



A tensor artificial viscosity for SPH

J. Michael Owen

Lawrence Livermore National Laboratory, A Division, MIS L-38, P.O. Box 808, Livermore, CA 94550, USA

Received 10 September 2003; received in revised form 14 November 2003; accepted 19 June 2004

Available online 13 August 2004

Abstract

A tensor generalization of the Monaghan-Gingold artificial viscosity for smoothed particle hydrodynamics (SPH) is formulated. We also develop a new limiter to restrict the effects of the artificial viscosity to the direction of shock travel, while maintaining enough of the artificial viscous dissipation to keep the SPH technique stable. We present a series of test problems comparing the new viscosity formalism with the standard method for fluid flows in two and three dimensions, such as shearing systems undergoing shocks. We also present tests of the conservation and transport of angular momentum, since total angular momentum conservation is not guaranteed with the new viscosity formalism.

© 2004 Elsevier Inc. All rights reserved.

1. Introduction

Smoothed particle hydrodynamics (SPH) was originally developed by Lucy [6] and Gingold and Monaghan [5] for use in modeling astrophysical phenomena, and has since been successfully applied to study a wide variety of phenomena. SPH is a meshless method of modeling hydrodynamics, wherein the material is discretized into a set of nodes or “particles” on which the Lagrangian hydrodynamic equations are solved. SPHs great strengths are its Lagrangian nature (allowing the resolution to follow the mass), its intrinsic robustness (SPH is not susceptible to mesh tangling in complex flows), its natural match with the well established N body schemes for modeling gravitational interactions, and its simplicity. This last point has probably been SPH’s major selling point: in its simplest form is literally as simple to implement SPH in three dimensions as in one. We will not discuss the SPH technique in depth here, but rather refer the reader to the review articles written by Monaghan [9] and Benz [2].

Although SPH’s simplicity is one of its great strengths, some of the simplifying assumptions which are typically made in extending SPH to two and three dimensions can compromise the quality of the solution in

E-mail address: mikeowen@llnl.gov (J.M. Owen).

the presence of multi-dimensional fluid flows. In particular, the most popular artificial viscosity employed in SPH calculations (based on the pioneering work of Monaghan and Gingold [10]) is extended to two and three dimensions as a scalar viscosity, essentially based on the divergence of the velocity. Like many other Lagrangian hydrodynamic schemes, SPH uses an artificial viscosity to spread the shock jump conditions over a few resolution scales, thereby stabilizing the technique in the presence of shocks. Usually a shock transition is locally a highly directional phenomenon, and it is desirable to have the numerical shock treatment take this local directionality into account and react accordingly. A scalar artificial viscosity does not possess the information to discriminate directionality though, and reacts to shocks as an isotropic local pressure increase. A further complication arises in the standard SPH artificial viscosity due to the fact that it is triggered by shearing flows as well as compressional flows. While there has been some work on patching up the standard viscosity in the presence of shears (Balsara [1]), this remains a problem.

In this paper we present a new method of extending the one-dimensional artificial viscosity originally developed by Monaghan and Gingold [10] to two and three dimensions, based on a tensor formalism which allows the viscosity to adapt to the dimensionality of the system being considered. The paper is organized as follows: in Section 2 we present the standard SPH evolution equations used in this work; in Section 3 we briefly summarize the traditional approach to extending the SPH artificial viscosity to higher dimensions; in Section 4 we present the new formalism for a tensor artificial viscosity; and finally Section 5 presents results from several test problems with analytic solutions (including a new problem propagating a shock over a shearing flow), as well as several problems designed to test the angular momentum conservation of the new tensor viscosity.

2. The SPH evolution equations

Before we begin, a word about notation. In this paper we use the convention that English subscripts denote node indices (m_i is the mass of node i), while Greek superscripts denote dimensional indices (r_i^α is the α th component of the position for node i). We use the summation convention for repeated Greek indices: $r_i^\alpha v_i^\alpha = \mathbf{r}_i \cdot \mathbf{v}_i$. We define $\langle F \rangle$ as the SPH interpolated value of F : $\langle F \rangle_i = \sum_j (m_j / \rho_j) F_j W_{ij}$. Consistent with our prior work on adaptive smoothed particle hydrodynamics (ASPH, cf. [15]), we fold the SPH smoothing scale h into our equations by preferring to write expressions in terms of “normalized” coordinates $\eta^\alpha = r^\alpha / h$. We do not present any results in this paper using ASPH because the improvements we seek in the form of the artificial viscosity are orthogonal to the resolution improvements addressed by ASPH. Everything presented in this paper is equally applicable to both SPH and ASPH, and by writing our equations with this convention both cases are covered.

The SPH discretizations of the Lagrangian conservation laws which we employ are

$$\rho_i = \sum_j m_j W_{ij}, \quad (1)$$

$$\frac{Dv_i^\alpha}{Dt} = - \sum_j m_j \left[\left(\frac{P_i}{\rho_i^2} + \frac{P_j}{\rho_j^2} \right) \frac{\partial W_{ij}}{\partial x^\alpha} + \Pi_{ij} \frac{\partial W_{\Pi_{ij}}}{\partial x^\alpha} \right], \quad (2)$$

$$\frac{Du_i}{Dt} = \sum_j m_j \left(\frac{P_i}{\rho_i^2} v_{ij}^\alpha \frac{\partial W_{ij}}{\partial x^\alpha} + \Pi_{ij} v_{ij}^\alpha \frac{\partial W_{\Pi_{ij}}}{\partial x^\alpha} \right), \quad (3)$$

where ρ is the mass density, m is the mass, v^α is the velocity, P is the pressure, and u is the specific thermal energy. Π is the artificial viscosity (written here appropriately for a scalar viscosity), and has units of P/ρ^2 . Therefore, Π is related to the artificial viscous pressure Q by $Q = \rho^2 \Pi$. W is the interpolation kernel; note

that in these equations we have explicitly separated out the interpolation kernel used with the artificial viscosity as W_{Π} . Usually the same interpolation kernel is used throughout these equations, but there is no reason that this need be so. The ij subscripts on Π_{ij} , W_{ij} , and $W_{\Pi_{ij}}$ denote that these terms are explicitly symmetrized for interactions between any pair of nodes i and j , so that $\Pi_{ij} = (\Pi_i + \Pi_j)/2$ and $W_{ij} = (W_i + W_j)/2$. This explicit pair-wise symmetrization of the forces in the momentum equation (2) enforces the global conservation of both linear and angular momentum.

In addition to the equations evolving the physical state variables, we also evolve the smoothing scale associated with each node inversely with the mass density, according to

$$\frac{Dh_i}{Dt} = -\frac{h_i}{v\rho_i} \sum_j m_j v_{ij}^\beta \frac{\partial W_{ij}}{\partial x^\beta} \tag{4}$$

in v dimensions.

We use two interpolation kernels in this paper. The third order B-spline typically used in SPH (see for instance [9]) is given by

$$W_3(\eta) = A_v \begin{cases} 1 - \frac{3}{2}\eta^2 + \frac{3}{4}\eta^3, & 0 \leq \eta \leq 1, \\ \frac{1}{4}(2 - \eta)^3, & 1 < \eta \leq 2, \\ 0, & \eta > 2, \end{cases} \tag{5}$$

where A_v is a normalization constant defined so that the volume integral $\int dV W_3 = 1$ in v dimensions. This cubic B-spline is only one member of the family of B-spline kernels (see [16] for the general definition), and another interesting kernel is the fifth order B-spline given by

$$W_5(\eta) = A_v \begin{cases} \frac{11}{20} - \frac{1}{2}\eta^2 + \frac{1}{4}\eta^4 - \frac{1}{12}\eta^5, & 0 \leq \eta \leq 1, \\ \frac{17}{40} + \frac{5}{8}\eta - \frac{7}{4}\eta^2 + \frac{5}{4}\eta^3 - \frac{3}{8}\eta^4 + \frac{1}{24}\eta^5, & 1 < \eta \leq 2, \\ \frac{81}{40} - \frac{27}{8}\eta + \frac{9}{4}\eta^2 - \frac{3}{4}\eta^3 + \frac{1}{8}\eta^4 - \frac{1}{120}\eta^5, & 2 < \eta \leq 3, \\ 0, & \eta > 3. \end{cases} \tag{6}$$

The penalty for using W_5 is computational expense. Since W_5 extends to $\eta = 3$ rather than $\eta = 2$ like W_3 , using W_5 means that each SPH node samples more neighbors than are required for the W_3 kernel. This difference can be substantial, particularly in three dimensions where the number of neighbors increases by roughly $(3/2)^3$, implying a corresponding threefold increase in run time. Nonetheless, we find that in many problems the quality of the answer improves with the use of the higher order kernel, so in some of the tests presented in this work (particularly those in which we are examining symmetry in detail) we use $W = W_5$ for the standard interpolation kernel, while we continue to use $W_{\Pi} = W_3$ since shock interactions in general do not benefit from higher order interpolation. In other cases we make the more standard choice of $W = W_{\Pi} = W_3$. For each test case we state which kernel choices are made.

3. The standard SPH artificial viscosity

The most commonly used form of artificial viscosity in SPH is based on the pioneering work of Monaghan and Gingold [10], wherein the authors derive a 1D pair-wise node form for the viscosity based on the bulk and Von Neumann–Richtmyer viscosities

$$\rho\Pi = \begin{cases} -C_l h c_s \frac{dv_x}{dx} + C_q h^2 \left(\frac{dv_x}{dx}\right)^2, & \frac{dv_x}{dx} < 0, \\ 0 & \frac{dv_x}{dx} \geq 0, \end{cases} \tag{7}$$

where C_l and C_q are constants, h the smoothing scale, dv_x/dx the (one-dimensional) gradient of the velocity, and c_s the sound speed. The pair-wise definition for this viscosity which these authors arrived at is

$$\rho_i \Pi_i = \begin{cases} -C_l c_i \mu_i + C_q \mu_i^2, & \mu_i < 0, \\ 0, & \mu_i \geq 0, \end{cases} \quad (8)$$

$$\mu_i = \frac{v_i \eta_i}{\eta_i^2 + \epsilon^2}, \quad (9)$$

where $v_{ij} = v_i - v_j$, $\eta_i = (r_i - r_j)/h_i$, and ϵ^2 is a small constant to prevent division by zero problems. Essentially μ_i is a pair-wise low order method of estimating dv/dx . The most important factor in this viscosity's successful application in SPH is its pair-wise definition, which ensures the viscosity will react to and damp individual velocity differences between nodes. A straightforward implementation of Eq. (7) using the SPH machinery to estimate $\langle dv_x/dx \rangle$ is possible, but because the SPH technique enforces smoothness on scales of order of the smoothing scale, such a viscosity will be insensitive to noise in the velocity on the scales $\leq h$ and therefore will fail to damp out high frequency signals on such small scales. A shock is essentially an unresolved step function in properties of the system, and as such is an excellent source of signals at or below the resolution scale. The primary purpose of the artificial viscosity is to smoothly spread out shocks into resolvable features, and this kind of failure is unacceptable.

