for the conservation of specific internal energy, where d/dt denotes a total Lagrangian derivative. In Equations (3) and (4), $\dot{\epsilon}_{mij}^p$ is the plastic deviatoric strain rate, P_m is the material pressure, and s_{mij} is the material deviatoric stress due to material response effects. Material-averaged quantities ρ , P , and s_{ij} are obtained from volume-weighted averages.

The system of conservation equations is closed with constitutive relations for the material pressure and deviatoric stress. The material pressure is given by an equation of state (EOS) relation,

$$P_m = P_m(\rho_m, \epsilon_m), \quad (5)$$

that currently can have various analytic forms: an ideal gas form for gases, Mie-Grüneisen and various polynomial forms for solids, the JWL (Dobratz and Crawford 1985) or BKW (Mader 1979) form for reaction product gases of high explosives, and the HOM (Mader 1979) form for partially reacted explosive materials. Equation (5) is scaled by the reaction product burn fraction if the material is a high explosive that is reacting via a program or Chapman-Jouguet volume burn algorithm (Mader 1979). A multi-material EOS is currently being formulated that allows materials in direct contact with one another (residing in the same computational cell) to relax at a finite rate toward pressure equilibrium.

The material deviatoric stress is obtained from strain rate decomposition plus the rate form of Hooke's stress-strain relation:

$$\dot{s}_{mij} = 2\mu_m (\dot{\epsilon}_{ij} - \dot{\epsilon}_{mij}^p) + \Omega_{ik} s_{mkj} - s_{mik} \Omega_{kj}, \quad (6)$$

where μ_m is the elastic shear modulus and Ω_{ij} is the rigid body rotation rate tensor, defined by

$$\Omega_{ij} \equiv \dot{R}_{ij} R_{ij}^t, \quad (7)$$

where R_{ij} is the rigid body rotation tensor. The Ω_{ij} terms in Equation (6) arise because constitutive laws based on Hooke's law are formulated from objective

stress rates as observed in the material frame rotated back to the laboratory frame (Davidson and Maudlin 1992). A material is assumed to yield when it reaches the Von Mises yield surface, defined by

$$F_m \equiv \frac{3}{2} s_{mij} s_{mij} - \sigma_m^2 = 0, \quad (8)$$

where σ_m is the material flow stress (yield strength). Given the yield surface F_m , associative flow defines $\dot{\epsilon}_{mij}^p$ as

$$\dot{\epsilon}_{mij}^p = \dot{\lambda}_m \left(\frac{\partial F_m}{\partial s_{mij}} \right), \quad (9)$$

where $\dot{\lambda}_m$ follows from the constraint $\dot{F}_m = 0$:

$$\dot{\lambda}_m = 3\dot{\epsilon}_{ij} s_{mij} / 2\sigma_m^2, \quad (10)$$

which states that a material stress state must remain on the yield surface during plastic flow. In addition to having the form of a simple constant, σ_m can also be obtained from either a Johnson-Cook (Johnson and Cook 1983) or Steinberg-Guinan constitutive relation (Steinberg, Cochran, and Guinan 1980).

PAGOSA also has a number of other important physical models supplementing the continuum mechanical models just discussed. Turbulence models in the form of a $k-s$ model (Daly 1992) and the variable density BHR model (Besnard, Harlow, and Rauenzahn 1987) are recent additions. A neutron transport model for reactor physics studies is also available via the 3-D multi-group, even-parity NIKE algorithm (Morel *et al.* 1991). The Forest-Fire (FF) and Multiple-Shock Forest Fire (MSFF) high explosive reactive burn models (Mader and Forest 1976; Mader 1979) are also available to approximate the initiation, shock desensitization, reactive growth, and detonation propagation of high explosives. The FF and MSFF models are based on burn rates derived from the experimental run distance to detonation dependence on shock pressure. A phenomenological "dynamic burn" model also exists, where burning of high explosive is dictated by an internal energy criterion (McCall and Bennion 1992).

