JOURNAL OF COMPUTATIONAL PHYSICS 20, 119–129 (1976)

Relativistic Fluid Dynamics Calculations with the Particle-in-Cell Technique*

F. H. HARLOW, A. A. AMSDEN, AND J. R. NIX

Theoretical Division, University of California, Los Alamos Scientific Laboratory, Los Alamos, New Mexico 87545

Received June 17, 1975

A technique is described for solving numerically the fully relativistic equations of fluid dynamics, for time-varying flows in two and three space dimensions. Based on the particle-in-cell technique, this extension has been developed especially for the investigation of high-energy collisions between atomic nuclei. Results are shown of a calculation in which an oxygen nucleus (^{16}O) impacts onto a uranium nucleus (^{238}U) at a speed that is 0.9516 times the vacuum speed of light.

INTRODUCTION

It has been suggested that the equations for relativistic fluid dynamics may furnish a very useful approximation for the analysis of high-energy collisions between atomic nuclei. The model is essentially that of a purely classical collision between droplets, with relativistic shocks, rarefactions, and fragmentation. Nuclear fluid is described by an equation of state in which pressure varies only as a function of rest-frame energy density and nucleon density. Viscosity, Coulomb energy, and surface tension are negligible, in contrast to the nonrelativistic circumstances in liquid–drop analyses of nuclear fission and low-energy collisions.

As a first step in the analysis of high-energy nuclear collisions with this classical model, we consider the axially symmetric case of an impact trajectory along the line of centers. There is some indication that coincident-pulse experimental techniques can distinguish collisions in the laboratory with this restrictive condition. Off-axis collisions will require fully three-dimensional solutions of the equations, which also are discussed briefly in this paper.

The technique that we use for solving the equations is an extension of the particle-in-cell (PIC) method [1, 2], which has been used widely for the analysis

* This work was performed under the auspices of the United States Energy Research and Development Administration.

of numerous nonrelativistic fluid flow problems involving large distortions in several space dimensions. This paper describes the numerical technique; the results of a series of calculations are presented elsewhere [3].

The Differential Equations

The basic variables in the frame of the observer (e.g., the laboratory frame) are the conserved quantities,

N, the number of nucleons per unit volume,
$\mathbf{M}$, the total momentum per unit volume,
E, the total energy per unit volume.

In the local rest frame of the material, the appropriate variables are

n, the number of nucleons per unit volume,
ϵ, the total energy per unit volume.

With the speed of light chosen to be unity throughout, the material velocity is $\mathbf{u}$, with cylindrical-coordinate components u and v in the r and z directions, respectively. Let

$$\gamma \equiv (1 - \mathbf{u} \cdot \mathbf{u})^{-1/2}. \tag{1}$$

We can relate quantities in the laboratory frame to those in the local rest frame of the material by the equations

$$N = \gamma n, \tag{2}$$

$$\mathbf{M} = \gamma^2(\epsilon + P)\mathbf{u}, \tag{3}$$

$$E = \gamma^2\epsilon + (\gamma^2 - 1)P, \tag{4}$$

in which the pressure P, is a function of ϵ and n, specified by the nuclear equation of state. For example, with n_0 denoting the normal rest-frame nucleon density and m_0 denoting the rest mass of a nucleon, we illustrate the numerical methodology with the following form for the equation of state:

$$P = \tfrac{2}{3}\epsilon + n_0 m_0 h(n/n_0). \tag{5}$$

With $\xi \equiv n/n_0$, a typical expression for $h(\xi)$ is given by

$$h(\xi) = -\tfrac{2}{3}\xi - 0.02470\,\xi^2 + 0.03592\,\xi^{8/3}. \tag{6}$$

In this nomenclature, the conservation equations describing the dynamics can be written

$$(\partial N/\partial t) + \nabla \cdot (\mathbf{u}N) = 0, \tag{7}$$

$$(\partial \mathbf{M}/\partial t) + \nabla \cdot (\mathbf{u}\mathbf{M}) = -\nabla P, \tag{8}$$

$$(\partial E/\partial t) + \nabla \cdot (\mathbf{u}E) = -\nabla \cdot (P\mathbf{u}). \tag{9}$$

Representation of the Fluid

The PIC-method calculations represent the material by a set of discrete particles, each following the motion of an element of fluid. To this extent, the numerical procedure is Lagrangian.

The spatial domain containing the fluid is subdivided into cells, each rectangular in cross section and toroidal in its revolution about the axis. The field variables are related to this mesh of cells, with the cell-centered values of energy, momentum, nucleon density, and velocity representing cellwise averages. The differential equations are approximated by finite-difference representations involving the cell-centered field variable averages, and in this respect, the numerical procedure is Eulerian.