As is demonstrated in Monaghan and Gingold [10], one-dimensional SPH simulations using the formalism of Eqs. (8) and (9) perform quite well, capturing shock transitions in $\sim 2h$ length scales and effectively preventing node interpenetration (where SPH nodes unphysically cross one and other). However, the question of how to best extend this formalism to two- and three-dimensional is not clear. Most SPH implementations simply replace the one-dimensional gradient dv_x/dx with the velocity divergence $\nabla \cdot \mathbf{v}$, which is accomplished by re-expressing μ as

$$\mu_i = \frac{v_{ij}^x \eta_i^x}{\eta_i^2 + \epsilon^2}. \quad (10)$$

This results in a scalar definition for the SPH artificial viscosity, and this form possesses a number of desirable properties: in the case of a simple planar shock it effectively reproduces the one-dimensional results; it is Galilean invariant; it maintains excellent symmetry properties (spherical problems remain spherical, etc.); it vanishes for rigid body rotation; when properly symmetrized it preserves linear and angular momentum to machine precision; it is conceptually simple and easy to code for any number of dimensions.

However, this particular form for the multi-dimensional viscosity also has some serious shortcomings. Since the viscosity per pair is a scalar, it always acts along the direction defined by the gradient of the interpolation kernel $\partial W_{ij}/\partial x^\alpha$. In SPH, this means the force acts along the line connecting the pair of nodes. However, in multiple dimensions this is not always the appropriate answer: preferably the viscosity should act in the direction of the shock (or at least the compression). Consider a spherical radial inflow problem, such as a cloud of gas collapsing due to gravitational instability or the spherical Noh problem. In this case, there should be a radial shock forming from the point of convergence and moving outward in Lagrangian coordinates. However, in an SPH simulation of such a flow, nodes which are flowing in along different radial lines will interact via the artificial viscosity and exert forces which are not aligned with the direction of convergence. This results in an unphysical heating of the gas during inflow, and distorts the solution. Furthermore, according to Eqs. (8) and (10), Π is positive for any pair of nodes for which $v_{ij}^\alpha r_{ij}^\alpha < 0$, i.e., nodes i and j are approaching one and other. This implies that the viscosity can be triggered by pure shearing flows, even though for such flows $\partial v^\alpha/\partial x^\alpha = 0$ and ideally there should be no artificial viscosity. This shearing activation of the viscosity is a serious problem for many astrophysical situations, since typically gas in rotational support with a gravitational field is undergoing rotational shear. In this situation the artificial

viscosity will damp the rotationally driven shear, resulting in an unphysical transport of linear and angular momentum.

In order to deal with the activation of the viscosity in shearing flows, Balsara [1] (also independently developed in [8]) proposed that the viscosity be multiplied by the correction term

$$f_{B_i} = \frac{|\langle \nabla \cdot \mathbf{v} \rangle_i|}{|\langle \nabla \cdot \mathbf{v} \rangle_i| + |\langle \nabla \times \mathbf{v} \rangle_i| + \epsilon^2 c_i / h_i}, \tag{11}$$

where c_i is the sound, and ϵ is a small number to avoid division by zero problems. As is noted in Balsara [1], the f_B correction term is an imperfect solution to the shearing viscosity problem: it is in fact only correct for either purely shear or completely shear free flows. In the pure shear case, f_B is zero and the viscosity is turned off. In the shear free case, f_B is one and the viscosity is unaffected. However, in any flow with both shearing and compressional components, f_B simply reduces the magnitude of the viscosity. In this situation the artificial viscosity is still allowed to place an unphysical drag on the shear while simultaneously the compressional activation of the viscosity is suppressed, potentially degrading the effectiveness of the viscosity in smoothly handling shock conditions. Nonetheless, f_B is necessary for any realistic use of the standard artificial viscosity in multi-dimensional SPH simulations, and we therefore employ it in all the test cases we present here of the standard viscosity. In order to maintain the pair-wise symmetry of interactions between nodes, we define the shear corrected symmetric viscosity as $\Pi_{ij} = (f_{B_i} \Pi_i + f_{B_j} \Pi_j) / 2$.

4. A tensor form of the artificial viscosity

The problems with the standard SPH artificial viscosity are due to the fact that for any given pair of interacting nodes it uses only the relative positions and velocities of those nodes in determining the viscous force which acts between them. However, in two or three dimensions when a pair of nodes are approaching one and other at some angle other than head on, there is no way to determine whether that closing represents a true compression or simply some shearing flow. Worse still, if there is a compression there is no way to determine the true orientation of that compression, and therefore the direction in which the viscosity should apply deceleration. The shear correction in Eq. (11) uses information from nodes surrounding i and j to try and detect shearing flows and moderate the shearing problem, but does not address the orientation problem. In order to provide a framework for specifying the directionality of the viscosity as well as deal more effectively with shearing flows, we will consider a tensor generalization of the one-dimensional viscosity: $\Pi \rightarrow \Pi^{\alpha\beta}$. The viscous acceleration and energy terms in the Lagrangian evolution equations appropriate for a tensor artificial viscosity are

$$\frac{Dv^\alpha}{Dt} = -\rho \frac{\partial \Pi^{\beta\alpha}}{\partial x^\beta}, \tag{12}$$

$$\frac{Du}{Dt} = -\rho \Pi^{\alpha\beta} \frac{\partial v^\beta}{\partial x^\alpha}. \tag{13}$$

The SPH discretizations for momentum and energy (Eqs. (2) and (3)) can be rewritten accordingly as

$$\frac{Dv_i^\alpha}{Dt} = - \sum_j m_j \left[\left(\frac{P_i}{\rho_i^2} + \frac{P_j}{\rho_j^2} \right) \frac{\partial W_{ij}}{\partial x^\alpha} + \Pi_{ij}^{\beta\alpha} \frac{\partial W_{\Pi_{ij}}}{\partial x^\beta} \right], \tag{14}$$

$$\frac{Du_i}{Dt} = \sum_j m_j \left(\frac{P_i}{\rho_i^2} v_{ij}^\alpha \frac{\partial W_{ij}}{\partial x^\alpha} + \Pi_{ij}^{\alpha\beta} v_{ij}^\beta \frac{\partial W_{\Pi_{ij}}}{\partial x^\alpha} \right). \tag{15}$$

Note that the viscous acceleration term in Eq. (14) is no longer necessarily aligned with the radial vector connecting an interacting pair of nodes. This is a desirable feature of the formalism since the acceleration should be applied in the direction of compression; however, this also implies that we have given up the guarantee of angular momentum conservation since the forces are no longer radial. Proper symmetrization of the pair-wise forces does ensure total linear momentum conservation, but now we must keep an eye on the total angular momentum evolution since conservation can no longer be guaranteed.

The form we propose for the artificial viscosity is

$$\Pi_i^{\alpha\beta} = -C_l h_i c_i \sigma_{ij}^{\beta\alpha} + C_q h_i^2 \sigma_{ij}^{\alpha\gamma} \sigma_{ij}^{\gamma\beta}, \quad (16)$$

where $\sigma^{\alpha\beta}$ is the gradient of the velocity $\partial v^\alpha / \partial x^\beta$, modulo the removal of any expansive components. Eq. (16) is essentially a direct tensor extension of the one-dimensional bulk and Von Neumann–Richtmyer viscosities in Eq. (7). Note that the linear term is an SPH form of the artificial viscosity proposed in Shashkov [3] for mesh based Lagrangian schemes.

It is important that we use a pair-wise definition for $\sigma_{ij}^{\alpha\beta}$, since as discussed in Section 3, it is the sensitivity to individual differences in nodal velocities that makes the standard SPH viscosity so successful. There is not enough information available from just the relative velocities and positions of a pair of nodes to determine all the components of $\partial v^\alpha / \partial x^\beta$, but for a given pair of nodes i and j we can view $v_{ji}^\alpha / x_{ji}^\beta$ as a slice through the full velocity gradient tensor along the line connecting the two nodes (recall that $v_{ji}^\alpha = v_j^\alpha - v_i^\alpha$ and $x_{ji}^\beta = x_j^\beta - x_i^\beta$). With this understanding, we use the following prescription to determine $\sigma_{ij}^{\alpha\beta}$ for a given node i interacting with a specific neighbor j .

- (1) Determine an average background $\overline{(\partial v^\alpha / \partial x^\beta)}_i$ at node i . There are several possible choices we could use; Appendix A gives the details of a pair-wise summed definition we use for the examples in this paper.
- (2) Find the rotational transformation $R_\sigma^{\alpha\beta}$, such that the x' axis in this rotated frame is aligned with the r_{ji}^α vector.¹
- (3) In this rotated frame, $v_{ji}^{\prime\alpha} / |r_{ji}|$ represents a pair-wise definition for the first column of $(\partial v^\alpha / \partial x^\beta)'$. We therefore construct a pair-wise $(\partial v^\alpha / \partial x^\beta)'_{ij}$ by replacing the first column of the background $\overline{(\partial v^\alpha / \partial x^\beta)}_i$ with $v_{ji}^{\prime\alpha} / |r_{ji}|$, and then rotate the result back to the lab frame as follows

$$v_{ji}^{\prime\alpha} = R_\sigma^{\alpha\beta} v_{ji}^\beta, \quad (17)$$

$$\overline{\left(\frac{\partial v^\alpha}{\partial x^\beta}\right)}'_i = R_\sigma^{\alpha\gamma} \overline{\left(\frac{\partial v^\gamma}{\partial x^\psi}\right)}_i R_\sigma^{\beta\psi}, \quad (18)$$

$$\left(\frac{\partial v^\alpha}{\partial x^\beta}\right)'_{ij} = \frac{v_{ji}^{\prime\alpha}}{|r_{ji}|} \delta^{1\beta} + \overline{\left(\frac{\partial v^\alpha}{\partial x^\beta}\right)}'_i (1 - \delta^{1\beta}), \quad (19)$$

$$\left(\frac{\partial v^\alpha}{\partial x^\beta}\right)_{ij} = R_\sigma^{\gamma\alpha} \left(\frac{\partial v^\gamma}{\partial x^\psi}\right)'_{ij} R_\sigma^{\psi\beta}, \quad (20)$$

where $\delta^{\alpha\beta}$ is the Kronecker delta function. Eqs. (18) and (20) are the similarity transforms rotating $\overline{(\partial v^\alpha / \partial x^\beta)}_i$ and $(\partial v^\alpha / \partial x^\beta)'_{ij}$ to/from the rotated frame, respectively, where we have explicitly made use of the fact that $(R_\sigma^{\alpha\beta})^{-1} = R_\sigma^{\beta\alpha}$.

¹ We use $'$ here to denote quantities in the rotated frame. Also note that the σ subscript in $R_\sigma^{\alpha\beta}$ is not an index, but rather distinguishes this rotational transformation.