NUMERICAL MODEL

The governing equations in the previous section are solved numerically in PAGOSA in an Eulerian (fixed) frame. The computational domain is partitioned in Cartesian geometry into fixed, logically-connected, orthogonal hexahedra. Orthogonality of the control volumes simplifies finite volume integrals of the equations that yield the finite difference approximations. Flow variables are laid out on the mesh in a standard "staggered-mesh" fashion, with each component of the velocity field residing at vertices and all other variables located at cell centers. The conservation equations are solved using second-order accurate finite difference approximations with a conventional Lagrangian/remap scheme.

The strength of the PAGOSA numerical algorithm lies in its treatment of the nonlinear advective terms which represent movement of material through the computational mesh. These terms introduce unwanted numerical dispersion and diffusion such as the unphysical spreading of thin material layers. This problem is addressed with a special treatment developed by D. Youngs (Youngs 1982), in which material interfaces are reconstructed from volume fraction data. The resulting interface geometry information is exploited for a more accurate estimate of material advection volumes to maintain more compact, discontinuous interfaces. This provides a Lagrangian-like representation of material interfaces with arbitrarily complex topology, and gives PAGOSA the ability to model realistically large fluid and material distortion.

In the Lagrangian phase, the conservation equations in the absence of advective terms are advanced explicitly in time. The computational mesh therefore remains embedded in the fluid during this phase. The velocity field is advanced with time-centered stresses and strain rates using a predictor-corrector method. Artificial viscosity, a tensor in general, is used to dissipate energy within shocks and to allow shock-capturing. The rigid body rotation rate tensor Ω_{ij} is approximated by the spin tensor W_{ij} , which is the antisymmetric part of $\dot{\epsilon}_{ij}$. Deviatoric stress states and plastic strain rates are estimated numerically by returning the stress state to the yield surface along the normal to the yield surface (radial return).

In the remap or advection phase, the hyperbolic advection equation

$$\frac{\partial \phi}{\partial t} + a_i \frac{\partial \phi}{\partial x_i} = 0 \quad (11)$$

is integrated forward in time, where the quantity ϕ represents any conserved (or nonconserved) quantity such as mass, momentum, internal energy, stress deviator, etc. The characteristic speed a_i is the local time-centered fluid velocity. A variant of van Leer's monotonic upwind scheme (van Leer 1977) discussed by Youngs (Youngs 1982) is currently used in PAGOSA. The algorithm uses a 3-point interpolation scheme for the ϕ gradient within a second-order Taylor series expansion, resulting in third-order accuracy (in the absence of limiting) on uniform meshes. Strang directional splitting (Strang 1968) is used to build the solution to Equation (11) successively along each coordinate direction. The directional sweep order is alternated on successive computational cycles to help mitigate the $\mathcal{O}(\delta t^2)$ cross-derivative error terms arising from directional-splitting of the advection operator.

Space restrictions prevent a more detailed discussion of the PAGOSA numerical model. The interested reader is referred to the excellent work of Youngs (Youngs 1982), Holian (Holian *et al.* 1990), and the comprehensive review article by Benson (Benson 1992).

DATA PARALLEL IMPLEMENTATION

PAGOSA is currently under active development on the LANL CM-200 and CM-5 parallel supercomputers developed by Thinking Machines Corporation (Hillis 1985; Thinking Machines Corporation 1991). The PAGOSA source code is written almost entirely in CM Fortran, the CM version of Fortran 90, which is the ANSI/ISO standard Fortran 90 (Adams *et al.* 1992) with function extensions such as FORALL and other intrinsics.

The CM-200 at LANL is a SIMD parallel supercomputer that has 2^{16} bit-serial processors and 2^{11} 64-bit Weitek floating-point units (FPUs) connected as a hypercube. Each bit-serial processor has 2^{10} Kbits of random access memory (RAM), providing a total memory of 8 GBytes. The Weitek FPUs have a 10-MHz clock, giving the CM-200 a theoretical peak speed of 40.9 GFlops (Flops = floating-point operations per second).

The CM-5 is a more flexible parallel supercomputer combining the attractive features of existing parallel architectures, including fine- and coarse-grained concurrence, MIMD and SIMD control, and fault tolerance (Thinking Machines Corporation 1991). It provides hardware support for both the data parallel and message-passing programs. The LANL CM-5 has 1024 processing nodes (PNs), each with a RISC microprocessor (SPARC technology), a network interface chip, a 64-bit bus, and 4 vector units (VUs) having 8 MBytes of RAM apiece. The VUs are capable of 32 MFlops of 64-bit floating-point performance, giving each PN a peak performance of 128 MFlops. A 1024-PN CM-5, then, has a peak performance of 131 GFlops operating on 32 GBytes of memory.