Subject to prescribed initial and boundary conditions, the numerical solution is developed in a sequence of steps or cycles, each representing a small increment δt, in problem time. Each cycle consists of three phases:

1. A Lagrangian phase, in which the conserved field variables for each cell are advanced to tentative new values, assuming that each cell momentarily follows the motion of the fluid.

2. A rezone phase, in which each cell is returned to its original position, and the transport across cell boundaries is calculated. Actually, the cell boundaries never move; in this phase the new particle coordinates are calculated and those that cross cell boundaries are accompanied by transport of the conserved field variables.

3. Conversion of the altered field variables into rest-frame values, in preparation for proceeding to the next cycle.

In creating the initial configuration for a computer calculation, each particle is assigned coordinates and the total number (noninteger) of nucleons it carries, this last remaining constant throughout the calculation. The initial configuration represents the two nuclei before collision, each with rest-frame number density of nucleons n_0, and a corresponding laboratory-frame number density, N_0. Particles are created in a uniform layout within each cell; the number of nucleons ν_k, carried by particle number k, is calculated for each particle lying in the fluid

to be the product of N_0 and the subcell volume associated with the particle; then, the coordinate of each particle is altered slightly to give a staggered array. No particles are created for subcells whose centers lie outside of either atomic nucleus, thus exploiting the PIC-method capability for delineating the edges of material regions.

The initial density of energy or momentum in each cell is calculated to be the specified uniform value for the nucleus times the ratio of the actual number of nucleons in the cell to the number that would be present if the cell were full. Thus, the density of each quantity is the true microscopic density only for completely full cells, and proper account will be required for the fact that they are not true microscopic densities for partially filled cells near the edges of the material.

Numerical Solution

Phase 1. The convective fluxes are ignored, and the number density of nucleons for each cell is held constant. Accordingly, for this phase, we write

$$\partial \mathbf{M}/\partial t = -\nabla P, \tag{10}$$

$$\partial E/\partial t = -\nabla \cdot (P\mathbf{u}). \tag{11}$$

Introducing finite-difference nomenclature, we label cell centers in the r direction with index i, and those in the z direction with index j. Tentative new values for this phase of the calculation are labeled with a tilde. Cell-edge quantities, labeled with half-integer indices, are always to be calculated as averages of the two adjacent cell-center quantities. Our equations thus become

$$\frac{(\tilde{M}_r)_i^j - (M_r)_i^j}{\delta t} = -\frac{1}{\delta r}(P_{i+1/2}^j - P_{i-1/2}^j), \tag{12}$$

$$\frac{(\tilde{M}_z)_i^j - (M_z)_i^j}{\delta t} = -\frac{1}{\delta z}(P_i^{j+1/2} - P_i^{j-1/2}), \tag{13}$$

$$\frac{\tilde{E}_i^j - E_i^j}{\delta t} = -\frac{1}{r_i\,\delta r}(P_{i+1/2}^j u_{i+1/2}^j r_{i+1/2} - P_{i-1/2}^j u_{i-1/2}^j r_{i-1/2}) - \frac{1}{\delta z}(P_i^{j+1/2} v_i^{j+1/2} - P_i^{j-1/2} v_i^{j-1/2}). \tag{14}$$

These equations are employed even for the partially filled cells along the material edge, despite the fact that in these cells the field variables do not represent the actual microscopic densities. The validity of this procedure can be demonstrated for those edge cells that are adjacent to a nearly vertical or horizontal interface.

For an oblique interface, the procedure may introduce an error into the flux to or from an edge cell, but with sufficiently fine resolution the error appears to be negligible.

The boundary condition at the edge of the material is $P = 0$, which is accomplished by setting the cell-edge pressure equal to zero whenever a cell with material is adjacent to an empty cell. Along the symmetry axis, $\partial P/\partial r = 0$, so that the pressure on the axis is equal to the pressure at the adjacent cell center.

Calculation of the tilde quantities from these equations for all cells with material completes the first phase of the cycle.

Phase 2. The particle movement takes place in identically the same way as in our previous PIC-method calculations [2]. Each particle is given an effective velocity for its motion by means of an area-weighted interpolation among the four nearest cell-center velocities. New coordinates for each particle are calculated appropriate to following the effective velocity for a time interval δt. If the particle still lies in the same cell, then there is no further effect from its motion. If it has crossed to a new cell, then it carries some number of nucleons, together with some momentum and energy. These must be subtracted from the donor cell and added to the acceptor cell.

