JOURNAL OF COMPUTATIONAL PHYSICS 120, 266–277 (1995)

Relativistic Hydrodynamics and Essentially Non-oscillatory Shock Capturing Schemes

A. DOLEZAL AND S. S. M. WONG

Department of Physics, University of Toronto, Toronto, Ontario, M5S 1A7 Canada

Received November 17, 1994; revised February 23, 1995

Numerical solutions of relativistic hydrodynamic equations are obtained with essentially non-oscillatory (ENO) finite differencing schemes. The method is explicit, conservative, consistent with the entropy condition, and high-order accurate in space and time. The present implementation is applicable to the most general, three-dimensional problems with an arbitrary equation of state. Numerical experiments, including computations of multi-dimensional flows, demonstrate that the method delivers sharp, non-oscillatory shock transitions without sacrificing high resolution of the smooth regions. This extends results already established for the Euler gas dynamics to the relativistic regime, suggesting the usefulness of ENO schemes for modelling relativistic nuclear collisions.

I. INTRODUCTION

Over the past decade, we have seen an increase of interest in computational relativistic fluid dynamics, stimulated by the ultrarelativistic nuclear collision experiments [1]. Presently, heavy nuclei are collided at kinetic energies as high as 200 GeV/A, corresponding to impact velocities within 10^{-5} of the speed of light. Since an exact theory describing such processes is not available, one has to resort to phenomenological models in order to obtain some testable predictions. An entire class of approaches takes advantage of the fact that large numbers of particles are involved. Instead of following the particles individually, such models concern themselves only with their resulting collective behaviour.

Hydrodynamic models are one example of such an approach. Here, the nuclei are considered as two droplets of a fluid whose properties are governed by a nuclear equation of state. The collision itself is then described in the language of relativistic hydrodynamics as a solution to an initial value problem. After the nuclei collide, compression shock waves form in the nuclear matter, and this presents a challenge for the numerical solution. In contrast to their non-relativistic counterparts, the relativistic equations are much more sensitive to overshoots produced by numerical oscillations. One apparent reason for that is the presence of the Lorentz factor, $(1 - v^2)^{-1/2}$, in the equations. Therefore, non-oscillatory handling of discontinuities is essential.

Traditionally, only particle-in-cell algorithms [2] or flux corrected transport methods [3, 4] (more recently [5]) have been applied. Unfortunately, these approaches are at best second-order accurate, and the oscillations are suppressed only at the cost of excessive smearing. The loss of resolution necessitates fine grids, increasing the demand both on the storage space and the computational time. As a result, three-dimensional codes may require a supercomputer to execute. What is needed is a numerical method which could deliver the non-oscillatory shock transitions while retaining high resolution.

One possible answer lies in the high-order accurate, essentially non-oscillatory (ENO) schemes. In this article, we wish to present their relativistic implementation. ENO schemes were designed as numerical approximations for hyperbolic systems of conservation laws. Instead of adaptively refining the computational grid near discontinuities, Harten, Engquist, Osher, and Chakravarthy [6, 7] proposed using a moving, *adaptive stencil*, which extracts the information about the flow from smooth regions. The algorithm generalized Godunov's scheme to an arbitrary order of accuracy in space and time. The ENO requirement relaxed the more stringent total variation diminishing (TVD) condition by allowing the presence of small oscillations on the order of truncation error. However, oscillations with a magnitude proportional to the size of the jump, commonly referred to as a Gibbs phenomenon, will not appear. Apart from correctly capturing the shocks, the scheme guarantees high resolution of the smooth regions of the flow. Even though mathematically pleasing, the algorithm was rather involved and complicated to program. It employed a cell reconstruction scheme to update the point values, and a Lax–Wendroff type time discretization, both of which become awkward to use in multi-space-dimensional problems [8, 9].

An efficient implementation was proposed by Shu and Osher [10, 11], who applied the moving stencil idea directly to the point values of the fluxes. They have also introduced convenient TVD Runge–Kutta type time discretizations. A comparison of these two approaches, termed "finite volume" and "finite differencing," can be found in Ref. [12]. The simplified, flux-based ENO schemes have been successfully tested on scalar models [10], Euler gas dynamics equations [11], Navier–Stokes equations [13], and the equations for the incompressible flow [14].

0021-9991/95 $12.00

It appears that these schemes should prove to be a viable alternative for the relativistic hydrodynamics also. As we shall see later, our results indicate that the high-order schemes can resolve most of the structures in the flow, without the need for fine or adaptive grids. At the same time, even though shock smearing is very mild, no spurious oscillations appear near the discontinuities. As a result, many calculations of interest, including three-dimensional ones, can be performed on most present-day workstations. This makes ENO algorithms extremely attractive.

In this article, we address mainly the general aspects of using the ENO method for initial value problems in relativistic hydrodynamics with an arbitrary equation of state. Applications to the relativistic nuclear collisions will be presented in a separate work. Because of the adaptiveness of the stencil, ENO schemes are highly nonlinear [15], and hence it is very difficult to obtain theoretical results regarding the stability and overall performance. Only through numerical experiments can one gain some confidence in, as well as uncover the limitations of, the method. To that aim, we present results of several numerical tests, including an actual three-dimensional calculation of a nuclear collison.

The paper is organised as follows. In the next section, the equations of relativistic hydrodynamics are introduced. An important part of their numerical solution is the reconstruction of local rest frame quantities. A simple method, which involves solving at most one non-linear equation, is given in Section III. The consistency conditions for the relativistic variables are also discussed there. In Section IV, we briefly explain the general principle of the ENO schemes. The characteristic decomposition, required by the numerical scheme, is presented in Section V. The remaining, miscellaneous issues associated with the implementation are discussed in Section VI. The last section contains the results of numerical experiments, after which a brief conclusion is given.

II. RELATIVISTIC HYDRODYNAMICS

Throughout this article, we will be using units where the speed of light is unity and, following the usual convention, we reserve the letter c for the speed of sound. The equations of the relativistic hydrodynamics express the conservation of energy, momentum, and the baryon number.[1] They are usually written in the compact, manifest-covariant form [16]

$$\partial_\nu T^{\mu\nu} = 0, \quad \partial_\nu n^\nu = 0. \tag{2.1}$$

Here, for the perfect fluid, $T^{\mu\nu} = U^\mu U^\nu(\varepsilon + p) - g^{\mu\nu}p$, is the energy-momentum tensor, and $n^\nu = U^\nu n$ is the 4-flux of baryon number. We have introduced the 4-velocity $U^\nu = \gamma(1, v_1, v_2, v_3)$, with $\gamma = (1 - v^2)^{-1/2}$ and $v = (v_1^2 + v_2^2 + v_3^2)^{1/2}$, and the Lorentz metric $g^{\mu\nu} = \mathrm{diag}(1, -1, -1, -1)$. The system is closed by the equation of state, $p = p(\varepsilon, n)$. Here p is the pressure, and ε and n, respectively, are the rest frame energy and the baryon number density.

For our purposes, it will be convenient to rewrite these equations in the form of a system of conservation laws

$$\frac{\partial \mathbf{u}}{\partial t} + \sum_{k=1}^{d} \frac{\partial \mathbf{f}^{(k)}(\mathbf{u})}{\partial x_k} = 0. \tag{2.2}$$

In general, $\mathbf{u} = (u_1, \ldots, u_n)$ is the vector of densities of the conserved variables, and $\mathbf{f}^{(k)} = (f_1^{(k)}, \ldots, f_n^{(k)})$ is the vector of their fluxes in the kth spatial direction. Let us write down $\mathbf{u}$ and $\mathbf{f}^{(k)}$ for the present case. We introduce the densities of relativistic energy

$$E = \gamma^2(\varepsilon + p) - p \tag{2.3a}$$

momentum

$$M_j = v_j\gamma^2(\varepsilon + p) \tag{2.3b}$$

and the baryon density

$$N = \gamma n. \tag{2.3c}$$

Define further the energy flux

$$\mathscr{E}_j = v_j(E + p) = M_j \tag{2.4a}$$

momentum flux

$$\mathscr{M}_{ij} = v_j M_i + \delta_{ij}p \tag{2.4b}$$

and the baryon number flux

$$\mathscr{N}_j = v_j N. \tag{2.4c}$$

here $\delta_{ij} = 1$ for $i = j$, and zero otherwise. One can verify that Eqs. (2.1) take on the form (2.2), with $\mathbf{u} = (E, M_1, M_2, M_3, N)$ and $\mathbf{f}^{(j)} = (\mathscr{E}_j, \mathscr{M}_{1j}, \mathscr{M}_{2j}, \mathscr{M}_{3j}, \mathscr{N}_j)$.