- (4) In order to construct $\sigma_{ij}^{\alpha\beta}$ we remove any expansive components from $(\partial v^\alpha/\partial x^\beta)_{ij}$. This is equivalent to requiring a convergent flow in one dimension: $dv_x/dx < 0$. We begin by decomposing $(\partial v^\alpha/\partial x^\beta)_{ij}$ into its symmetric and antisymmetric components $S_{ij}^{\alpha\beta}$ and $A_{ij}^{\alpha\beta}$

$$S_{ij}^{\alpha\beta} = \frac{1}{2} \left[\left(\frac{\partial v^\alpha}{\partial x^\beta} \right)_{ij} + \left(\frac{\partial v^\beta}{\partial x^\alpha} \right)_{ij} \right], \tag{21}$$

$$A_{ij}^{\alpha\beta} = \frac{1}{2} \left[\left(\frac{\partial v^\alpha}{\partial x^\beta} \right)_{ij} - \left(\frac{\partial v^\beta}{\partial x^\alpha} \right)_{ij} \right]. \tag{22}$$

The information about any expansion or compression of the fluid volume is contained in the symmetric component $S_{ij}^{\alpha\beta}$. Identify the rotational transformation $R_S^{\alpha\beta}$ which diagonalizes $S_{ij}^{\alpha\beta}$

$$D_{ij}^{\alpha\beta} = R_S^{\alpha\gamma} S_{ij}^{\gamma\psi} R_S^{\beta\psi}, \tag{23}$$

limit the diagonal elements of $D_{ij}^{\alpha\beta}$ to compressional components only

$$\tilde{D}_{ij}^{\alpha\alpha} = \min(0, D_{ij}^{\alpha\alpha}), \tag{24}$$

(note that in Eq. (24) we are not using the summation convention, each component is limited separately). Finally, reconstruct the expansion free gradient tensor

$$\sigma_{ij}^{\alpha\beta} = R_S^{\alpha\gamma} \tilde{D}_{ij}^{\gamma\psi} R_S^{\beta\psi} + A_{ij}^{\alpha\beta}. \tag{25}$$

In order to ensure that the viscous term is always dissipative, we only apply the viscosity if nodes i and j are approaching one and other ($v_{ij}^\alpha r_{ij}^\alpha < 0$) and if the viscous work term in Eq. (15) is positive. In Appendix C we give the explicit formulas for $R_\sigma^{\alpha\beta}$ and $R_S^{\alpha\beta}$ in two and three dimensions. Note that using the prescription described here, this viscosity reduces exactly to the standard SPH viscosity given in Eqs. (7)–(9) in one dimension.

For the moment using the approximation that $\sigma^{\alpha\beta} \sim \partial v^\alpha/\partial x^\beta$, it is simple to demonstrate that this viscosity has several desirable properties. For instance, consider a two-dimensional flow field of the form

$$\mathbf{v} = \begin{pmatrix} f(x) + g(y) \\ 0 \end{pmatrix}, \tag{26}$$

so that $g(y)$ represents a purely shearing component of the velocity. Plugging the gradient of this velocity field into Eqs. (16) and (12), the corresponding viscous acceleration is (neglecting all terms other than the velocity derivatives)

$$\frac{D\mathbf{v}}{Dt} \propto \begin{pmatrix} f'' + 2f'g'' + f'g'' \\ 0 \end{pmatrix} \tag{27}$$

(where ' indicates a spatial derivative). Several properties of the tensor viscosity are immediately obvious from this simple analysis. First, in the presence of a planar shear the tensor viscosity will not generate any acceleration away from the plane of the shear, unlike the scalar form of the previous section. This sort of failure is referred to as “mode conversion” in Margolin [7]. This is an important property for systems in which the angular momentum evolution is significant, because mode conversion is essentially a spurious local transport of linear and angular momentum (irrespective of the fact that the total momenta are conserved.) Second, if $f(x) = 0$ so there is no compressive term, the tensor viscosity automatically goes to zero despite the presence of the shearing term $g(y)$. If there is a compressional signal and the viscosity is active, the $f'g''$ term in the x acceleration will cause the viscosity to damp the shear by decelerating against it in the

direction of the shear velocity component. This shear contribution to the acceleration is an error, but nonetheless the tensor viscosity is better behaved in the presence of shears than the standard scalar form, without the use of an external limiter such as f_B . These results are not dependent on the alignment of the shear with the x direction as we have considered here for simplicity, but are true for planar shear at any angle. It is also possible to demonstrate analytically that the tensor viscosity will not distort cylindrically or spherically symmetric flow fields, such as radial inflow or outflow.

4.1. A limiter to restrict overuse of the viscosity

The tensor artificial viscosity gives us the ability to discriminate different compressional signals in constructing the viscous acceleration, but we do not always want to have the viscosity active for any compression. For instance, in the case of homogeneous radial inflow we do not want the viscosity to activate and heat the system before it shocks at the point of convergence, since such unphysical preheating of the fluid will distort the results. Ideally we would like to restrict the artificial viscosity to be active only across shocks. One of the strong signatures of a shock is that the velocity divergence diverges as we approach a shock: $\partial v^\alpha / \partial x^\alpha \rightarrow -\infty$. Therefore, the gradient of the velocity divergence should point directly away from a shock, and we can use $-\nabla(\nabla \cdot \mathbf{v}) = -\partial^2 v^\beta / (\partial x^\alpha \partial x^\beta)$ as a good indicator of the direction toward any local shock. We use this idea to construct a directional limiter as follows. Measure a local value at node i for $\partial^2 v^\beta / (\partial x_\alpha \partial x_\beta)_i$ (the details of how to do this are discussed in [Appendix B](#).) Define a unit vector \hat{s}_i^α

$$\mathbf{g}_i^\alpha \equiv \frac{\partial^2 v^\beta}{\partial x^\gamma \partial x^{\beta_i}}, \quad (28)$$

$$\hat{s}_i^\alpha = \frac{\mathbf{g}_i^\alpha}{(\mathbf{g}_i^\beta \mathbf{g}_i^\beta)^{1/2} + \epsilon c_i / h_i^2}, \quad (29)$$

where ϵ is a small number to prevent division by zero errors. We then form a limiter for node i as the dyadic product of \hat{s}_i^α

$$f_i^{\alpha\beta} = \hat{s}_i^\alpha \hat{s}_i^\beta. \quad (30)$$

The limiter $f_i^{\alpha\beta}$ is used to define the fully symmetrized tensor viscosity as

$$\Pi_{ij}^{\alpha\beta} = \frac{1}{2} \left(f_i^{\alpha\gamma} \Pi_{ij}^{\gamma\beta} + f_j^{\alpha\psi} \Pi_{ij}^{\psi\beta} \right). \quad (31)$$

The symmetrized, limited viscosity of Eq. (31) is the expression we use in the SPH momentum and energy equations (14) and (15).

Applying the dyadic product of \hat{s}_i^α to $\Pi_{ij}^{\alpha\beta}$ is equivalent to taking the dot product of the viscous acceleration and \hat{s}_i^α , and applying the result in the direction of \hat{s}_i^α . Therefore, Eq. (31) ensures that only the components of the viscosity that are aligned with $\nabla(\nabla \cdot \mathbf{v})$ are allowed to be active.

5. Test problems

5.1. Cylindrical and spherical converging Noh problem

The Noh shock tests are a classic and notoriously difficult set of hydrodynamic test problems, following the self-similar evolution of an infinite strength shock [14]. An initially pressureless fluid is set in uniform motion converging upon the origin of a planar, cylindrical, or spherical coordinate system. The planar case consists of two streams of gas colliding at the origin, resulting in a pair of planar shocks propagating back

up the colliding streams of material. In the cylindrical and spherical cases the fluid is converging radially upon the origin, resulting in expanding cylindrical and spherical shock fronts, respectively. The tests presented here use a $\gamma = 5/3$ gamma-law gas with unit initial mass density and unit radial motion: $\rho_0 = 1$, $v_{r_0} = -1$. For these conditions the shock speed is $v_s = 1/3$, and the post-shock density enhancement is $\rho_s = 4$ (planar), $\rho_s = 16$ (cylindrical), and $\rho_s = 64$ (spherical).

Fig. 1 presents the results of modeling the cylindrical Noh problem at time $t = 0.6$ (corresponding to a shock position of $r_s = 0.2$) with both the standard SPH viscosity (“Scalar Q”) and the new tensor formalism (“Tensor Q”). The simulations are performed in the $(x \geq 0, y \geq 0)$ two-dimensional quadrant, with reflecting boundary conditions along the $x = 0$ and $y = 0$ lines. We initialize the SPH nodes on concentric rings with constant mass nodes, arranged so that an initial mass density of $\rho_0 = 1$ is established out to a radius of $r_{\max} = 1$. By establishing the nodes on concentric rings around the origin we can easily track how well the cylindrical symmetry of the system is maintained, since the nodes on a particular ring should all remain at the same radius as they move inward and pass through the shock, at which point they stagnate. We seed 60 such rings out to $r_{\max} = 1$ (with a corresponding azimuthal spacing equal to the radial spacing, resulting in a total of 2799 nodes) and initialize the SPH smoothing scales h such that there are

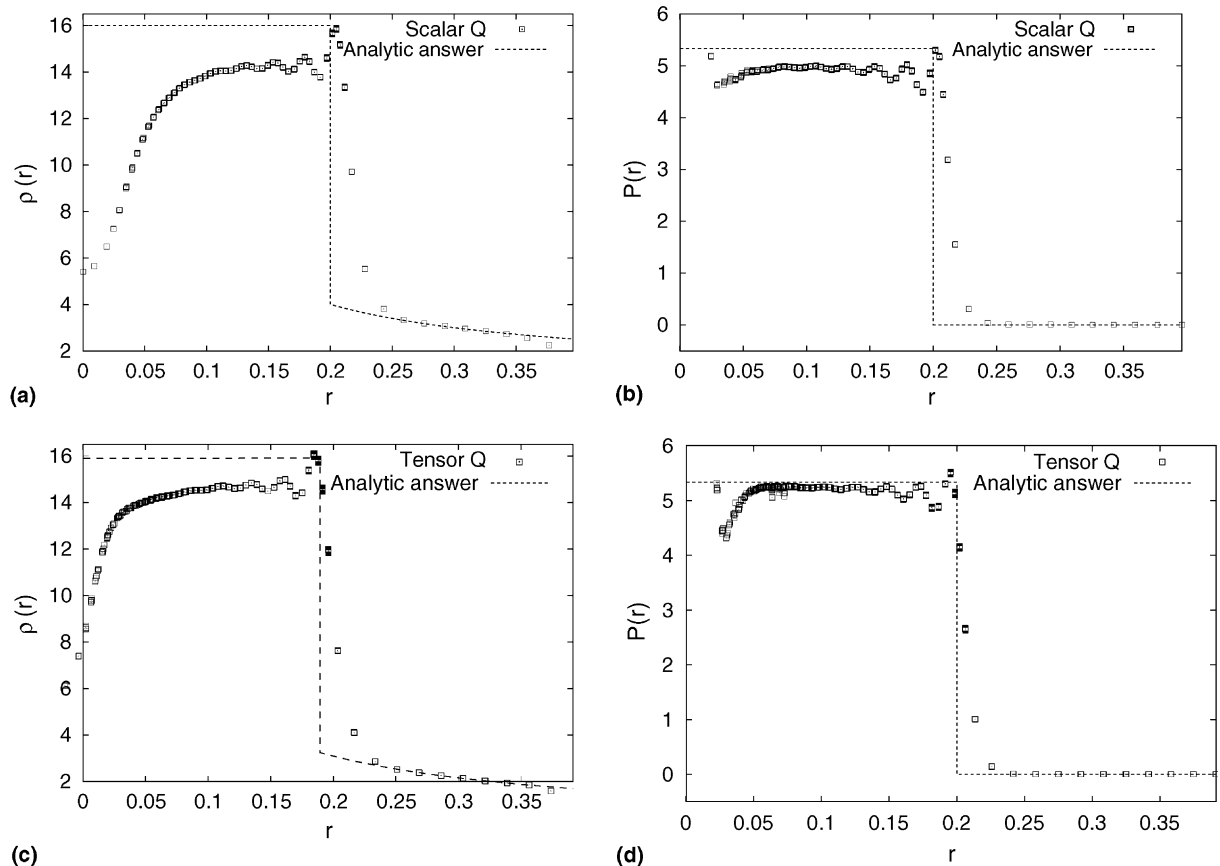


Fig. 1. Radial profiles of the mass density and pressure in the cylindrical Noh problem compared with the analytic answer. (a) is the mass density and (b) is the pressure using the standard (scalar) SPH artificial viscosity, while (c) and (d) are the corresponding results using the tensor viscosity.