To accomplish the convective transport, the code converts all densities to cell-wise totals before any particles are moved. With cell volume τ_i, equal to $2\pi r_i\, \delta r\, \delta z$, the conversion is given simply by

$$\bar{N}_i^j = \tau_i N_i^j, \qquad \bar{E}_i^j = \tau_i \tilde{E}_i^j,$$
$$(\bar{M}_r)_i^j = \tau_i (\tilde{M}_r)_i^j, \qquad (\bar{M}_z)_i^j = \tau_i (\tilde{M}_z)_i^j,$$

in which the bar designates a cell-wise total. If particle number k crosses a cell boundary, then ν_k is subtracted from $\bar{N}$ for the donor cell and added to $\bar{N}$ for the acceptor cell. If the totals of the donor cell quantities just before the movement of that particular particle are $\bar{N}_d$, $\bar{E}_d$, $(\bar{M}_r)_d$, and $(\bar{M}_z)_d$, then the amounts of energy and momentum carried by the particle, to be subtracted from the donor cell and added to the acceptor cell, are

$$\delta E_k = \bar{E}_d \nu_k / \bar{N}_d, \tag{15}$$
$$\delta (M_r)_k = (\bar{M}_r)_d\, \nu_k / \bar{N}_d, \tag{16}$$
$$\delta (M_z)_k = (\bar{M}_z)_d\, \nu_k / \bar{N}_d. \tag{17}$$

After all the particles have been moved, the bar quantities are converted back to densities by dividing each by τ_i, thereby concluding the calculations for the second phase of the cycle.

Phase 3. The two previous phases have obtained the new field variable values in the laboratory frame, but to proceed with these calculations again in the next cycle requires updating the field variables in the local rest frame of the material. For this purpose, Eqs. (2)–(4) must be solved for n, ϵ, and $\mathbf{u}$, using the newly updated values of N, E, and $\mathbf{M}$. Two cases must be distinguished, the "interior" cells and the "edge" cells. We define an interior cell as any fluid cell whose eight neighbors are also all fluid cells, and an edge cell as any fluid cell with at least one empty cell among its eight neighbors.

Consider the interior cells first. With

$$\begin{aligned} A &\equiv N/n_0\,, \\ B &\equiv E/(n_0 m_0), \\ C &\equiv \mathbf{M}\cdot\mathbf{M}/(n_0 m_0)^2, \\ x &\equiv \mathbf{u}\cdot\mathbf{u}, \\ y &\equiv \epsilon/(n_0 m_0), \end{aligned}$$

the equations to be solved for x, y, and ξ in terms of the known quantities A, B, and C, become

$$A^2(1-x) = \xi^2, \tag{18}$$

$$B(1-x) = y(1+\tfrac{2}{3}x) + xh(\xi), \tag{19}$$

$$C(1-x)^2 = x[\tfrac{5}{3}y + h(\xi)]^2. \tag{20}$$

These three equations can be combined to give

$$\xi = A\left\{1 - \frac{C(5-2\xi^2/A^2)^2}{[5B+3h(\xi)]^2}\right\}^{1/2}. \tag{21}$$

Numerical solution can be accomplished efficiently in most cases by successive substitution into the right side of Eq. (21), with $\xi = 0$, initially. The procedure converges slightly more slowly than Newton's method, but avoids ambiguities of multiple roots. With ξ determined by this procedure, the rest of the unknown quantities are calculated easily from the equations

$$\begin{aligned} x &= 1 - (\xi^2/A^2), \\ y &= [B(1-x) - xh(\xi)]/(1+\tfrac{2}{3}x), \\ n &= n_0\xi, \\ \epsilon &= n_0 m_0 y, \\ P &= \tfrac{2}{3}\epsilon + n_0 m_0 h(\xi), \\ \mathbf{u} &= \mathbf{M}[x/(\mathbf{M}\cdot\mathbf{M})]^{1/2}. \end{aligned}$$

The second case is that of the edge cells. For these we do not know the actual volume occupied by the fluid, which is necessary for Eqs. (2)–(4), in which N, $\mathbf{M}$, and E are the *microscopic* densities in the fluid, not the smeared density across the whole cell that our finite-difference field variables denote. If N_{micro} were known, then its value would be given by

$$N_{\text{micro}} = \xi n_0 (1 - x)^{-1/2},$$

in which we continue the same nomenclature as used in inverting the equations for the interior (full) cells. Then, we would also have

$$E_{\text{micro}} = EN_{\text{micro}}/N,$$
$$\mathbf{M}_{\text{micro}} = \mathbf{M}N_{\text{micro}}/N.$$

Because N_{micro} is unknown at this stage, an additional equation is required. To supply it, we specify $\xi = 1.0$, which, as in previous nonrelativistic PIC-method codes, states that the local-rest-frame number density is the "normal" density of nuclear fluid. Accordingly, the equations for x and y become, after appropriate alterations to Eqs. (3) and (4),

$$B(1 - x)^{1/2}/A = y(1 + \tfrac{2}{3}x) + xh(1), \tag{22}$$
$$C(1 - x)/A^2 = x[5y/3 + h(1)]^2. \tag{23}$$

Again, these equations can be combined into a single one, in this case for x, and again are solved best by successive substitution.