III. FLUX UPDATE

The structure of the fluxes (2.4) embodies the main difficulty in the relativistic hydrodynamics. The flux dependence on the conserved densities, $\mathbf{f}^{(k)}(\mathbf{u})$, is *implicit*, through the velocity and the local rest frame densities ε and n. A discretized version of Eqs. (2.2) will allow us to compute the values of the conserved densities, E, M_j, and N, at time $t + \Delta t$ from the values of the conserved densities and fluxes at time t. However, in order to

[1] There may be additional conserved quantities, such as the electric charge, lepton family numbers, etc., but we shall not discuss these more general cases here.

update the fluxes, we need first the new values of ε, v_j, and n. They may be found from (2.3) in the following way.

Suppose first that $N \neq 0$. Equation (2.3c) gives

$$v^2 = 1 - \frac{1}{\gamma^2} = 1 - \frac{n^2}{N^2}. \tag{3.1}$$

By combining (2.3a) and (2.3b) we find

$$\varepsilon = E - \sqrt{v^2 M^2} \tag{3.2}$$

and

$$v^2 = \frac{M^2}{(E + p(\varepsilon, n))^2}. \tag{3.3}$$

Substitute (3.1) into (3.2), then substitute the result into (3.3), and finally substitute that result into (2.3c) to give

$$n^2 = N^2 \left\{ 1 - \frac{M^2}{[E + p(E - \sqrt{M^2(1 - n^2/N^2)}, n)]^2} \right\}, \tag{3.4}$$

This equation can be solved numerically for n. Having found n, we compute ε from (3.2) and (3.1)

$$\varepsilon = E - \sqrt{M^2(1 - n^2/N^2)} \tag{3.5}$$

and finally the velocity components as

$$v_j = \frac{M_j}{E + p(\varepsilon, n)}. \tag{3.6}$$

In the case $N = 0$, which by Eq. (2.3c) implies $n = 0$, we substitute Eq. (3.2) into (3.3) to give

$$v^2 = \frac{M^2}{[E + p(E - \sqrt{v^2 M^2}, 0)]^2}. \tag{3.7}$$

We solve this equation numerically for v^2, then use (3.2) to get ε, and finally (3.6) to give the velocity components.

We note that the reconstruction of ε, v_j, and n from E, M_j, and N is unique. Indeed, suppose that there are two solutions, ε, v_j, n and ε', v_j', n'. One can always transform into a local rest frame, where $M_j = 0$. Such a transformation has no effect on the invariant ε and n. Equation (2.3b) then implies that the rest frame velocity is zero, which in turn, by (2.3c) and then (2.3a), implies $n = n'$ and $\varepsilon = \varepsilon'$.

On the other hand, there may exist unphysical, mutually incompatible values of E, M_j, and N, such that Eqs. (2.3) cannot be satisfied for any ε, v_j, and n. Even though ENO schemes guarantee the absence of the zeroth-order oscillations, smaller oscillations and thus non-physical overshoots, might still appear in principle. For this reason, the algorithm must include provisions to handle such events. In practice, these non-physical overshoots are extremely rare with ENO schemes.

With the non-relativistic Euler gas dynamics, the consistency condition is straightforward; the energy and mass densities must be positive. In the relativistic regime, the physical domain of the allowed values of energy, momentum, and baryon density is more complicated, as two additional constraints emerge. The first one expresses the causality limit on the allowed signal speed, $v^2 < 1$. The second constraint results from the equivalence of mass and energy. In relativistic thermodynamics, the rest mass of the gas constituents is included in the internal energy. Hence, even at zero temperature, we have $\varepsilon > \varepsilon_0(n) > 0$, where $\varepsilon_0(n)$ is the ground-state energy density (defined implicitly by $T(\varepsilon, n) = 0$). We remark that the baryon density—in contrast to the mass density—*may* be negative, corresponding to antimatter. However, on physical grounds we expect that an initial state with positive n everywhere (which we shall always assume) cannot develop into a state where there is a region with $n < 0$ (no spontaneous matter-antimatter separation). Thus when a negative N appears in a numerical approximation, it also indicates an overshoot.

As a practical matter, the consistency check may be done in four states: (1) Is it $N \geq 0$? (2) Does (3.4) have a solution within the domain $0 \leq n \leq N$? (3) Compute ε by (3.5). Is it $\varepsilon \geq \varepsilon_0(n)$? (4) Compute v_j by (3.6). Is it $v^2 < 1$? If, at any stage, the test fails, we do not proceed to the next stage and we terminate the computation. At stage 2, we also have to ensure that no domain error occurs in the course of solving (3.4). This more technical issue will be addressed in Section VI.

IV. ENO SCHEMES

Let us now briefly describe the numerical algorithm. For the details and proofs we refer to the original works [6, 7, 10, 11]. For simplicity, we start with an initial value problem for a one-dimensional scalar conservation law

$$\frac{\partial u}{\partial t} + \frac{\partial f(u)}{\partial x} = 0, \quad u(x, 0) = u_0(x). \tag{4.1}$$

On a regular computational mesh, $x_j = j\,\Delta x$, $t_n = n\,\Delta t$, denote the numerical approximation to $u(x_j, t_n)$ as u_j^n. ENO methods belong to the class of conservative schemes, for which the explicit (time forward) version reads

$$u_j^{n+1} = u_j^n - \frac{\Delta t}{\Delta x}(\hat{f}_{j+1/2}^n - \hat{f}_{j-1/2}^n). \tag{4.2}$$

Here, we assume stability under a Courant–Friedrichs–Lewy (CFL) restriction, $\Delta t/\Delta x < C/|f'|$, for a suitable C. The numerical flux function $\hat{f}_{j+1/2}^n \equiv \hat{f}(u_{j-k}^n, \ldots, u_{j+k+1}^n)$ is consistent with the physical flux in the sense $\hat{f}(u, \ldots, u) = f(u)$.

In that case, the Lax–Wendroff theorem [17] ensures that, if the numerical approximation converges boundedly almost

everywhere, its limit is the weak solution of (4.1). In other words, the integral conservation law will be satisfied across the discontinuities, as expressed by the Rankine–Hugoniot condition [18]. On the other hand, the theorem does not guarantee that the numerical approximation will always converge to the physical solution for which the entropy increases across the shock. It is well known that the weak solutions to (4.1) are not unique [18]. The entropy condition must be ensured separately.

A simple way to construct numerical flux functions is by decomposing the physical flux into positive and negative speed components, $f = f^+ + f^-$ with $f^{+\prime} > 0$, and $f^{-\prime} < 0$, where the prime denotes a derivative with respect to u. The numerical fluxes are first constructed separately for f^+ and f^-, and in the end one puts $\hat{f}_{j+1/2} = \hat{f}^+_{j+1/2} + \hat{f}^-_{j+1/2}$. For an rth order accurate scheme, the values $\hat{f}^{\pm}_{j+1/2}$ are determined from an $(r-1)$th order interpolation polynomial of $f^{\pm}$ on an adaptively chosen r-point stencil around x_j.

The following point lies at the heart of the ENO schemes; the stencil selection is based on the direction of the signal velocity and local smoothness. The stencil is built up in an inductive manner. Its initial point is chosen so as to achieve the proper upwinding, i.e., x_j for f^+ and x_{j+1} for f^-. Then, at each inductive step, a new point is added to the existing stencil, where a choice has to be made between extending the stencil to the left or to the right. The selection picks the one on which the flux is smoother, in the sense of difference tables. To choose the next point to be added to an existing k-point stencil, one compares the absolute magnitudes of the kth undivided differences of the flux, on the two candidate stencils.

For the decomposition, we always use the Lax–Friedrichs building blocks, $f^{\pm}(u) = \frac{1}{2}(f(u) \pm \alpha u)$, where $\alpha > \max|f'(u)|$. We shall denote such a scheme, involving a global α, as ENO-LF. The amount of numerical viscosity may be reduced by constructing the building blocks *locally*, where for each j we define a local set $f^{\pm}_i = \frac{1}{2}(f_i \pm \alpha_{j+1/2} u_i)$ for $i = j-k, \ldots, j+k+1$, with

$$\alpha_{j+1/2} = \max_{u_j \le u \le u_{j+1}} |f'(u)|. \tag{4.3}$$

The stencil selection and polynomial interpolation are then applied at each point to its local set $\{f^{\pm}_i\}$. The resulting local LF scheme will be denoted as ENO-LLF.