three nodes per h . We therefore effectively only have 20 resolution scales in the radial direction in this problem, which is why the mass density profiles are somewhat low. We could of course run a much more highly resolved problem, but in real simulations we typically have to characterize structures with even worse resolution than this, so it is interesting to consider how the technique behaves with marginal resolution. Additionally, since we are examining the symmetry of this system in detail we use the fifth order B spline for the interpolation kernel while continuing to use the third order cubic B Spline for the artificial viscosity: $W = W_5$ and $W_{II} = W_3$. In this and all subsequent tests we set the linear and quadratic coefficients in the viscosities (Eq. (8) and (16)) to $C_l = 1$ and $C_q = 3/4$.

In Fig. 1 we plot all of the nodes vs. radius, so that each of the points in these plots are actually the result of overlapping many SPH nodes at each radius. For instance, each of the 3–4 apparent plotted points in the shock transition at $r_s \sim 0.2$ actually consist of roughly 50 overlapping points. It is clear that both the standard SPH viscosity and the new tensor viscosity do an excellent job of maintaining the symmetry of the problem. On the downside both of these simulations show the so-called “wall heating” problem near the origin, where the fluid becomes overheated and stagnates prematurely, resulting in a temperature overshoot and density undershoot at the origin. This problem is well known, and as Noh [14] points out can be dealt through the introduction of an artificial heat conduction. We do not consider such remedies here though, and simply tolerate the wall heating. In both the density and pressure plots, it is clear that the results with the tensor viscosity do a better job of matching the analytic post shock conditions. The primary reason for this improvement is the $\nabla(\nabla \cdot \mathbf{v})$ limiter described in Section 4.1. In the pre-shock region the fluid is undergoing a homologous, non-shocking compression due to the convergent nature of the flow. Both the scalar and tensor viscosities pick up this motion and seek to decelerate the fluid against this apparent azimuthal compression. However, the limiter detects that the velocity divergence is aligned radially toward the origin and suppresses any acceleration components that are not aligned with this radial direction. Since the tensor viscosity does not activate in the radial direction until the shock is encountered, the viscosity is effectively deactivated in the pre-shock compression and the gas is not prematurely heated. This is why the tensor viscosity results show a slightly higher compression of the post-shock material, and calculate the post shock pressure correctly.

The improvements in the post-shock values for the density and pressure in the spherically symmetric Noh problem (Fig. 2) realized with the tensor viscosity are even more pronounced, though the symmetry suffers somewhat near the origin. In these simulations we initialize the SPH nodes on a uniform lattice of $40 \times 40 \times 40$ points in the unit box between (0,0,0) and (1,1,1) and throw away any points with $r > 1$, leaving a total of 33512 nodes. For efficiency we initialize two particles per smoothing scale, and use the cubic interpolation kernel $W = W_{II} = W_3$. We place reflecting boundary conditions on the $x = 0$, $y = 0$, and $z = 0$ planes, thereby allowing us to model one quadrant of a converging sphere of gas. Although the SPH nodes are seeded on a very different geometry than that of the spherically converging system we are modeling, the tightness of the radial profiles in Fig. 2 (where we have again plotted all 33512 nodes against radius) indicate both viscosities are doing an excellent job of maintaining the physical symmetry of the system. The slight broadening of the tensor viscosity profiles in the post shock region seem to be due to slight inaccuracies in calculating $(\partial v^\alpha / \partial x^\beta)_i$ and $\nabla(\nabla \cdot \mathbf{v})_i$. It is clear that the $\nabla(\nabla \cdot \mathbf{v})$ limiter is playing an important role in the tensor viscosity results, preventing the preheating of the infalling gas and allowing the simulation to better match the analytic shock position and post-shock profiles for the mass density and pressure. The spherical pre-shock compression is stronger than the cylindrical compression in the previous simulations, which is why the unphysical pre-shock heating is stronger in the standard scalar viscosity simulation. This error results in the shock being in the wrong position and the corresponding post shock compression being too low.

Comparison of Figs. 1 and 2 shows that the results with the tensor viscosity for the three-dimensional spherically symmetric problem are very close to those of the two-dimensional cylindrical problem: the spherical and cylindrical post-shock density and pressure profiles approach the analytic solutions in similar ways, and the shock position is accurately captured. This is what we expect of the tensor formalism, since it has the flexibility to adapt to the dimensionality of the problem. The results using the scalar viscosity suffer with




Fig. 2. Radial profiles of the mass density and pressure in the spherical Noh problem compared with the analytic answer. (a) is the mass density and (b) is the pressure using the standard (scalar) SPH artificial viscosity, while (c) and (d) are the corresponding results using the tensor viscosity.

increasing dimensionality because the scalar formalism only knows about the divergence of the velocity, and therefore does not have enough information to discriminate any directionality or dimensionality about a compression. It is also worth pointing out that we would expect that applying the $\nabla(\nabla \cdot \mathbf{v})$ limiter to the standard viscosity will not result in as much improvement as we see when it is used with the tensor viscosity. The acceleration due to the scalar viscosity is always aligned with the line connecting any pair of interacting nodes. Therefore, for nodes interacting at some arbitrary angle between 0° and 90° from the radial direction, the acceleration has some non-zero component in the radial direction and the limiter will not deactivate this component. Therefore, there would still be some amount of unphysical heating of the pre-shock material when using the standard scalar viscosity with the $\nabla(\nabla \cdot \mathbf{v})$ limiter.

5.2. Shearing planar Noh problem

In this section we examine a very challenging test problem: propagating a shock over a shearing flow. In this test we set up an initially pressureless two-dimensional gamma law gas with an initial velocity field

$$\mathbf{v} = \begin{pmatrix} v_0 \cos 2\pi y \\ -1 \end{pmatrix}, \quad (32)$$

in a box with periodic boundary conditions in the x direction and a reflecting boundary condition along the $y = 0$ line. This system should evolve as the planar Noh problem in the y direction, but now the infinite strength shock will be propagating over a shearing flow. Ideally we expect the shock front and post-shock conditions to evolve independently of the shearing component of this flow field (since the shock is perpendicular to the shearing flow), but as we increase the level of the shear this becomes an increasingly challenging problem. Though the setup of this test case is somewhat contrived, it is arguably a more relevant multi-dimensional test case than the idealized tests we normally consider (such as the standard convergent Noh problems in Section 5.1). In almost all interesting multi-dimensional physical systems there are several components to the velocity field, not all of which are aligned with the direction of any shocks that are present. Propagating a shock over a shearing flow is a particularly relevant test case for many astrophysical problems: consider for instance a rotating gaseous disk forming in a gravitational well. Typically the supporting rotation curve of such a disk consists of a radially shearing flow, and we must follow the evolution of the angular momentum as additional material falls into the system and interacts with the disk via shocks. This type of scenario is typical of problems such as galaxy formation, galactic collisions, and star formation.

In Fig. 3 we compare the result of modeling the shearing Noh problem using the scalar viscosity to a model using the tensor form. The SPH nodes are initially seeded on a 40×100 lattice in the box between $(0,0)$ and $(0.4,1)$, and we set the magnitude of the shear velocity to $v_0 = 1$ in Eq. (32). We initialize smoothing scales such that there are two nodes per h , and employ the cubic B spline interpolation kernel $W = W_H = W_3$. Periodic boundary conditions are enforced at the $x = 0$ and $x = 0.4$ lines, while a reflecting boundary condition is placed on the $y = 0$ line. Fig. 3 plots the density and pressure of all nodes in these simulations vs. y position at time $t = 0.6$, along with the expected analytic profiles for the planar Noh problem. Clearly the tensor viscosity results are improved over the scalar version. As we discussed in Section 4 the shearing motion incorrectly contributes some additional deceleration/heating in the tensor viscosity, but the tensor form does not allow the shear to contribute any spurious acceleration in the y direction, unlike the scalar viscosity. This is an example of a problem where the shear correction f_B (Eq. (11)) for the standard viscosity simply cannot capture the correct behavior. Since there are both shearing and compressional components to this flow field, the viscosity must be active to capture the shock condition but inevitably the unphysical shear contribution to the viscosity is also allowed to occur. Even worse, with f_B reducing the scalar viscosity magnitude it is potentially reducing the effectiveness of the viscosity at capturing the shock condition and allowing some interpenetration of the nodes in the y direction. The combination of the scalar viscosity not being able to restrict acceleration to the y direction and the shear correction f_B reducing the viscosity's effectiveness against the shock explain the evident disorder that arises in the profiles of Figs. 3(a) and (b). The tensor viscosity is able to distinguish the compression in the y direction and operate appropriately in that direction, which in combination with the fact that the $\nabla(\nabla \cdot \mathbf{v})$ limiter allows the viscosity to operate at full strength in the y direction explains how the tensor simulation is able to do a relatively good job of capturing the Noh solution.

In Fig. 4 we plot results for simulations where we set the magnitude of the shear component to $v_0 = 5$. In this case the shearing flow is much stronger than the compressional component, making this a very challenging case. Clearly the scalar viscosity simulation is breaking down badly in this problem, with the shear induced viscosity component strongly heating the fluid and severely distorting the solution. The scalar viscosity results have also almost completely lost the symmetry of the problem, with the nodes becoming so mixed that the original rows in y are randomized. The tensor viscosity results are also degraded from what we see in the $v_0 = 1$ results in Fig. 3, though the planar Noh solution is still recognizable in contrast to the results using the scalar viscosity. The source of the degradation in this test can be traced to the difficulty measuring $\partial v^\alpha / \partial x^\beta$ and $\nabla(\nabla \cdot \mathbf{v})$ accurately to use in the viscosity formalism.