The PAGOSA parallel array dimensions (those that are spread across the processors) correspond to the (x, y, z) coordinates. The other dimension, needed for multiple materials, is "serial", or an in-processor dimension that exists entirely within the memory space of each processor. Elements of all arrays with identical spatial coordinate indices therefore reside within the same physical processor. The CM Fortran compiler currently lays data out on the machine by allocating a portion of each parallel array to each processor, termed the "subgrid" because it represents a subvolume of the mesh. The programmer will have the option of more direct control over the subgrid in the future, much like domain decomposition is user-determined on current MIMD machines.

Communication of data is almost entirely nearest-neighbor, termed "NEWS" on the CM (for North-East-West-South), and is accomplished with the Fortran 90 CSHIFT and EOSHIFT functions. The few exceptions are the global reductions (reducing the data in an array to one value, such as a maximum) used

for diagnostics and time step restrictions. Communication requirements are only nearest-neighbor because PAGOSA is based on a time-explicit finite-difference formulation using a logically-connected data structure mapped onto a regular Cartesian mesh. The explicit nature minimizes the number of needed global reductions and the regular data structure eliminates the need for non-local communication, both of which degrade performance on the CM. The PAGOSA algorithms are highly-parallelized and well-suited for the CM, as current performance analysis indicates that >95% of the computational time is spent in parallel operations on the CM, only 10-20% of which is needed for communication. A portion of this parallel computation time, however, is currently spent performing useless work as discussed below.

Algorithms that may not parallelize easily are those that require logical branching. The Fortran 90 WHERE construct, used in conjunction with masks, however, significantly alleviates this problem, and enables parallel array operations to be conditional. Legislating that different operations be performed on different portions of an array is also allowable without breaking parallelization. This feature is used in PAGOSA to great advantage, for example, in the application of boundary conditions and in the upwind differencing aspects of the advection calculation.

A component of the PAGOSA algorithm that did not parallelize easily is the treatment of cells containing more than one material ("mixed cells"). Despite the fact that the percentage of mixed cells at any given time is usually quite small, the current PAGOSA algorithm treats every cell as though it were mixed, with memory allocated accordingly, even though the cell may contain only one material. A large load-balance penalty is therefore incurred frequently, since for problems with many materials most processors are doing useless work a large fraction of the time. This also results in inefficient memory usage for problems with large (> 10) numbers of materials. It is surprising that, in spite of this load-balance problem, PAGOSA performs well on the CM relative to MESA (Holian *et al.* 1990), a similar 3-D flow model written for conventional vector machines. This improved performance has prompted other groups to follow suit, using the same basic data parallel implementation (Olson and Kimsey 1992).

NUMERICAL EXAMPLE

The modeling capabilities of PAGOSA are nicely illustrated with a simulation of a problem important to the oil industry known as oil well perforation. Well holes are typically lined with steel pipe and/or concrete casing that usually must be perforated with tiny high-explosive charges ("oil well perforators") prior to pumping. Perforation allows production of oil from

specific depths predetermined from logging data. The perforators are inserted into the well hole inside "carrier tubes", and then detonated when the tube has been lowered to the prescribed depth. They are designed to make clean holes in the casing and to penetrate several inches outward into the surrounding oil-bearing strata. By modeling the perforation process, PAGOSA can be used to study perforator performance, i.e., hole size and penetration depth, as a function of perforator design and layout, tubing/casing geometry, rock formation, and other design parameters.

Consider two perforator charges aimed horizontally in opposite directions inside a steel carrier tube that has been inserted down-hole in a main oil well casing. The perforators are very similar to a current industrial design, with a conical copper liner surrounded by high explosive and a steel case. The carrier tube is positioned flush against one side of the well casing. Each charge is point-detonated at the apex of the high explosive layer surrounding the conical copper liner. Energy release in the detonated explosive then causes the liners to converge and form shaped-charge jets that perforate the steel carrier tube and casing, and penetrate into the surrounding rock.