The numerical flux construction is then combined with the multi-step TVD Runge–Kutta time discretization introduced in Ref. [10]. As a result, we obtain an explicit scheme which has a high-order accuracy in space and time. While a general proof of stability is not available (for the third and higher order schemes), all the numerical experiments suggest that the resulting schemes are stable and satisfy the entropy condition.

The generalisation to systems (2.2) follows the characteristic projection idea due to Roe [19]. Define the Jacobi matrix, $\mathbf{A} = \partial \mathbf{f}/\partial \mathbf{u}$, and the matrix of its right eigenvectors, $\mathbf{R}$, such that $\Lambda = \mathbf{R}^{-1}\mathbf{A}\mathbf{R}$ is diagonal. The systems for which $\mathbf{A}$ is diagonalizable and has real eigenvalues are called hyperbolic. Denote further the values of $\mathbf{A}$, $\mathbf{R}$, and Λ at $\mathbf{u} = \frac{1}{2}(\mathbf{u}_j + \mathbf{u}_{j+1})$ by a suffix $j + \frac{1}{2}$. At each point, we project the variables and fluxes onto the local characteristic space, $\mathcal{U}_i = \mathbf{R}^{-1}_{j+1/2}\mathbf{u}_i$ and $\mathcal{F}_i = \mathbf{R}^{-1}_{j+1/2}\mathbf{f}_i$, for $i = j-k, \ldots, j+k+1$. In terms of these characteristic fields, the equations locally decouple, and we are in a position to apply the scheme separately to each of the characteristic fields. The role of f' in the construction of the Lax–Friedrichs blocks is now played by the corresponding eigenvalues of $\mathbf{A}$. We may set $\mathcal{F}^{\pm}_i = \frac{1}{2}(\mathcal{F}_i \pm \mathbf{K}\mathcal{U}_i)$, where $\mathbf{K}$ is a diagonal, component-wise, global major of Λ. Alternatively, we may define $\mathbf{K}_{j+1/2}$ locally, in lieu of (4.3). Having constructed the numerical fluxes for the characteristic variables, $\hat{\mathcal{F}}_{j+1/2}$, we carry out the inverse transformation and obtain the numerical flux functions for the original physical components as $\hat{\mathbf{f}}_{j+1/2} = \mathbf{R}_{j+1/2}\hat{\mathcal{F}}_{j+1/2}$. The scheme will be denoted as CH-ENO-LF, and its local version as CH-ENO-LLF (CH for characteristic).

A somewhat simpler way is to apply the scheme directly on the components, with building blocks $\mathbf{f}^{\pm}(\mathbf{u}) = \frac{1}{2}(\mathbf{f}(\mathbf{u}) \pm \boldsymbol{\alpha}\mathbf{u})$. Here the matrix $\boldsymbol{\alpha}$ is chosen such that the eigenvalues of the Jacobi matrices $\partial \mathbf{f}^+/\partial \mathbf{u}$ and $\partial \mathbf{f}^-/\partial \mathbf{u}$ are, respectively, positive and negative. One possible choice is $\boldsymbol{\alpha} = \mathbf{R}\mathbf{K}\mathbf{R}^{-1}$, with $\mathbf{K}$ as above. Again, a local version with an $\boldsymbol{\alpha}_{j+1/2}$ is also possible, and we end up with two algorithms, CW-ENO-LF and CW-ENO-LLF (CW for component-wise). Despite the simplicity of implementation and computational speed of the CW method, we shall see that the CH method usually gives better results, especially for higher orders, and is therefore preferable.

The main advantage of the flux-based ENO algorithms is the relative ease to implement them in multi-dimensions: the scheme is simply applied separately to the fluxes for each of the directions, $\mathbf{f}^{(k)}$. When constructing the stencils, one computes the difference table for the particular direction, keeping the remaining variables fixed. We stress that this is *not* the usual operator splitting method. In the time-update step, the difference expressions for all the fluxes appear on the right-hand side simultaneously.

V. CHARACTERISTIC DECOMPOSITION

For the characteristic projections we need to diagonalize the Jacobi matrices $\partial \mathbf{f}^{(k)}/\partial \mathbf{u}$. The algebraic calculations are straightforward, but tedious, especially in multi-dimensions. They can be carried out using, for example, the symbolic algebra system REDUCE [20]. For the sake of simplicity, we shall confine our discussion to the one-dimensional case

here. The results for the full three-dimensional problem can be found in Ref. [21].

As we have mentioned, the fluxes, $\mathscr{E}$, $\mathcal{M}$, and $\mathcal{N}$, depend on the conserved variables, E, M, and N, implicitly through ε, v, and n. Therefore, in order to compute the Jacobi matrix for the flux functions, we need first the elements of another Jacobi matrix,

$$\mathbf{U} = \frac{\partial(\varepsilon, v, n)}{\partial(E, M, N)} \tag{5.1}$$

whose inverse can be computed directly from Eq. (2.3). After much algebra, the nine elements of the flux function Jacobi matrix come out as follows:

$$\begin{aligned}
&\frac{\partial \mathscr{E}}{\partial E} = 0, \quad \frac{\partial \mathscr{E}}{\partial M} = 1, \quad \frac{\partial \mathscr{E}}{\partial N} = 0, \\
&\frac{\partial \mathcal{M}}{\partial E} = Q\left(w(c^2 - v^2) - \frac{n}{\gamma^2}\frac{\partial p}{\partial n}\right), \\
&\frac{\partial \mathcal{M}}{\partial M} = Qv\left(2w(1 - c^2) + \frac{n}{\gamma^2}\frac{\partial p}{\partial n}\right), \\
&\frac{\partial \mathcal{M}}{\partial N} = \frac{Qw}{\gamma^3}\frac{\partial p}{\partial n}, \quad \frac{\partial \mathcal{N}}{\partial E} = -\frac{Qnv}{\gamma w}\left(w(1 + c^2) - n\frac{\partial p}{\partial n}\right), \\
&\frac{\partial \mathcal{N}}{\partial M} = \frac{Qn}{\gamma w}\left(w(1 + v^2c^2) - nv^2\frac{\partial p}{\partial n}\right), \\
&\frac{\partial \mathcal{N}}{\partial N} = Qv\left(w(1 - v^2c^2) - \frac{n}{\gamma^2}\frac{\partial p}{\partial n}\right).
\end{aligned} \tag{5.2}$$

Here, we have introduced the enthalpy density, $w = \varepsilon + p$, the speed of sound

$$c^2 \equiv \left(\frac{\partial p}{\partial \varepsilon}\right)_s = \frac{\partial p}{\partial \varepsilon} + \frac{n}{w}\frac{\partial p}{\partial n}, \tag{5.3}$$

where s is the entropy per baryon, and defined a common factor $Q = 1/w(1 - v^2c^2)$.

Our next task is to diagonalize the flux matrix, as given by (5.2). The eigenvalues, arranged in increasing size, are

$$\lambda_1 = \frac{v - c}{1 - vc}, \quad \lambda_2 = v, \quad \lambda_3 = \frac{v + c}{1 + vc}. \tag{5.4}$$

In contrast to the Newtonian gas dynamics case, the flow velocity and the speed of sound are added relativistically. The corresponding right eigenvectors are

$$\mathbf{r}^{(1)} = \begin{pmatrix} \gamma(1 - vc) \\ \gamma(v - c) \\ \dfrac{n}{w} \end{pmatrix}, \quad \mathbf{r}^{(2)} = \begin{pmatrix} \gamma\dfrac{\partial p}{\partial n} \\ \gamma v\dfrac{\partial p}{\partial n} \\ \dfrac{n}{w}\dfrac{\partial p}{\partial n} - c^2 \end{pmatrix}, \tag{5.5}$$

$$\mathbf{r}^{(3)} = \begin{pmatrix} \gamma(1 + vc) \\ \gamma(v + c) \\ \dfrac{n}{w} \end{pmatrix}.$$

To find the left eigenvectors, form a 3×3 matrix which has the three elements of $\mathbf{r}^{(j)}$ as its jth column. The left eigenvectors will be the rows of its inverse. We find

$$\begin{aligned}
\mathbf{l}^{(1)} &= \left(\frac{\gamma}{2c^2w}\left(-n\frac{\partial p}{\partial n} + cw(v + c)\right),\right. \\
&\qquad \left.\frac{\gamma}{2c^2w}\left(vn\frac{\partial p}{\partial n} - cw(1 + vc)\right), \frac{1}{2c^2}\frac{\partial p}{\partial n}\right) \\
\mathbf{l}^{(2)} &= \left(\frac{\gamma n}{c^2 w}, -\frac{\gamma n v}{c^2 w}, -\frac{1}{c^2}\right) \\
\mathbf{l}^{(3)} &= \left(\frac{\gamma}{2c^2w}\left(-n\frac{\partial p}{\partial n} - cw(v - c)\right),\right. \\
&\qquad \left.\frac{\gamma}{2c^2w}\left(vn\frac{\partial p}{\partial n} + cw(1 - vc)\right), \frac{1}{2c^2}\frac{\partial p}{\partial n}\right).
\end{aligned} \tag{5.6}$$

The matrix built of the right eigenvectors is invertible as long as $c \neq 0$, as its determinant is equal to $-2c^3$. The eigenvalues may diverge for $vc = \pm 1$. However, within the physical, causal domain, where $|v| < 1$ and $c < 1$, this cannot happen and the system is, indeed, hyperbolic.