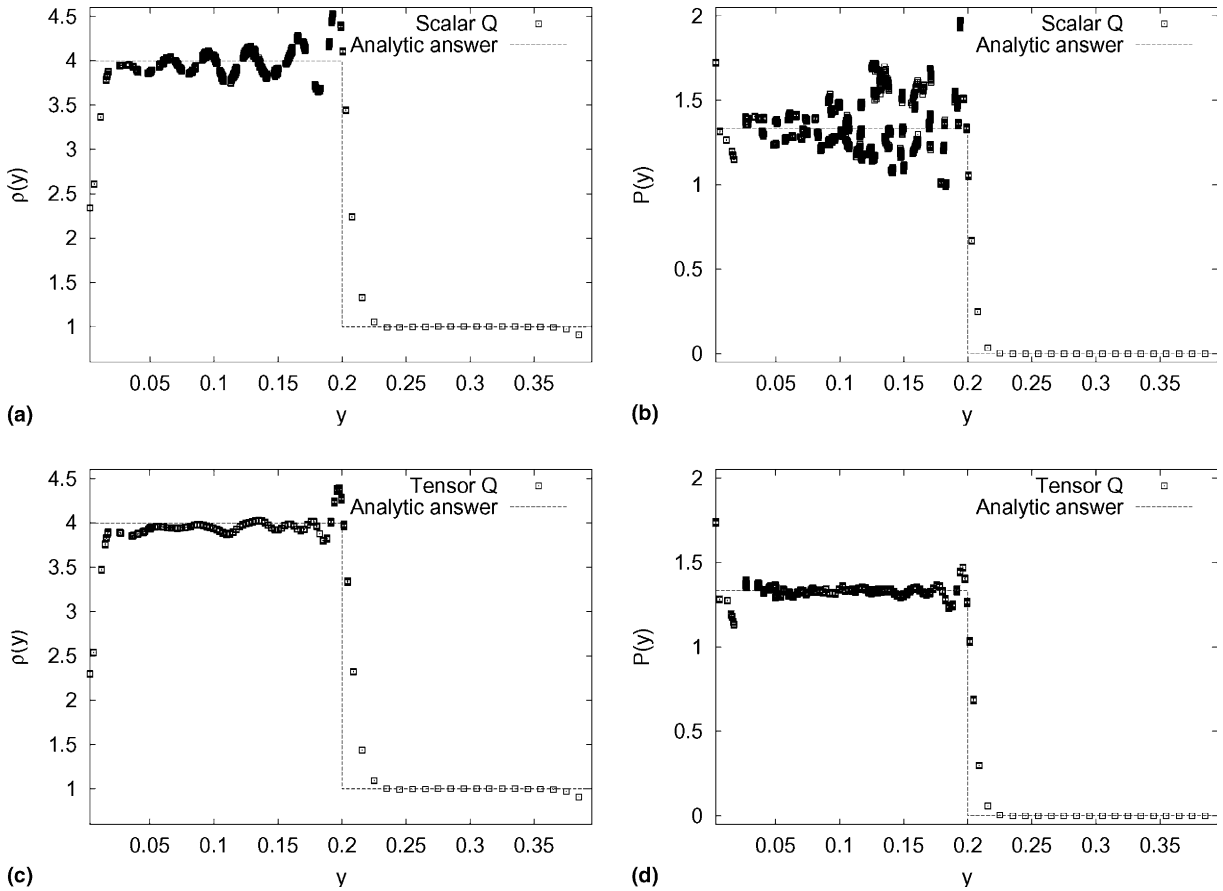


Fig. 3. Profiles of the mass density and pressure vs. y position for the shearing planar Noh shock problem (with $v_0 = 1$) at $t = 0.6$. (a) is the mass density and (b) is the pressure for the simulation using the standard scalar viscosity, while (c) and (d) are the corresponding results using the tensor viscosity.

5.3. Statically rotating disk with arbitrary pressure and rotational support

In this test case (and the next) we examine the angular momentum conservation properties of the tensor viscosity in the context of a gravitationally driven rotational shear. We consider these problems primarily as a test of the angular momentum conservation of simulations using the tensor viscosity, since (as discussed in Section 4) total angular momentum conservation is not guaranteed when the tensor viscosity is active. We also examine how well the analytic radial profiles are maintained, since any spurious activation of the artificial viscosity should cause deviation from the initial supporting profiles.

In this test case we establish a gaseous disk in rotational and pressure equilibrium with an externally imposed fixed gravitational potential. The potential is modeled as a softened point potential

$$\Phi(r) = -\frac{GMm}{(r^2 + r_s^2)^{1/2}}, \tag{33}$$

where G is the gravitational constant, M the mass of central gravitating mass, m the mass of the SPH node, and r_s the softening scale. We assume a polytropic equation of state $P = f_p K \rho^\gamma$, where K is the polytropic

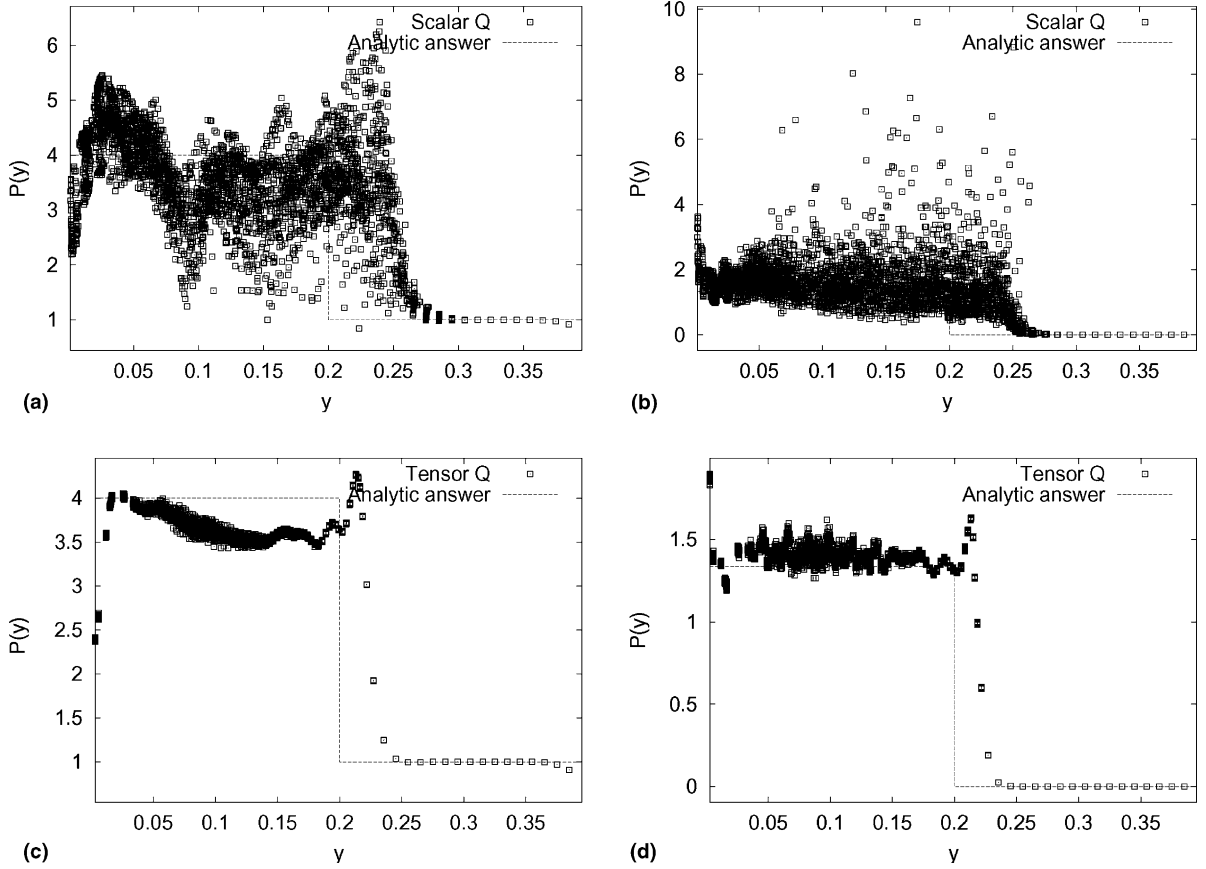


Fig. 4. Profiles of the mass density and pressure vs. y position for the shearing planar Noh shock problem at $t = 0.6$. In these simulations we have boosted the magnitude of the shear by choosing $v_0 = 5$, so that the shear component now has a higher peak velocity than the compressional component. (a) is the mass density and (b) is the pressure for the simulation using the standard scalar viscosity, while (c) and (d) are the corresponding results using the tensor viscosity.

constant and γ is related to the polytropic index n by $\gamma = (n + 1)/n$. f_p is an arbitrary parameter in the range $f_p \in [0, 1]$, which we use to select the amount of pressure support. The polytropic constant and radial mass density profile necessary to establish equilibrium are

$$K = \frac{GM}{3r_s \rho_0^{1/2}}, \quad (34)$$

$$\rho(r) = \left\{ \frac{GM(\gamma - 1)}{K\gamma(r^2 + r_s^2)^{1/2}} \right\}^{1/(\gamma-1)}, \quad (35)$$

where ρ_0 is the mass density at $r = 0$ (assumed to be the origin of the external potential in Eq. (33)). The supporting orbital velocity is

$$v^2 = (1 - f_p)^2 \frac{GM r^2}{(r^2 + r_s^2)^{3/2}}. \quad (36)$$

Since the system is initialized in equilibrium it should statically rotate forever with no radial evolution. However, the velocity field in Eq. (36) creates a radial shear, so any shear activation of the viscosity will unphysically cause angular momentum to be transported outward and force the system to deviate from this radial equilibrium.