The inversion procedures for interior and edge cells described above are those that are normally used in Phase 3. There are, however, some difficulties that can occur, which require special consideration. Foremost among these is the occasional occurrence of a computational cell for which there is only a spurious root with $x > 1$. Such a circumstance is physically impossible, corresponding to material speeds exceeding that of light. Nevertheless, the calculations occasionally do lead to such cells, especially at the material edge. This situation can be anticipated whenever $C > B^2$. The circumstances are apparently always associated with stagnation of the fluid relative to the computational mesh, in which case the energy flux vanishes, but the momentum flux, which can enhance the value of C, remains appreciable. A decrease in the time step per cycle helps to avoid the difficulty, but whenever it nevertheless occurs, the computer code bypasses the inversion and lets the fluid "coast" with the same speed as in the neighboring fluid cells.

In rare cases, usually when the density, n, is very small, the successive substitution method in Eq. (21) does not converge, usually indicated at an early step along the way by a negative value for the quantity within the brace. In these circum-

stances, the difficulty can be circumvented by an alternative successive substitution equation obtained by solving Eq. (21) for the ξ appearing in the linear term of $h(\xi)$, that is, the first term on the right side of Eq. (6).

Likewise, there are several alternative ways to combine Eqs. (22) and (23) for solution by successive substitution, depending on whether or not the root for x is very close to unity. For usual circumstances, with x not close to unity, we use

$$x = C\left[\frac{3 + 2x}{5B + 3Ah(1)(1 - x)^{1/2}}\right]^2.$$

Evidence of trouble with this solution procedure is usually in the form of x exceeding unity at an early stage of the iteration. If this occurs, we replace the equation by one in which the expression is rearranged to solve for the x occurring in $(1 - x)^{1/2}$.

Extension to Three Dimensions

The analysis of off-axis collisions requires fully three-dimensional calculations, which can be accomplished by means of a relatively simple extension of the above procedures. A Cartesian mesh is used, with velocity components u, v, and w in the x, y, and z directions, respectively. Correspondingly, we require three components of momentum in the laboratory frame, for which the calculations in Phase 1 are directly analogous to those in Eqs. (12) and (13). Likewise, the changes of energy in the laboratory frame are accomplished by means of a direct extension of Eq. (14).

The target particles themselves now all carry the same (noninteger) number of nucleons, and the same holds in the projectile, for which, however, the number of nucleons per particle will usually differ from the number per particle in the target.

The movement of these particles in Phase 2 is accomplished by a volume-weighted average of velocities in the adjacent cells, much like the area-weighted average used in most previous PIC-method studies. The adjustments to cellwise totals of energy and momentum are calculated by means of the same fractions described for the two-dimensional case.

The inversion in Phase 3 is independent of the number of dimensions. The same difficulties encountered in the axially-symmetric calculations can be expected for the three-dimensional studies, and can be overcome by the same procedures.

Perhaps the most difficult aspect of fully three-dimensional calculations is the relatively large amount of computer time required. Even accounting for the plane of symmetry that is present for nucleus–nucleus collision studies, it can be expected that the computer time per problem will increase over the two-dimensional studies

at least in proportion to the number of finite-difference cells required to resolve the third linear dimension. In addition, the full investigation of each nuclear system and incident energy will require five or so runs, with different displacements of the projectile's incoming vector from the center of the target.

Thus, we conclude that fully three-dimensional calculations can be accomplished by means of conceptually simple extensions of the techniques described in this paper, but at the cost of considerably greater computing time per problem.