An important issue is whether the eigenvalues change when we move in the direction of the corresponding right eigenvectors. We recall that the kth characteristic field is called genuinely nonlinear if

$$\mathbf{r}^{(k)} \cdot \nabla\lambda_k \equiv \frac{\partial \lambda_k}{\partial E} r_1^{(k)} + \frac{\partial \lambda_k}{\partial M} r_2^{(k)} + \frac{\partial \lambda_k}{\partial D} r_3^{(k)} \neq 0. \tag{5.7}$$

Otherwise, it is called linearly degenerate [18]. The gradients of the eigenvalues are complicated and not very illuminating. We present only the final results for the directional derivatives

$$\mathbf{r}^{(1)} \cdot \nabla\lambda_1 = -\mathbf{r}^{(3)} \cdot \nabla\lambda_3 = -\frac{n\dfrac{\partial c}{\partial n} + w\dfrac{\partial c}{\partial \varepsilon} + c(1-c^2)}{w\gamma^3(1-vc)^3} \quad (5.8a)$$

$$\mathbf{r}^{(2)} \cdot \nabla\lambda_2 = 0. \quad (5.8b)$$

Just as in the non-relativistic case, the second field is always linearly degenerate. The other two fields are either both genuinely nonlinear or both linearly degenerate.

In the theory of relativistic shock waves [22], an important role is played by the quantity

$$\Delta \equiv \left(\frac{\partial^2 p}{\partial(w/n^2)^2}\right)_s. \quad (5.9)$$

"Normal" matter is characterised by $\Delta > 0$, corresponding to the situation where only compression shocks are thermodynamically stable. A straightforward computation yields

$$\Delta = \frac{2cn^4}{w(1-c^2)^3}\left(n\frac{\partial c}{\partial n} + w\frac{\partial c}{\partial \varepsilon} + c(1-c^2)\right). \quad (5.10)$$

Comparison with (5.8a) indicates how the shock stability condition might be recovered from the direction taken here by analysing the behaviour of characteristics. We shall not pursue this point further.

VI. IMPLEMENTATION NOTES

In this section, we address the remaining questions pertaining to the actual implementation of the ENO schemes.

Lax–Friedrichs Blocks. Let us start with the choice of the matrix $\mathbf{K}$ for the Lax–Friedrichs blocks. It follows from Eq. (5.4) that $\lambda_j < 1$, since the speed of light (=1 in the units used) is the ultimate signal speed. Setting $\mathbf{K} = \mathbf{1}$ will thus always work for CH-ENO-LF and, when using CW-ENO-LF, the unit matrix may also be used for $\boldsymbol{\alpha}$. As a result, for the component-wise version, the eigenvectors are not needed at all.

As for the local LF algorithms, we can take

$$\mathbf{K}_{j+1/2} = \max(|\Lambda_j|, |\Lambda_{j+1}|) \quad (6.1)$$

so long as we know that Δ from Eq. (5.9) does not change sign on a straight line connecting $\mathbf{u}_j$ with $\mathbf{u}_{j+1}$ (in particular, if it always stays positive). If Δ changes sign, we may take $\mathbf{K}_{j+1/2} = \mathbf{1}$, just to be on the safe side.

The signal speed is also related to the CFL restriction imposed on the size of the time step. All the numerical tests we have performed indicate that for the d-dimensional case, the forward version remains stable as long as $\Delta t/\Delta x < 1/d$ (for any spatial order). This restriction remains valid for the second- and third-order Runge–Kutta time discretizations. The fourth-order scheme becomes more restrictive, $\Delta t/\Delta x < 2/3d$ (see Ref. [11]).

Characteristic Decomposition. The eigenvalues and eigenvectors, given by (5.4)–(5.6), are functions of w, v_j, n, c, and $\partial p/\partial n$. The average matrices, $\mathbf{R}_{j+1/2}$ and $\Lambda_{j+1/2}$, may be computed at average values of these variables, $\frac{1}{2}(w_j + w_{j+1})$, etc. This slightly modifies the definition presented in Section IV, where they are taken at the averages of the conserved variables, E, M_j, and N here. Our simplified version produces a faster code, since it avoids the need for solving Eq. (3.4) repeatedly.

In the Euler gas dynamics, one can employ the Roe linearization [23]. The advantage of using the Roe matrix is that single shocks are "recognised" and resolved exactly. Unfortunately, in relativistic domain this is not practicable. Even though we know from theory that the matrix must exist [24], we are unable to construct it in a way which is computationally feasible.

Flux Update. For the equation of state of a relativistic ideal gas of massless particles, $p = \frac{1}{3}\varepsilon$, Eq. (3.4) leads to a quartic equation which can be solved analytically. However, for anything more complicated, the resulting equation will have to be solved numerically. This may cost a large amount of computational time, especially for schemes with high-order accuracy in time. This is because the rth order Runge–Kutta TVD time discretization is an r-step method and, hence, the local rest frame quantities must be reconstructed r times during each complete time update.

To solve Eq. (3.4) or (3.7), we used the classic bisection method, which always converges. The relation $0 < n < N$ defines the limits between which the solution must lie. In the bisection process, the argument of p could fall outside the physical domain, $\varepsilon > \varepsilon_0(n)$, as mentioned in Section IV. As a remedy, we replace p in (3.4) or (3.7) by $\tilde{p}(\varepsilon, n) = p(\max(\varepsilon, \varepsilon_0(n)), n)$, thus formally extending the definition of pressure for $\varepsilon < \varepsilon_0(n)$. If there exists a physical solution, the method will still converge to it, while domain errors will be avoided in the process. We may, however, arrive at a non-physical solution. This is not a problem either, since any such false solution will be discarded during the subsequent checking steps, outlined in Section IV.

We have found that 30–60% of the total computational time is spent in updating ε, v_j, and n. Therefore, if a particular equation of state admits a faster converging algorithm, it should be used. On the other hand, the flux update stage provides an excellent opportunity for parallelism, since the flux updates of different grid points can be done independently. We have done some preliminary testing and the results are encouraging.

Polar and Cylindrical Coordinates. In polar coordinates, the equations look formally the same as their

Cartesian counterparts, except for the additional $1/r$ terms. As a result, one-dimensional algorithms may readily be adopted for 2D problems with polar symmetry. One sensitive point is perhaps handling of the $r = 0$ boundary. Since for $r \to 0$, it must be $v_r \to 0$, and thus $M_r \to 0$, we have a fixed boundary condition for M_r. During the stencil selection stage, we want to treat the points around $r = 0$ on an equal footing with the other points. To this aim, we formally add auxiliary points to the left of $r = 0$, satisfying mirror boundary conditions: $E(-r) = E(r)$, $M_r(-r) = -M_r(r)$, and $N(-r) = N(r)$. The number of the added points depends on the width of the stencil, i..e, on the spatial order of the used scheme. Similarly, one can implement cylindrical coordinates for 3D problems with axial symmetry.

Multiple-Dimensional Cases. A serious problem with multi-dimensions is the increased storage requirements. Among the possible compromises we would like to point out one. We have found that the difference tables are actually cheap to compute (relative to the total CPU time) so that we may choose not to store the values—instead, recompute them when they are needed again at the stencil selection stage. For a third-order scheme, this decreases the memory requirements by 60%, while the computational time increases by less than 4%.