Fig. 5 shows the radial profiles of modeling one case of this disk to time $t = 10$, where we have set $G = 1$, $M = 1$, $r_s = 1/2$, $f_p = 1/2$, $\rho_0 = 1$, and selected a polytropic index $n = 2$ so that $\gamma = 3/2$. These are two-dimensional simulations where we initialize the nodes on rings, each ring consisting of constant mass nodes spaced azimuthally to match the local radial spacing. The radial spacing of the rings is chosen to conform to the desired initial density profile. Each of these simulations incorporates roughly 5000 SPH nodes, which are seeded out to a maximum radius of $r_{\max} = 3$. The smoothing scales are initialized with two nodes per h , and we employ the fifth order interpolation kernel so that $W = W_5$ and $W_{II} = W_3$. The analytic solution for the disk is of course infinite, but since the mass density falls off as $\rho \propto r^{-2}$ the material in the disk rapidly becomes negligible with increasing radius. Small deviations from the analytic profiles are evident near the cutoff radius, but overall both simulations do an excellent job of maintaining the static solution for this system. In this figure we have plotted all the nodes in the simulations vs. radius, so we can directly see

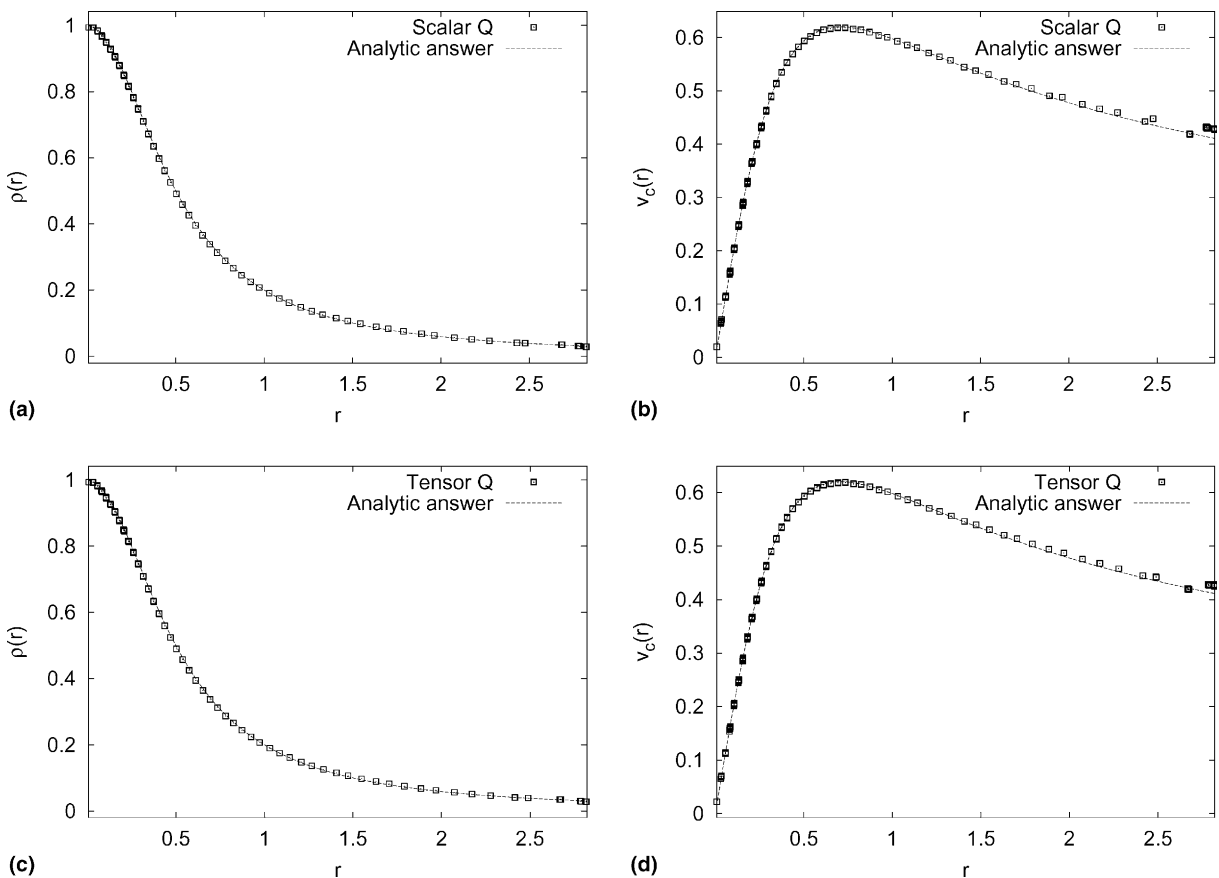


Fig. 5. Radial profiles of the mass density and rotational velocity in the statically rotating disk simulations at $t = 10$. (a) is the mass density and (b) the rotational velocity for the simulation using the scalar viscosity, while (c) and (d) are the corresponding results for the tensor viscosity.

how well symmetry is being maintained by inspecting these plots and verifying that the results of all the nodes at a particular radius share the same value for density and rotational velocity.

Fig. 6 plots the time history of the total angular momentum in the two simulations. Measuring the peak to peak relative error as

$$\frac{\Delta L_z}{L_z} = \frac{\max(L_z(t)) - \min(L_z(t))}{L_z(t=0)}, \quad (37)$$

we find that $\Delta L_z/L_z = 2 \times 10^{-7}$ in the scalar viscosity run, while for the tensor viscosity $\Delta L_z/L_z = 3 \times 10^{-4}$. This is of course something of a null test for the scalar viscosity since that formalism conserves angular momentum perfectly by construction: the small variation in the scalar viscosity result is due to the accumulation of round-off error and time integration errors in the orbital motion of the gas. Clearly the tensor viscosity simulation has conserved the total angular momentum in this problem very accurately, as we would expect since it has not caused deviation from the initial equilibrium profile. By time $t = 10$ the gas at radius $r = 1/2$ has completed roughly two rotations. We have pushed these simulations out much further in time, and eventually the errors in time integrating the orbital motion cause the results to become progressively noisier about the expected analytic profiles. However, we do not see any net angular momentum transport or resulting distortion of the disk profiles.

5.4. Gravitationally driven collapse of a rotating fluid

For our final set of test cases we consider an initially pressureless gamma-law gas cloud born in solid body rotation in an external potential given by Eq. (33), but now rotating too slowly to be supported. This system rapidly collapses in the gravitational potential, shocks, and forms a hot, dense, rotating disk. In this case we do not have an analytic solution for what the final collapsed disk should look like, so instead we run many versions of this problem with increasing resolution to see if the results converge with resolution. This problem poses a stringent test for the angular momentum evolution, since the angular momentum will be transported over large ranges in radius and the artificial viscosity will be strongly activated as the material collapses to form the dense central disk of material. Therefore, it is interesting to look at not only whether the total angular momentum is conserved, but also how it is transported amongst the mass of the system as it evolves.

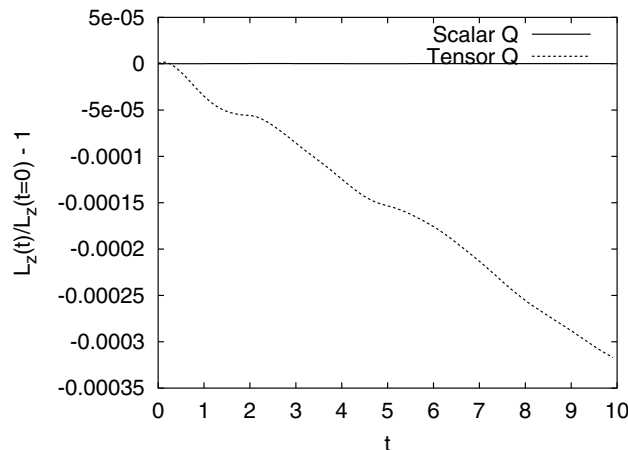


Fig. 6. Time evolution of the error in the z component of the total angular momentum in the statically rotating disk simulations.

We begin with a series of two-dimensional simulations. The imposed potential is identical to that used in the tests presented in Section 5.3 ($M = G = 1$, $r_s = 1/2$). The gas cloud is initialized as a $\gamma = 5/3$ gas at constant density $\rho_0 = 1$ out to a radius of $r_{\max} = 10$, with an initial solid body rotation such that the rotation at the outer edge is 20% of the critical support value. The gas disk is initialized as rings of SPH nodes with constant radial spacing and azimuthal spacing equal to the radial. We perform simulations of this system at four different resolutions (250, 1000, 4000, and 16,000 SPH nodes) with both the scalar and tensor viscosities, for a total of eight simulations. Note that each simulation represents a doubling of the linear resolution and a quadrupling of the mass resolution compared to the prior. In each case we initialize the smoothing scale such that there are two nodes per h , and use the cubic B spline $W = W_H = W_3$. The simulations are run to a final time $t = 100$.

Table 1 shows the total peak to peak error in the z component of the angular momentum in each of the eight simulations, and Fig. 7 shows the time evolution of the error $L_z(t)/L_z(t=0) - 1$. This figure is a bit busy since there are so many simulations plotted here, but the point is that none of the simulations violate total angular momentum conservation by more than a few percent: the lowest resolution tensor viscosity simulations show a total peak to peak error of 5% in the angular momentum, while the highest resolution simulation shows an error closer to 1%. These results seem acceptable, with the error in the total angular momentum similar to what is typically seen in the total energy.

Fig. 8 shows the radial profiles for the mass density and circular velocity of the collapsed disk material at the final time, $t = 100$ in each simulation. We identify the material in the disks as any SPH node with mass density $\rho_i \geq 20$. In order to generate these profiles, the nodes in the disk are sorted in radius from the center of the potential, binned in bins of constant number of nodes, and then the average values and standard deviations are computed for each bin. Several interesting trends are apparent in these plots. The mass density profiles in the scalar viscosity runs (Fig. 8(a)) tend to decrease with increasing resolution. This implies that the angular momentum is increasingly transported away from the disk with decreasing resolution, thereby allowing mass to settle into a lower orbit. The mass density profile in the 4000 and 16,000 node scalar viscosity runs seem to have converged. However, the circular velocity profiles in Fig. 8(b) do not show such a clear trend. The lowest resolution simulation (250 nodes) clearly loses nearly all of its angular momentum, and therefore shows relatively little rotational motion. The 1000 and 4000 node results demonstrate a trend toward larger rotational motion with increasing resolution, but then the 16,000 node simulation contradicts this, coming in closer to the 1000 than 4000 node result. All of the rotation curves in the scalar viscosity runs show a relatively large amount of scatter, evidenced by the large standard deviations shown in Fig. 8(b). The results with the tensor viscosity show much better agreement as we increase resolution, and much less scatter in the velocity bins. In fact, with the exception of the rotation curve of the lowest resolution tensor viscosity run, the tensor viscosity profiles are almost insensitive to resolution. However, the rotational velocities in the runs using the tensor viscosity are much lower than those seen in the simulations using the scalar viscosity, indicating the disk material winds up with less angular momentum.

Fig. 9 shows the time evolution of the mass and angular momentum transport in the disk. In order to try and gauge the amount of angular momentum transport occurring during these simulations we use the following prescription: at each time slice we determine which nodes meet the criterion $\rho \geq 20$ for membership in the disk; we then sum up the amount of angular momentum in this set of nodes at the time in question

Table 1

Peak to peak error in the z component of the total angular momentum $\Delta L_z/L_z$ (Eq. (37)) in the two-dimensional collapsing disk simulations

	250 Nodes	1000 Nodes	4000 Nodes	16,000 Nodes
Scalar viscosity	5.6×10^{-6}	5.3×10^{-3}	2.1×10^{-3}	1.3×10^{-4}
Tensor viscosity	5.3×10^{-2}	5.3×10^{-2}	2.2×10^{-2}	1.4×10^{-2}

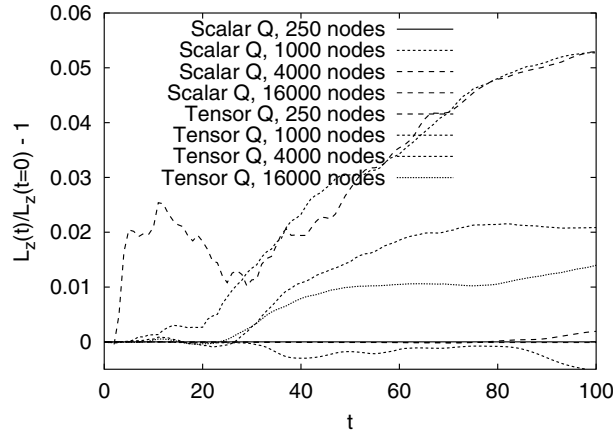


Fig. 7. Time evolution of the error in the z component of the total angular momentum in the two-dimensional collapsing disk simulations. The instantaneous angular momentum error is defined as $L_z(t)/L_z(t=0) - 1$.

and the amount of angular momentum those nodes were born with at $t = 0$; the ratio of these total angular momenta is then what is plotted in Figs. 9(b) and (d). Examining the growth of the mass in the disk it is clear that the disk has finished assembly by time $t = 40$. The low resolution (250 and 1000 nodes) scalar viscosity runs predict too much mass accumulating in the disk, consistent with the idea that these low resolution simulations transport too much angular momentum away from collapsing material and thereby allow too much mass to accrete onto the disk. The 4000 and 16,000 node scalar viscosity results seem to converge on the mass in the disk. The final disk mass in the tensor viscosity results does not show as much dependence on resolution, and is consistent with (though slightly higher than) the high resolution scalar viscosity result, indicating that roughly 8% of the material in the simulation winds up in the disk.