Inner diameters and thicknesses of the carrier tube and casing are 1.22/0.175 and 4.89/0.30 inches, respectively. The conical copper liners in the tiny perforators are 75 mils thick, having a base-to-apex height of 0.75 inch and a base radius of 0.34 inch. The oil-bearing rock is modeled as quartz because of its similar mass density, and the high explosive in the perforators is cyclotol. A simple elastic-plastic material response model is used for the copper and steel. All materials are assumed to be described by a Mie-Grüneisen equation of state, except for the perforator high explosive, which follows a JWL form.

The PAGOSA model of this oil well perforation scenario is shown in Figures 1 and 2 at the times of 0, 10, and 60 μ s after charge initiation. At 10 μ s, the high explosive has converged the conical liners into jets that have already perforated the carrier tube. The bottom charge aimed at the left has almost penetrated the casing, while the top charge aimed at the right is beginning to jet across the casing interior to the other side. By 60 μ s, both perforators have pierced the casing and are making their way into the rock. This particular scenario results in successful perforation for both charges.

The quantitative accuracy needed for this complex simulation is dictated in part by adequate resolution of all relevant time and length scales. This resolution places CPU speed and memory capacity requirements that would be excessively prohibitive were it not for today's massively parallel supercomputers. The PAGOSA oil well perforator calculation was performed




Figure 1: Side view of a PAGOSA oil well perforator simulation, looking into an oil well cut in half by a vertical plane. Shown are material interfaces for the perforator liners and covers, perforator tubing encasement, and oil well casing at times of 0, 10, and 60 μ s after detonation of the perforator charges.

in 10.3 CPU hours on 512 processing nodes of the CM-5 (Thinking Machines Corporation 1991) in the Advanced Computing Laboratory (ACL) at LANL. Reliable modeling of the jet formation process dictates a 0.5 mm zone dimension, and results in a total of 1.9×10^6 cells for the computational domain and a total memory requirement of 3 GBytes.

SUMMARY

The data parallel algorithms developed and implemented in the PAGOSA high-speed flow model have made possible accurate simulations of complex 3-D flows within the SIMD paradigm of parallel supercomputers such as the Connection Machine CM-200 and CM-5. Performance competitive with and frequently




Figure 2: A corresponding down-hole view of the PAGOSA oil well perforator simulation shown in Figure 1.

surpassing conventional vector supercomputers is realized. The data parallel programming model embodied in PAGOSA provides an easy, portable framework for further enhancement of physics models and hydrodynamic algorithms. The large, scalable memory offered by parallel machines provides the capacity for adequately resolved simulations of realistic 3-D systems, rendering parallel flow models like PAGOSA attractive candidates for parameter studies and systems design.

ACKNOWLEDGEMENTS

The authors wish to thank J. Fowler and F. Ortega for providing the *Eulerian Viewer* graphics software used in visualization of PAGOSA data sets, M. Grey for generating the figures in this paper, R. Karpp for assistance in the oil well perforator simulation, and the LANL ACL and on-site TMC staff for the hardware and software support necessary to carry out the PAGOSA simulations.