VII. NUMERICAL EXPERIMENTS

Before we begin, let us discuss our choice of the equation of state. The simplest relativistic equation of state is probably that of an ideal gas of massless particles, $p = \frac{1}{3}\varepsilon$. The problem is that it is already too simple, since $\partial p/\partial n$ vanishes (cf. Eqs. (5.5) and (5.6)). On the other hand, realistic nuclear equations of state which one would use in the modelling of relativistic heavy ion collisions are too complicated for numerical tests. Therefore, we decided to use a model equation of state which contains all the qualitative ingredients of any realistic situation, while retaining relative simplicity.

To ensure the thermodynamic consistency, we start from

$$p = AT^4 + B\mu^2. \tag{7.1}$$

Here, T is the temperature, μ is the chemical potential, and A and B are adjustable parameters. In all our calculations, we use the values $B = \frac{1}{4}$, $A = \frac{1}{3}$. To express p in terms of n and ε, we can compute the former directly, as $n = (\partial p/\partial\mu)_T = 2B\mu$. To calculate the energy density we find first the entropy density $\sigma = (\partial p/\partial T)_\mu = 4AT^3$, and then use the Gibbs–Duhem relation $\varepsilon = T\sigma - p + \mu n = 3AT^4 + B\mu^2$. Expressing T and μ in terms of ε and n and substituting the results into (7.1) produces

$$p = \frac{\varepsilon}{3} + \frac{n^2}{6B}. \tag{7.2}$$

The speed of sound squared is given by Eq. (5.3), which yields

$$c^2 = \frac{1}{3} + \frac{2n^2}{8B\varepsilon + n^2}. \tag{7.3}$$

One can check that the causality domain, $c^2 < 1$, happens to coincide with the physical domain, $T > 0$, and is expressed by the inequality $\varepsilon > \varepsilon_0(n) = n^2/4B$.

Our one-dimensional test-cases are mostly Riemann problems, where the initial conditions are described by two constant states, one to the left and one to the right of the origin, respectively labelled by L and R. For the model equation of state, the exact solutions to the Riemann problems can be found in a closed (but not analytic) form. We have run the second-, third-, and fourth-order schemes (in time and space) on most examples. To facilitate the comparison, all tests were run on meshes with 200 points, unless indicated otherwise.

EXAMPLE 1. *We start with a Riemann shock-tube problem with*

$$(v_L, \varepsilon_L, n_L) = (0, 100, 5), \quad (v_R, \varepsilon_R, n_R) = (0, 10, 2). \tag{7.4}$$

The solution is a familiar succession of a shock, contact discontinuity, and rarefaction wave, separated by regions of constant flow. In Fig. 1, the results at $t = 4$ for the second-order CH-ENO-LLF scheme are shown as diamonds. The full line is the exact solution. While there are no oscillations and no overshoot, the discontinuities and the corners of the rarefaction wave are smeared. The situation improves, if we use a fourth-order CH-ENO-LLF, as shown in Fig. 2. There are still no spurious oscillations, and only a minute overshoot in pressure and density at the shock. The results for the component-wise version, CW-ENO-LF, are shown in Fig. 3. While the overall resolution remains very good, slight oscillations appear near the contact discontinuity in the central plateau. A similar noise in an Euler gas dynamics calculation was also reported in [7]. The noise in the component-wise approach may still be acceptable, as a trade-off for the computational speed and the ease of programming. The noise appears only in high-order schemes. A second-order CW-ENO scheme (not shown) has no oscillations at all.

EXAMPLE 2. *We use a perturbed Riemann problem to document the capacity of the ENO method to transport smooth* structures across discontinuities. The left state is kept the same as in our first example, but in the right state we perturb the baryon density by a sinusoidal wave, $n_R = 2 + 0.3 \sin(5x)$, while keeping the pressure constant. In Fig. 4, the results for *the baryon density are shown at $t = 3$ for the second- (top),* third- (center), and fourth-order (bottom) CH-ENO-LLF. For comparison, the solid line is a third-order CH-ENO-LLF with 2000 points, which we may take as the exact solution. Apart from fine resolution of the structures across the discontinuities,

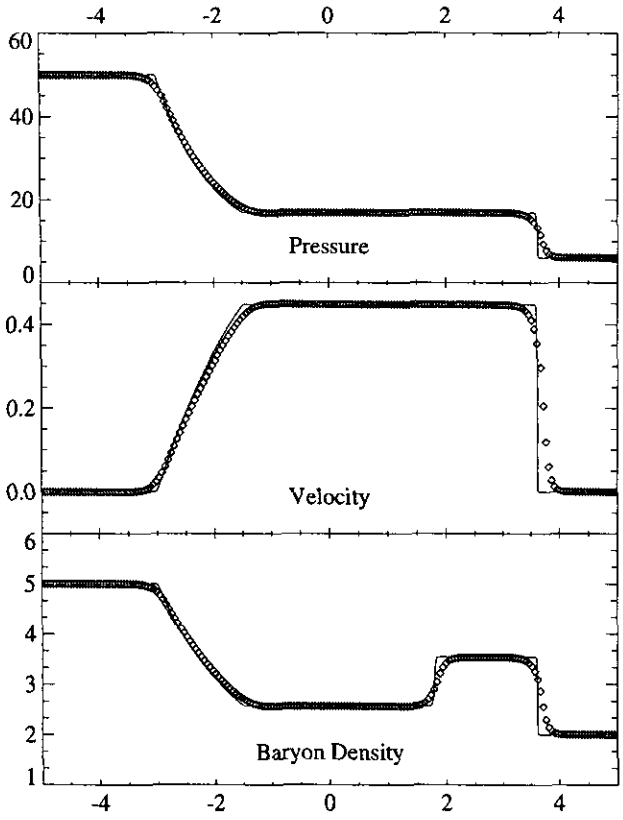

FIG. 1. Relativistic shock tube problem (Example 1), with second-order CH-ENO-LLF, 200 points, $\Delta t/\Delta x = 0.6$. The full line corresponds to the exact solution.

we can also see the improvement with the increasing order of the method.

EXAMPLE 3. This is an ultra-relativistic colliding slab Riemann problem

$$(v_L, \varepsilon_L, n_L) = (1 - 10^{-12}, 5, 1),$$
$$(v_R, \varepsilon_R, n_R) = (-(1 - 10^{-12}), 5, 1). \tag{7.5}$$

It represents a one-dimensional version of the kind of initial conditions which appear in relativistic nuclear collision calculations. Our test uses truly extreme velocities (here, $\gamma \approx 700{,}000$, well above the range needed in applications) just to demonstrate that there are no constraints to the allowed values of γ.

We present only the results for the third- (Fig. 5), and fourth-order (Fig. 6) CH-ENO-LLF schemes. One small but persistent problem is the dip in baryon density at $x = 0$. It always consists of only two points and does not spread further. The fourth-order scheme has some spurious oscillations in pressure and density; however, there is still no overshoot in velocity. The fast onset of the correct solution is shown in Fig. 7, where we present a surface plot of the baryon density as a function of time and space. This computation was performed on a grid with 50 spatial points and 50 time slices, using a third-order CH-ENO-LLF. Here, we can also see how the dip in the middle forms.

In our numerical code, we have not included any special provision to treat extremely large velocities. As a result, we will eventually hit the numerical precision barrier, when the velocity becomes indistinguishable from unity. A remedy for this extreme situation is simple. Instead of using the velocity itself, one may work with the spatial components of the 4-velocity $U_j \equiv v_j\gamma$, using $\gamma = (1 + U^2)^{1/2}$. Nothing seems to indicate that the ENO algorithm itself imposes any limit to the allowed values of γ.