The evolution of the angular momentum transport in the disk is strikingly different in the tensor viscosity simulations as compared with the scalar viscosity results. As we increase the resolution the tensor viscosity runs seem to be resolving a higher and higher peak of transport of angular momentum into the high density disk material at early times, which then rapidly falls off so that the mass in the disk varies between having twice and one times the angular momentum it was born with. The simulations using the scalar viscosity show a very different trend, with angular momentum being transported into the high density material throughout the formation of the disk, and then slowly being transported away after the disk is assembled. The scalar viscosity runs also never converge on the evolution of this quantity, and similarly to the trend noted in the rotation curves in Fig. 8(b) the pattern is not monotonic with increasing resolution – the 16,000 node result falls between the 1000 and 4000 node simulations. By contrast the 1000, 4000, and 16,000 node tensor viscosity simulations agree fairly well on the evolution of the angular momentum in the disk.

Granted that there is a significant amount of angular momentum transport evident in Figs. 9(b) and (d), the question of how well these simulation maintain their axisymmetry naturally arises. If the mass distribution were to remain strictly axisymmetric, then the viscosity represents the only mechanism whereby angular momentum can be transported. However, if the rotating disks develop significant non-axisymmetries (such as a bar mode) then we would expect torquing on this structure to provide a strong mechanism for angular momentum transport. In order to investigate the degree of axisymmetry in the simulations, we take a second moment of the mass distribution according to

$$J^{z\beta} = \frac{\sum_i m_i x_i^z x_i^\beta}{\sum_i m_i}, \quad (38)$$

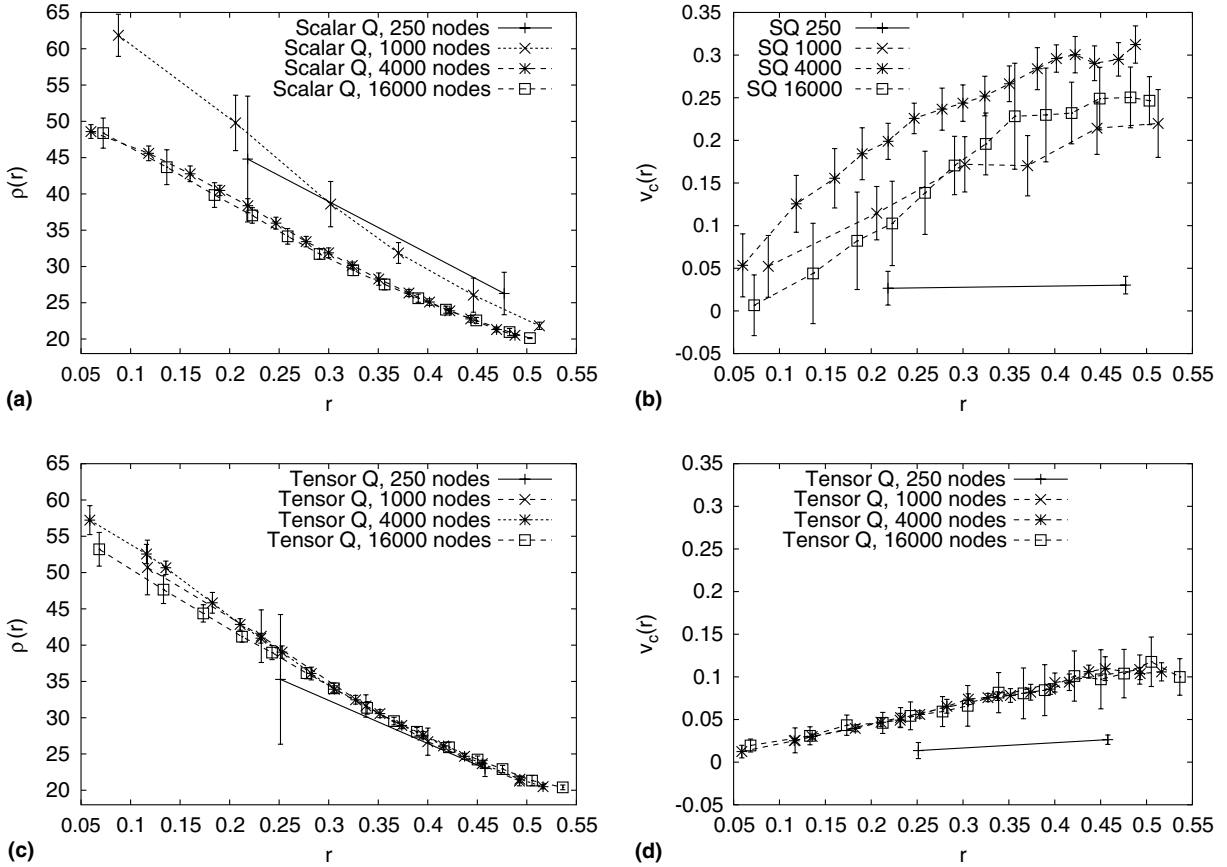


Fig. 8. Radial profiles of the mass density and rotational velocity in the two-dimensional collapsing disk simulations at $t = 100$, where the disk is selected as any nodes with mass density $\rho \geq 20$. Nodes are binned radially and the profiles averaged – the errorbars show the 1σ distribution of values in each radial bin. (a) is the mass density and (b) the rotational velocity for the simulations using the scalar viscosity, while (c) and (d) are the corresponding results using the tensor viscosity.

where the position x^α is with reference to the center of the imposed potential. In the two-dimensional case we can associate the eigenvalues of $J^{\alpha\beta}$ (defined here as j_1 and j_2) as the squares of the axes of an ellipse (denoted as h_1 and h_2) fitting the mass distribution: $h_1^2 = j_1$ and $h_2^2 = j_2$ (see the discussion in Appendix B of [15]). Therefore, we can define a shape factor F_s as the ratio

$$F_s = \frac{\min(h_1, h_2)}{\max(h_1, h_2)}, \tag{39}$$

which provides a simple metric of how round the mass distribution is. If the system maintains perfect axisymmetry then $F_s = 1$, while for varying degrees of elongation $F_s < 1$.

Fig. 10 plots $F_s(t)$ for each of the two-dimensional collapsing disk simulations, both for the total mass in the system as well as just the mass in the disks. Clearly the disks themselves maintain axisymmetry fairly well (with $F_s(t) \sim 0.95$). The total mass shows a bit more elongation (with the extreme example being the 250 node scalar viscosity case), but this seems to be primarily due to the granularity of mass being ejected from the system and traveling to large radii. The long moment arm of this material at large radius enhances its contribution to $J^{\alpha\beta}$, but the poor sampling of the angles represented by the small number of

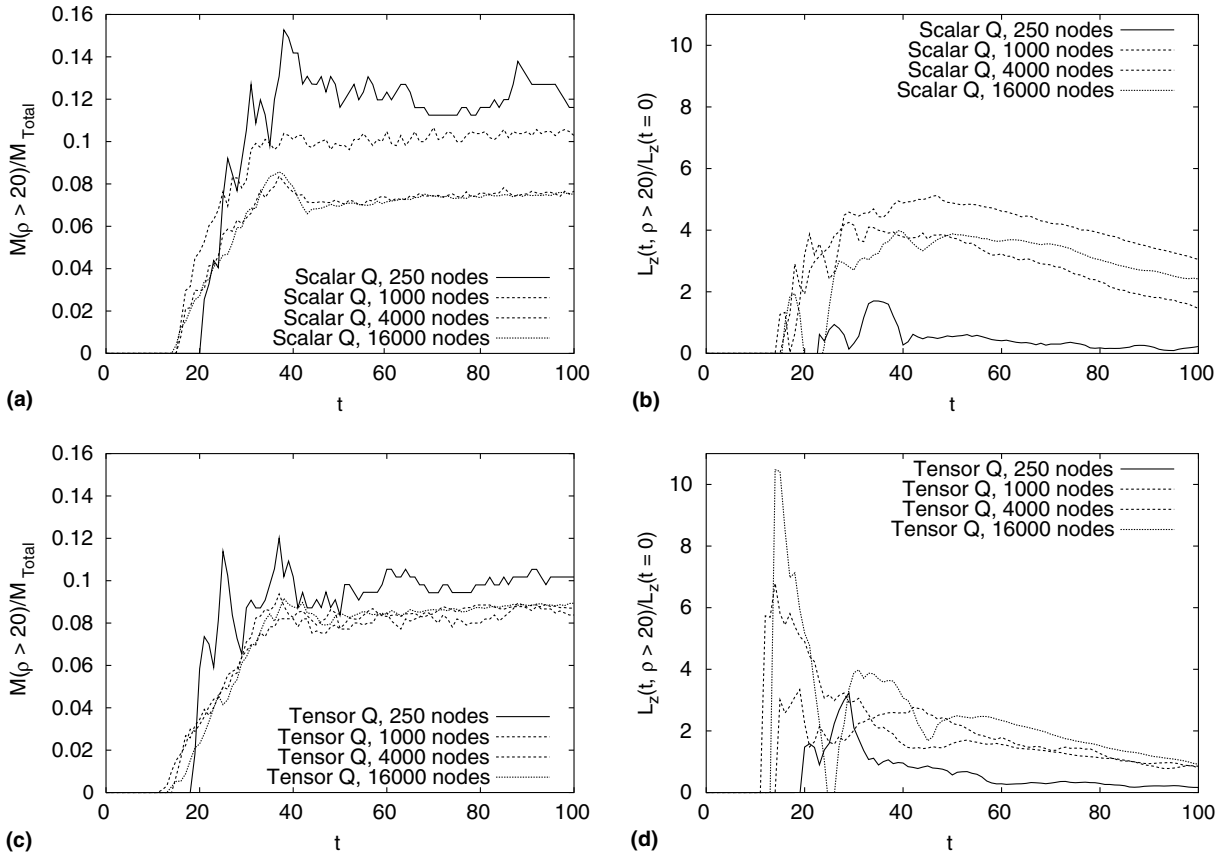


Fig. 9. Time evolution of the mass and angular momentum ratio in the two-dimensional collapsing disk simulations, where the disk is selected as any nodes with mass density $\rho \geq 20$. (a) is the mass and (b) the $L_z(t)/L_z(t=0)$ ratio for the simulations using the scalar viscosity, while (c) and (d) are the corresponding results using the tensor viscosity.

nodes that escape to these large distances leads to an artificial apparent elongation of the mass distribution. Even with this effect all but the lowest resolution scalar viscosity simulation maintain $F_s(t) > 0.9$ for all mass throughout the runs, so we do not see evidence that these simulations deviate significantly from axisymmetry.