REFERENCES

- Adams, J. C.; W. S. Brainerd; J. T. Martin; B. T. Smith; and J. L. Wagener. 1992. *Fortran 90 Handbook: Complete ANSI/ISO Reference*. McGraw-Hill, New York, NY.
- Benson, D. J. 1992. "Computational Methods in Lagrangian and Eulerian Hydrocodes." *Computer Methods in Applied Mechanics and Engineering*, vol. 99 (Sept.): 235-394.
- Besnard, D.; F. Harlow; and R. Rauenzahn. 1987. "Conservation and Transport Properties of Turbulence with Large Density Variations." LA-10911-MS. Los Alamos National Laboratory, Los Alamos, NM (February).
- Daly, B. 1992. "Turbulence in MESA-2D." LA-12257-MS. Los Alamos National Laboratory, Los Alamos, NM (March).
- Davidson, R. F. and P. J. Maudlin. 1992. "A Continuum Code Investigation of Stress Integration Using Exact and Approximate Material Rotation." LA-12389-MS. Los Alamos National Laboratory, Los Alamos, NM (November).
- Dobratz, B. M. and P. C. Crawford. 1985. "LLNL Explosives Handbook." UCRL-52997 (Change 2). Lawrence Livermore National Laboratory, Livermore, CA (January).
- Gardner, D. R.; D. D. Cline; and C. T. Vaughan. 1992. "Implementation of a Single Version of PAGOSA on MIMD Hypercubes." SAND92-0640. Sandia National Laboratory, Albuquerque, NM (March).
- Gardner, D. R. and H. E. Fang. 1992. "Three-Dimensional Shock Wave Physics Simulations on Massively Parallel Supercomputers." SAND92-0404C. Sandia National Laboratory, Albuquerque, NM (July).
- Gardner, D. R.; D. D. Cline; C. T. Vaughan; R. Krall; and M. Lewitt. 1992. "The Performance of PAGOSA on Several SIMD and MIMD Parallel Computers." SAND92-1452. Sandia National Laboratory, Albuquerque, NM (October).
- Gardner, D. R.; D. D. Cline; and C. T. Vaughan. 1992. "The Performance of PAGOSA on Several MIMD Parallel Computers." SAND92-1880C. Sandia National Laboratory, Albuquerque, NM (November).
- Hillis, W. D. 1985. *The Connection Machine*. The MIT Press, Cambridge, MA.
- Holian, K. S.; D. A. Mandell; T. F. Adams; F. L. Adessio; J. R. Baumgardner; and S. J. Mosso. 1990. "MESA: A 3-D Computer Code for Armor/Anti-Armor Applications." In *Proceedings of the Supercomputing World Conference* (San Diego, CA, June 16-19).
- Johnson, G. R. and W. H. Cook. 1983. "A Constitutive Model and Data for Metals Subjected to Large Strains, High Strain Rates, and High Temperatures." In *Proceedings of the Seventh International Symposium on Ballistics* (The Hague, The Netherlands, April) 541-547.
- Mader, C. L. 1979. *Numerical Modeling of Detonations*. University of California Press, Berkeley, CA.
- Mader, C. L. and C. A. Forest. 1976. "Two-Dimensional Homogeneous and Heterogeneous Detonation Wave Propagation." LA-6259. Los Alamos National Laboratory, Los Alamos, NM (June).
- McCall, G. H. and S. T. Bennion. 1992. Private communication. Los Alamos National Laboratory, Los Alamos, NM (December).
- Morel, J. E.; J. M. McGhee; L. A. Olvey; and G. W. Claborn. 1991. "Solution of the 3-D Even-Parity S_N Equations on the Connection Machine." In *Proceedings of the International Topical Meeting on Advances in Mathematics, Computations, and Reactor Physics* (Pittsburgh, PA, April 28-May 2). American Nuclear Society, LaGrange Park, IL, vol. 1, sec. 1.1, 2-1-2-11.
- Olson, M. A. and K. D. Kimsey. 1992. "Calculation of Elastic-Plastic Wave Propagation on the Connection Machine." BRL-TR-3360. Ballistic Research Laboratory, Aberdeen Proving Ground, MD (June).
- Steinberg, D. J.; S. G. Cochran; and M. W. Guinan. 1980. "A Constitutive Model for Metals Applicable at High-Strain Rates." *Journal of Applied Physics*, vol. 51 (March): 1498-1504.
- Strang, G. 1968. "On the Construction and Comparison of Difference Schemes." *SIAM Journal of Numerical Analysis*, vol. 5 (Sept.): 506-517.
- Thinking Machines Corporation. 1991. *The Connection Machine CM-5 Technical Summary*. Cambridge, MA (Oct.).
- van Leer, B. 1977. "Towards the Ultimate Conservative Difference Scheme. IV. A New Approach to Numerical Convection." *Journal of Computational Physics*, vol. 23 (March): 276-299.
- Youngs, D. L. 1982. "Time-Dependent Multi-Material Flow with Large Fluid Distortion." In *Numerical Methods for Fluid Dynamics*, K. W. Morton and M. J. Baines, eds. Academic Press, NY, 273-285.