EXAMPLE 4. Our two-dimensional test case is based on our shock tube example. One might call it a shock cylinder problem, for the initial conditions read

$$\begin{aligned} \varepsilon &= 100, n = 5 \quad \text{for } x^2 + y^2 < 2.5^2 \\ \varepsilon &= 10, \quad n = 2 \quad \text{for } x^2 + y^2 \geq 2.5^2 \end{aligned} \tag{7.6}$$

and $v = 0$ always. Instead of the left and right states, now we have "inside" and "outside" ones. Because of the problem's symmetry, we only have to calculate one quadrant using "per-

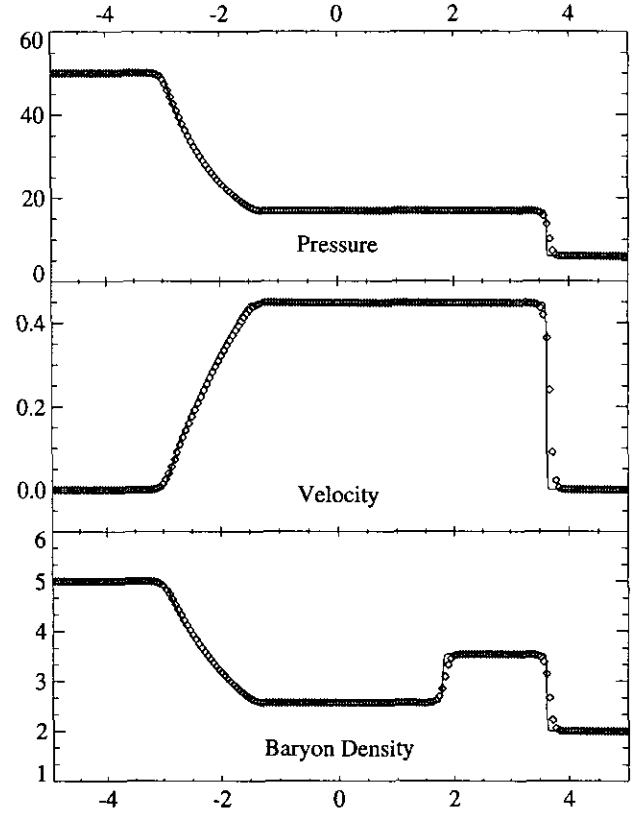

FIG. 2. Same as Fig. 1, except with fourth-order characteristic ENO-LLF and $\Delta t/\Delta x = 0.5$.

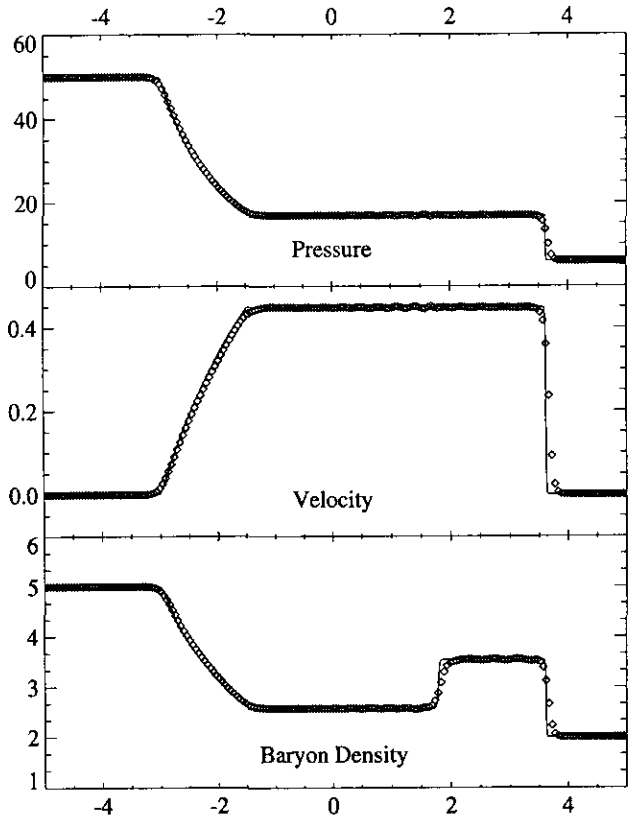

FIG. 3. Same as Fig. 1, except with fourth-order component-wise ENO-LLF and $\Delta t/\Delta x = 0.5$.

pendicular-periodic'' boundary conditions. Apart for testing the non-oscillatory behaviour in two dimensions, we are also interested in the degree of isotropy of the numerical solution.

The results for a two-dimensional, third-order CH-ENO-LLF on a 200 × 200 grid are shown in Fig. 8. A radial cut at $x = 0$ and $t = 3$ is shown in diamonds. To check the correctness of the two-dimensional computation, we also present the results obtained using a third-order CH-ENO-LLF in polar coordinates with 2000 points, shown in a full line. We see a very good agreement, with the exception of a few points near the origin, where the two-dimensional solution decays too rapidly. The flow structure is similar to the one-dimensional cases. Again, there is a shock, contact discontinuity, and rarefaction wave. However, the radial profile of the rarefaction wave is different, and the flow is not constant in the intermediate regions, since the radial fluxes vary as $1/r$.

To check the isotropy (grid aliasing) of the method, we inspect cuts done at different angles with respect to the gridlines. In Fig. 9, the error in baryon density is plotted, for cuts performed at angles $\theta = 0°$, 20°, 30°, and 45°. Since they are all very close to one another, there is no point in identifying them by different line styles. The vertical dotted lines indicate the positions of the top and bottom of the rarefaction wave, the contact discontinuity, and the shock. The oscillations on some of the lines were caused by the slicing algorithm, which uses a simple nearest-neighbour approximation.

EXAMPLE 5. Since the main motivation for our present work is modelling relativistic nuclear collisions, we present results of such a computation. The case chosen is a stationary target S–Au collision. The laboratory-frame kinetic energy of the incident sulphur nucleus is $E = 10$ GeV per nucleon, with the corresponding Lorentz factor $\gamma = 10.7$, while the gold nucleus is at rest in the laboratory. We assume that the collision is central, that is, the impact parameter is zero. As a result, the problem is axially symmetric and we can perform the calculation in cylindrical coordinates.

The computation was done in the equal-speed reference frame. We used a third-order CH-ENO-LF method on a grid with 600 points along the collision axis direction and 300 points in the radial direction. For the nuclear equation of state we took the model from Ref. [25]. Figure 10 shows snapshots of the baryon density contours at levels $n = 0.1, 0.2, ..., 1.1$ fm^{-3}, taken at the central slice. The baryon density in the initial nuclei (nuclear ground state density) is $n_0 = 0.15$ fm^{-3}.

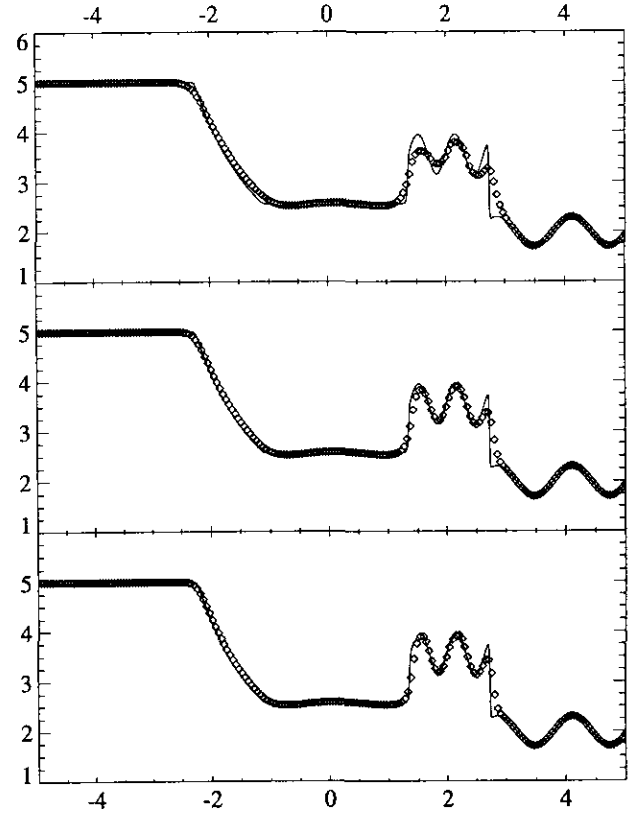

FIG. 4. Perturbed relativistic shock tube problem (Example 2), plot of baryon density with second- (top), third- (center), and fourth-order (bottom) CH-ENO-LLF, 200 points, $\Delta t/\Delta x = 0.5$. The solid line is a third-order CH-ENO-LLF with 2000 points.

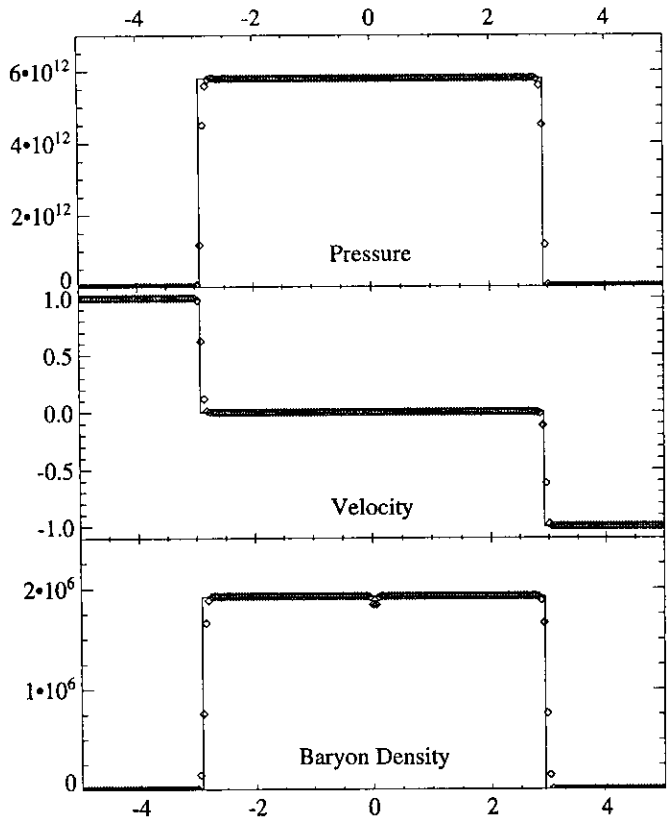

FIG. 5. Ultrarelativistic colliding slab (Example 3), with third-order CH-ENO-LLF, 200 points, $\Delta t/\Delta x = 0.5$. The full line is the exact solution.