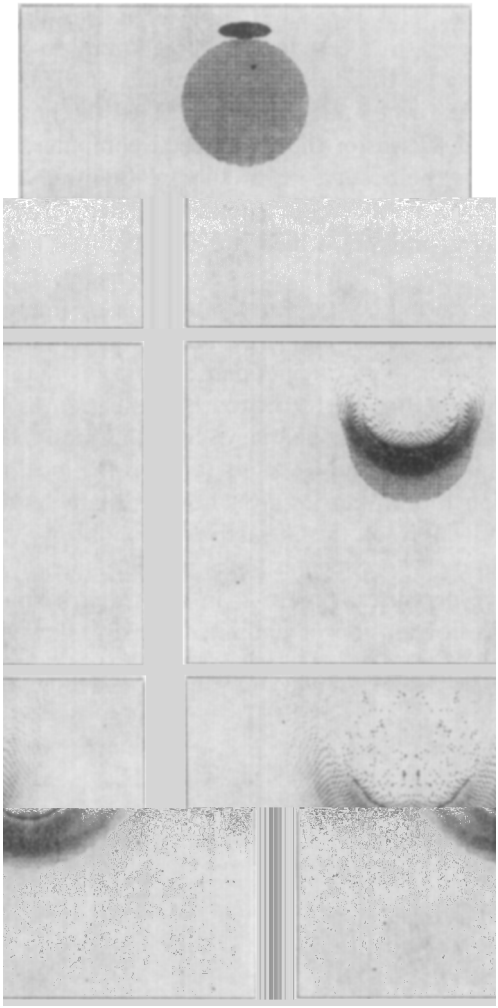

FIG. 1. Several stages in a relativistic PIC-method calculation of the impact of an oxygen (^{16}O) nucleus onto a uranium (^{238}U) target. Tensions (negative pressures) are ignored.

EXAMPLE

Examples of nucleus–nucleus impact performed by this relativistic PIC-method extension are described elsewhere [3]. Here, we show a single case to illustrate the type of studies that can be done.

Figure 1 depicts several stages in the impact of an oxygen (^{16}O) nucleus onto a uranium (^{238}U) target, at a laboratory bombarding energy of 2.10 GeV per projectile nucleon, corresponding to a projectile speed that is 0.9516 times the vacuum speed of light. This typical example of a two-dimensional calculation employed a computing mesh 35 cells in the r direction by 50 cells in the z direction, scaled such that there were 10 cells across the radius of the target. The calculation was considered complete when target debris reached the bottom boundary of the mesh, and it required 82 sec on the CDC 7600 computer. Compression of the particle spacing shows the occurrence of shocks, followed by a visible rarefying as the nuclear fluid flies apart. The computer code summarizes the energy and angular distributions of the material motion, which are the primary data for comparison with experiments.

At the high energies of interest, the viscosity, Coulomb energy, and surface tension of the nuclear material are negligible, and have been omitted. Also, no explicit artificial viscosity is required for these calculations, since the PIC method supplies sufficient dissipation as long as the perturbed parts of the fluid are not stagnant relative to the mesh, the effective kinematic viscosity being estimated by $s\,\delta x/2$, in which s is the local fluid speed and δx refers to either δr or δz.

The capability to sustain tension (negative pressure), however, is of possible significance. In the laboratory, an effect of the attractive nuclear force is the condensation of nuclear material into clusters such as alpha particles. The calculation in Fig. 1 ignores any negative pressures. Figure 2, in contrast, is the same calculation except for an allowance for tension of unlimited magnitude, and the

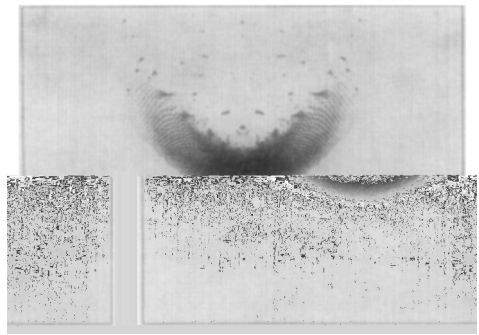

FIG. 2. Corresponding to the final frame of Fig. 1, this calculation is identical except for the allowance of negative pressures of unlimited magnitude.

presence of clumping is clearly evident. The dimensions of the clumps, however, are somewhat smaller than the computational cells, so that detailed correspondence with physical reality is not to be expected. The problem time in Fig. 2 corresponds to that of the final frame of Fig. 1.

REFERENCES

1. F. H. HARLOW, The particle-in-cell method for numerical solution of problems in fluid dynamics, *Proc. Symp. Appl. Math.* **15** (1963), 269.
2. A. A. AMSDEN, The particle-in-cell method for the calculation of the dynamics of compressible fluids, Los Alamos Scientific Laboratory Report LA-3466, 1966.
3. A. A. AMSDEN, G. F. BERTSCH, F. H. HARLOW, AND J. R. NIX, Relativistic hydrodynamic theory of heavy-ion collisions, *Phys. Rev. Letters*, **35** (1975), 905.