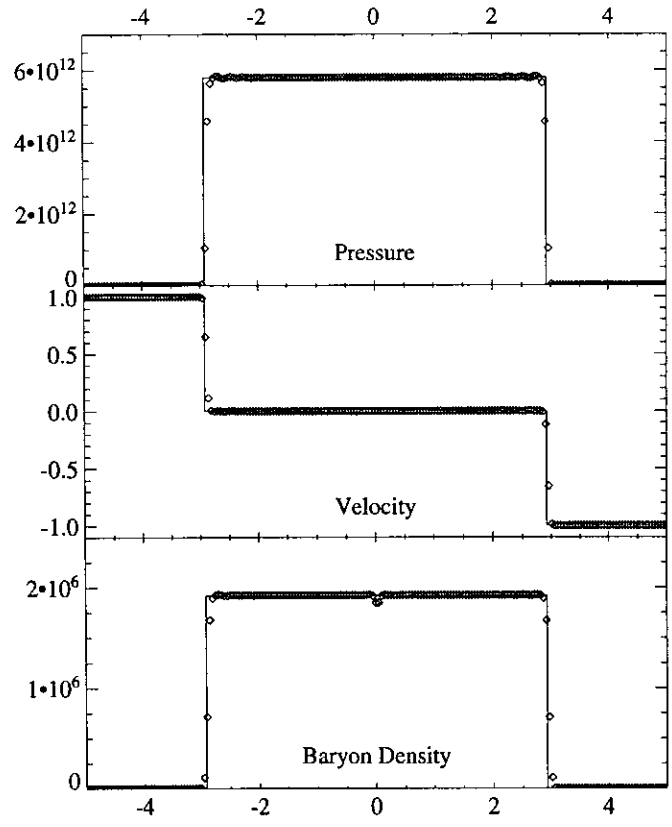

FIG. 6. Same as Fig. 5, except with fourth-order CH-ENO-LLF.

In the initial state, the two spherical nuclei are Lorentz-contracted along the collision axis. After the impact, two compression shock waves form (at $t = 1, 2$), giving rise to a hot and dense region in between. After the outer edge of the sulphur nucleus crosses the shock, a fast expansion of the highly excited matter into vacuum begins ($t = 3, 4$). The same happens when the other shock reaches the outer edge of the gold nucleus ($t = 5, 6$). Note also the two crescents of cold, "spectator" nuclear matter on the sides, surviving until $t = 8$.

Taking the cylindrical calculation as an "exact solution" for comparison, we present in Fig. 11 the results for the same problem computed in Cartesian coordinates, using a third-order CH-ENO-LF on a three-dimensional mesh with $75 \times 75 \times 75$ points. While there is obviously certain loss in the resolution, the global features of the flow are borne out rather well. This gives us confidence in applying three-dimensional ENO schemes for non-central (peripheral) collisions, where we can no longer reduce the dimensionality by using cylindrical coordinates.

All calculations were carried out on a 99 MHz HP9000/735 workstation, where, for example, the 3D nuclear collision problem took less than six hours to complete.

VIII. CONCLUSIONS

We have implemented the high-order, finite differencing ENO schemes to obtain numerical solutions to the equations of relativistic hydrodynamics of a perfect fluid. The present algorithm is applicable to any number of space dimensions and arbitrary equation of state. We have performed numerical tests on schemes up to the fourth-order accuracy in time and space and found consistently good results. This extends results previously established for the Euler gas dynamics. The overshoot-

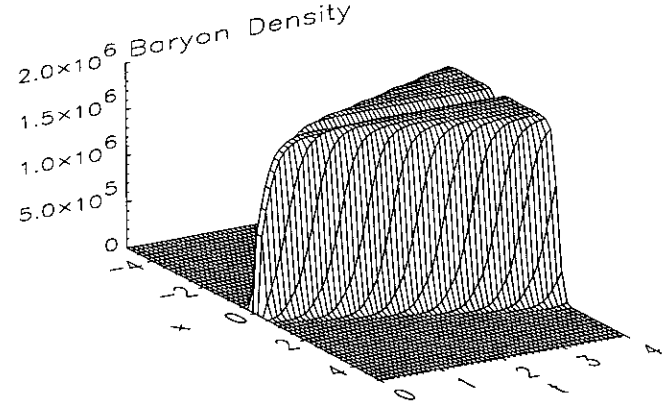

FIG. 7. Surface plot of the baryon density history (colliding slab problem of Example 3), with third-order CH-ENO-LLF with 50 spatial points, 50 time levels, and $\Delta t/\Delta x = 0.5$.

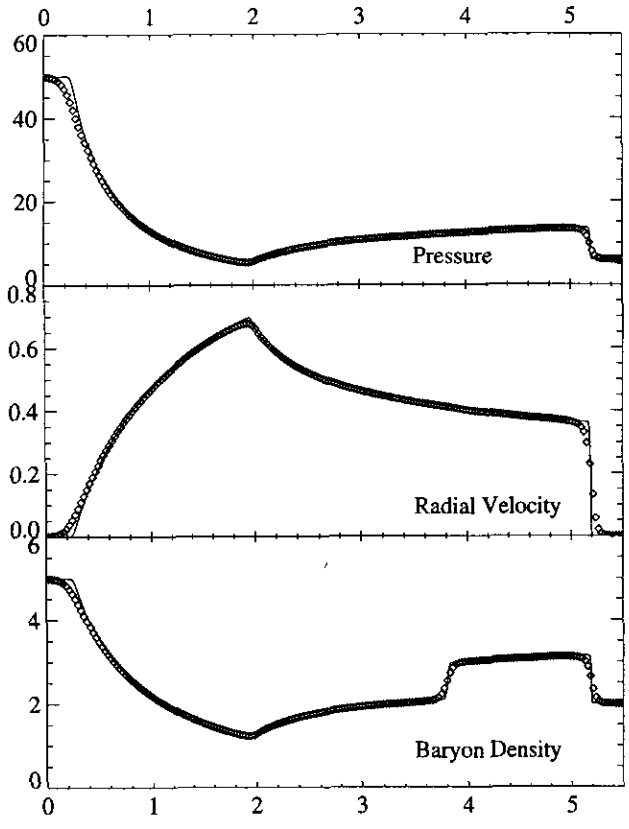

FIG. 8. The "shock cylinder" problem (Example 4). Cut at $x = 0$ is shown for a solution obtained using third-order CH-ENO-LLF in two-dimensional Cartesian coordinates with 200×200 points and $\Delta t/\Delta x = \Delta t/\Delta y = 0.3$. The solid line was computed in polar coordinates using third-order CH-ENO-LLF with 2000 points.

free, non-oscillatory shock handling in the ENO schemes is an asset that is even more important in the relativistic regime. This property is carried out completely from the Newtonian case. Because of the high resolution, we believe that the ENO schemes will provide an effective approach to numerical solu-

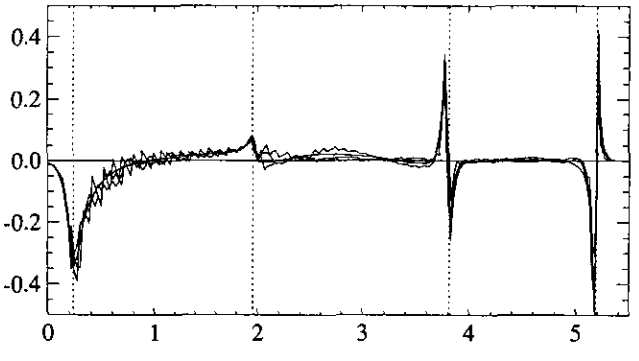

FIG. 9. Error in baryon density in the "shock cylinder" problem (Example 4), for radial cuts done at angles $\theta = 0°$, 20°, 30°, and 45° with respect to the gridlines. All the angles are shown in full line.

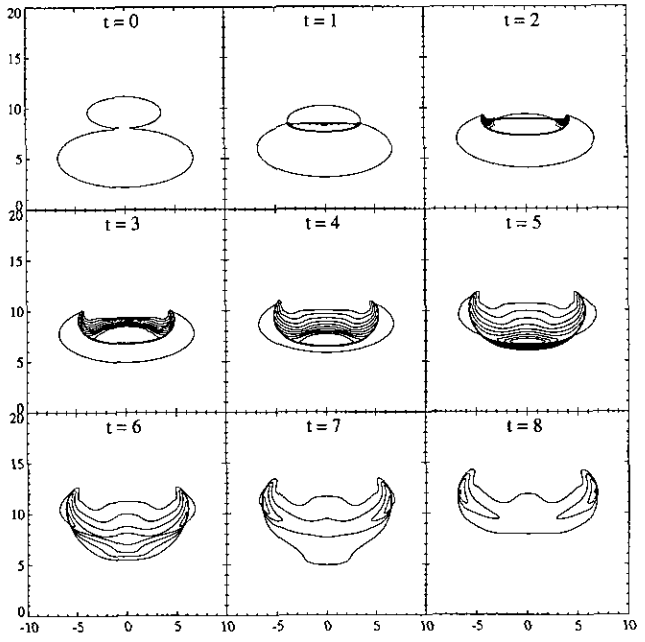

FIG. 10. Central S–Au collision at 10 GeV/A incident energy (Example 5). The results were obtained using third-order CH-ENO-LF in cylindrical coordinates with 600 points along the collision axis (vertical) and 300 points in the radial direction with $\Delta t/\Delta z = \Delta t/\Delta r = 0.3$. The contours of baryon density at levels $n = 0.1, 0.2, \ldots, 1.1$ fm^{-3} are shown, as viewed from an equal-speed frame. The time and distance are given in fm (10^{-15}m).

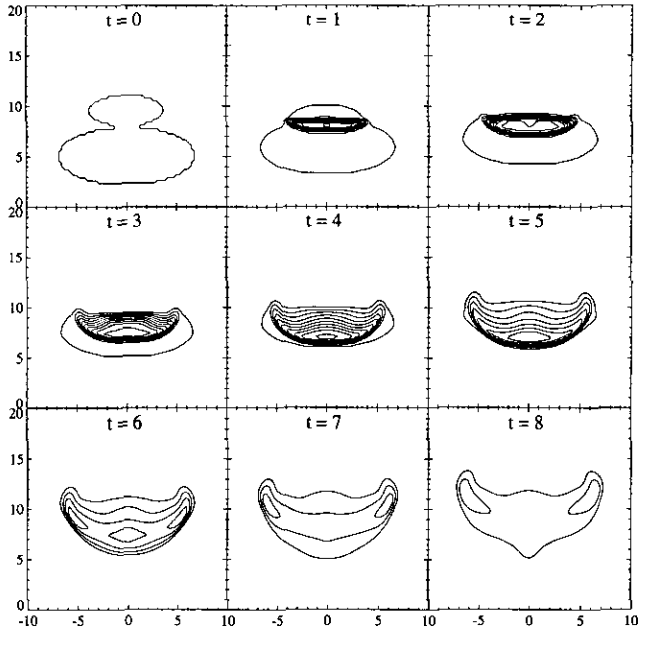

FIG. 11. The same problem as in Fig. 10, computed using third-order CH-ENO-LF in three-dimensional Cartesian coordinates with $75 \times 75 \times 75$ points and $\Delta t/\Delta x_j = 0.2$ for all spatial directions.

tions of hydrodynamic models of relativistic nuclear collisions. The feasibility of such an application is demonstrated by the nuclear collision calculations presented.

We have tested only the "standard" finite differencing ENO scheme by Shu and Osher [11]. Modified algorithms exist, for example, the artificial compression method for sharpening of contact discontinuities by Yang [26], a modified stencil selection algorithm with an adjustable biasing towards central differencing [27], or hybrid methods [28], combining central differencing in the smooth regions with the ENO method near discontinuities. We do not expect any difficulty in implementing these modified schemes, where they are desired.

In closing we remark that extension to a viscous flow or to a flow with heat transfer is by no means a simple matter. In contrast to their Newtonian counterparts, relativistic Navier–Stokes equations are second-order and nonlinear in *time*, as, e.g., terms proportional to $\partial v_j/\partial t$ appear on the right-hand side of (2.3) [16]. An entire new class of difficulties arises with such a situation.

ACKNOWLEDGMENTS

This work has been supported in part by the Natural Sciences and Engineering Research Council of Canada. One of the authors (A.D.) wishes to thank the Ministry of Colleges and Universities of the Province of Ontario for the financial support through an Ontario Graduate Scholarship and Professor Chi-Wang Shu of Brown University for helpful comments and suggestions.

REFERENCES

1. L. P. Csernai, *Introduction to Relativistic Heavy Ion Collisions* (Wiley, Chichester, New York, 1994).
2. I. N. Mishustin, V. N. Russkikh, and L. M. Satarov, *Sov. J. Nucl. Phys.* **54,** 260 (1991).
3. H. Stöcker and W. Greiner, *Phys. Rep.* **137,** 277 (1986).
4. R. B. Clare and D. H. Strottman, *Phys. Rep.* **141,** 177 (1986).
5. W. Schneider *et al., J. Comput. Phys.* **105,** 93 (1993).
6. A. Harten, *J. Comput. Phys.* **49,** 357 (1983).
7. A. Harten, B. Engquist, S. Osher, and S. R. Chakravarthy, *J. Comput. Phys.* **71,** 231 (1987).
8. A. Harten, in *Proceedings, International Conference on Nonlinear Hyperbolic Systems, Saint-Etienne, 1986,* edited by C. Carasso *et al.*, Lecture Notes in Mathematics (Springer-Verlag, Berlin, 1987), p. 23.
9. J. Casper and H. L. Atkins, *J. Comput. Phys.* **105,** 92 (1993).
10. C.-W. Shu and S. Osher, *J. Comput. Phys.* **77,** 439 (1988).
11. C.-W. Shu and S. Osher, *J. Comput. Phys.* **83,** 32 (1989).
12. J. Casper, C.-W. Shu, and H. Atkins, *AIAA J.* **32,** 1970 (1994).
13. C.-W. Shu *et al., Appl. Numer. Math.* **9,** 45 (1992).
14. E. Weinan and C.-W. Shu, *J. Comput. Phys.* **110,** 39 (1994).
15. A. Harten, in *Wave Motion: Theory, Modelling and Computation: Proceedings, Conference in Honor of the 60th Birthday of Peter Lax,* edited by A. J. Chorin and A. J. Majda (Springer-Verlag, New York, 1987), p. 147.
16. L. D. Landau and E. M. Lifshitz, *Fluid Mechanics* (Pergamon, London, 1959).
17. P. D. Lax and B. Wendroff, *Commun. Pure Appl. Math.* **13,** 217 (1960).
18. P. Lax, *Hyperbolic Systems of Conservation Laws and the Mathematical Theory of Shock Waves* (Society for Industrial and Applied Mathematics, Philadelphia, 1972).
19. P. Roe, in *Proceedings, 7th Int. Conf. Numer. Methods Fluid Dyn., Stanford and NASA/Ames, June 1980,* edited by W. C. Reynolds and R. W. MacCormack (Springer-Verlag, New York/Berlin, 1981), p. 354.
20. A. C. Hearn, *REDUCE User's Manual, Version* 3.5, Rand publication CP78, 1993.
21. A. Dolezal, Ph.D. thesis, University of Toronto, 1995.
22. K. Thorne, *Astrophys. J.* **179,** 897 (1973).
23. P. Roe, *J. Comput. Phys.* **43,** 357 (1981).
24. A. Harten, *J. Comput. Phys.* **49,** 151 (1983).
25. A. Dolezal and S. S. M. Wong, *Phys. Lett. B,* **294,** 303 (1992).
26. H. Yang, *J. Comput. Phys.* **89,** 125 (1990).
27. C.-W. Shu, *J. Sci. Comput.* **5,** 127 (1990).
28. E. Harabetian and R. Pego, *J. Comput. Phys.* **105,** 1 (1993).