Taken as a whole, these comparisons indicate that the tensor viscosity is more active in the formation of the disk material than the scalar viscosity. It is likely that that shear correction f_B effectively suppresses the scalar viscosity throughout the evolution of the disk and may well over-suppress the viscosity, allowing some degree of unphysical interpenetration of the SPH nodes. The scatter in the rotational velocity in these simulations could be a symptom of this. The tensor viscosity should be allowed to exert its full deceleration in the radial direction, though it also is likely experiencing some degree of shear induced enhancement such as we see in the strong shear case of the shearing Noh problem in Fig. 4. The tensor viscosity being more active radially explains why the tensor results predict a more rapid transport of angular momentum into the disk material at early times, why the rotation curves show much less scatter in the tensor simulations, and also why the final rotation of the disk is lower in the tensor than the scalar viscosity simulations. Unfortunately we do not know what the correct answer is for this system. The tensor viscosity results seem to converge to a consistent answer for all the simulations with more than 250 nodes, while the angular

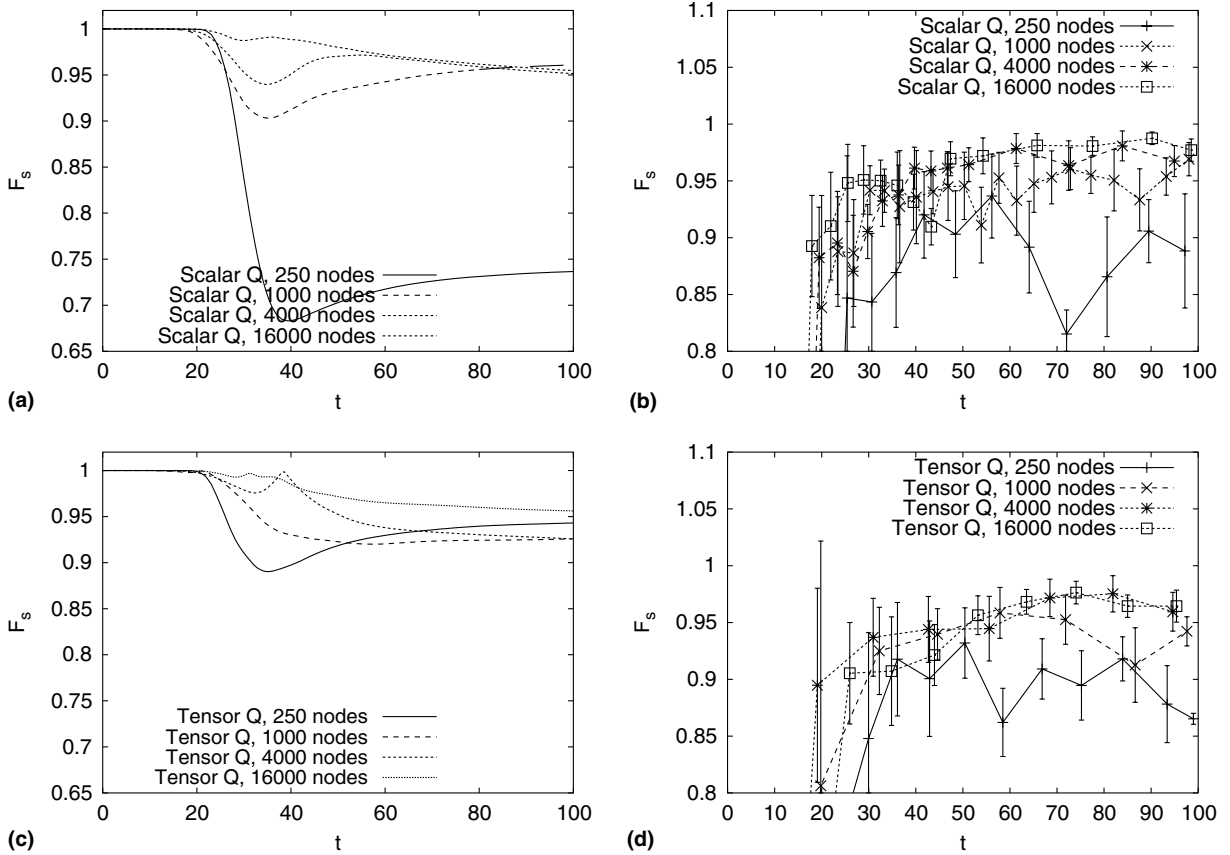


Fig. 10. Time evolution of the ratio of the minor to major axes for the fitted ellipse to the mass in the two-dimensional collapsing disk simulations (see Eq. (39)). (a) and (b) show the results for the scalar viscosity simulations, where the ellipse is either fitted to all of the mass (a) or restricted to just the mass in the disk (b). (c) and (d) are the corresponding results for the simulations with the tensor viscosity. Note that for the disk measurements in (b) and (d) we bin the curves in time and plot the binned results with errorbars showing the standard deviation.

momentum and velocity information in the scalar viscosity simulations have not converged even in the 16,000 node simulation. This could be an indication that the tensor viscosity results are correct, or that the tensor viscosity simulations are incorrectly (but consistently) transporting too much angular momentum away from the collapsed disk material.

In our final set of tests we repeat the collapsing cloud angular momentum test in three dimensions. In this case we initialize an initially spherical cloud in the same potential used in the previous two-dimensional tests, again with initial constant unit density and solid body rotation about the z -axis equal to 20% of the support velocity at the outer radius $r_{\max} = 10$. We exploit the symmetry of this system by establishing a reflecting boundary condition on the $z = 0$ plane and only creating a hemisphere of gas in the $z > 0$ volume. We run two different resolutions of this problem using the scalar and tensor viscosity formalisms, for a total of four simulations. The low-resolution simulations initialize the SPH nodes on a $20 \times 20 \times 10$ lattice above the $z = 0$ plane, and then we throw away any nodes with $r > r_{\max} = 10$ resulting in a total of 2112 nodes. The high-resolution simulations use the same procedure with a $40 \times 40 \times 20$ lattice, resulting in a total of 16,777 nodes. Note that the higher-resolution simulations double the linear resolution and increase the mass

resolution by a factor of eight over the low-resolution models. As in the previous two-dimensional tests we initialize smoothing scales with two nodes per h , and employ the cubic B spline interpolation kernel $W = W_H = W_3$. This system differs from the two-dimensional tests in that now we have gas falling in from above the plane of rotation, which will interact, shock, and shear with the forming disk at all angles from the rotational axis.

Table 2 and Fig. 11 test the error in the total angular momentum conservation of L_z , just as we did with the two-dimensional tests. Because we employ a reflecting boundary condition on the $z = 0$ plane there is no point in testing the x and y components of the angular momentum, since these will always be identically zero. The tensor viscosity simulations conserve the angular momentum in this test to $\sim 3\%$ (low-res) and $< 1\%$ (high-res), which is consistent with what we see in the two-dimensional models.

Fig. 12 plots the radial profiles of the mass density and fluid circular velocity at the final time $t = 100$. In these plots we define the radius as the distance from the z -axis, so each bin represents a cylindrical shell slice through the system. We now find that all the simulations agree quite well on the mass density profile, even demonstrating similar amounts of scatter in each radial bin. Likely some of this scatter is due to the fact that in three dimensions we are averaging over nodes at different z heights above the plane of the disk, and therefore each radial bin is sampling qualitatively different material (unlike the cleaner test in two dimensions, when all of the material is confined to the plane of the disk.) The circular velocity profiles are not converged for either viscosity, though we again note the scalar viscosity runs at each resolution find a somewhat higher rotational velocity with more scatter in each radial bin as compared with the results using the tensor viscosity. Comparing this with the results noted in the two-dimensional simulations in Fig. 8 suggests that the trends in the three-dimensional rotational velocity curves are similar to what we find in two-dimensional, and we simply have not used enough SPH nodes to achieve convergence in the tensor viscosity results.

Fig. 13 plots the time evolution of the mass of the disk and the angular momentum transport in three dimensions, calculated just as we did for the two-dimensional simulations in Fig. 9. The scalar viscosity simulations demonstrate a fairly large overshoot in the disk mass at $t \sim 35$ (compared with the tensor viscosity simulations), after which time the disk mass converges to roughly the same value in all simulations. It is possible this initial overshoot in the mass is due to over-suppression of the scalar viscosity by f_B , allowing nodes accreting onto the disk to penetrate too deeply into the body of the disk before they are halted and pushed back out. Consistent with the rotation curves in Fig. 12(b), the time evolution of the angular momentum transport does not converge for any of the three-dimensional simulations. There is a clear trend for the scalar viscosity to transport more angular momentum into the disk as it is assembled, and both formalisms show that the amount of angular momentum in the disk increases with increasing resolution. It is encouraging that once the disk is formed the tensor formalism loses angular momentum from the disk at the same rate as the shear suppressed scalar viscosity, despite the fact that we do not explicitly suppress the shear contribution to the tensor viscosity.

Fig. 14 plots the evolution of the shape factor $F_s(t)$ test for axisymmetry, precisely as we did for the two-dimensional case in Fig. 10. Since we expect axisymmetry about the z -axis we only use the (x,y) components of the positions in Eq. (39), thereby ignoring the compression in the z direction for these plots. The disk

Table 2

Peak to peak error in the z component of the total angular momentum $\Delta L_z/L_z$ (Eq. (37)) in the three-dimensional collapsing disk simulations

	2113 Nodes	16,777 Nodes
Scalar viscosity	1.6×10^{-4}	8.4×10^{-5}
Tensor viscosity	3.1×10^{-2}	8.8×10^{-3}

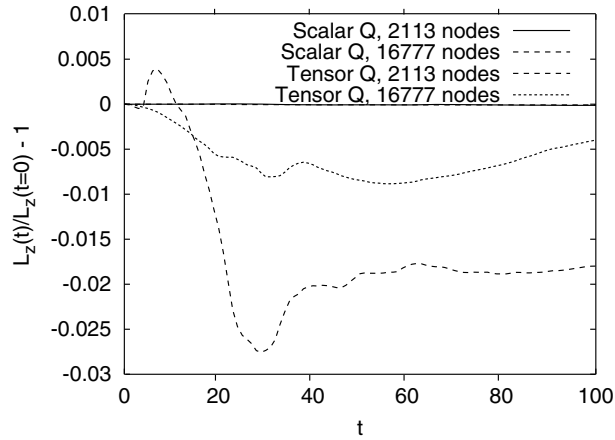


Fig. 11. Time evolution of the error in the z component of the total angular momentum in the three-dimensional collapsing disk simulations.

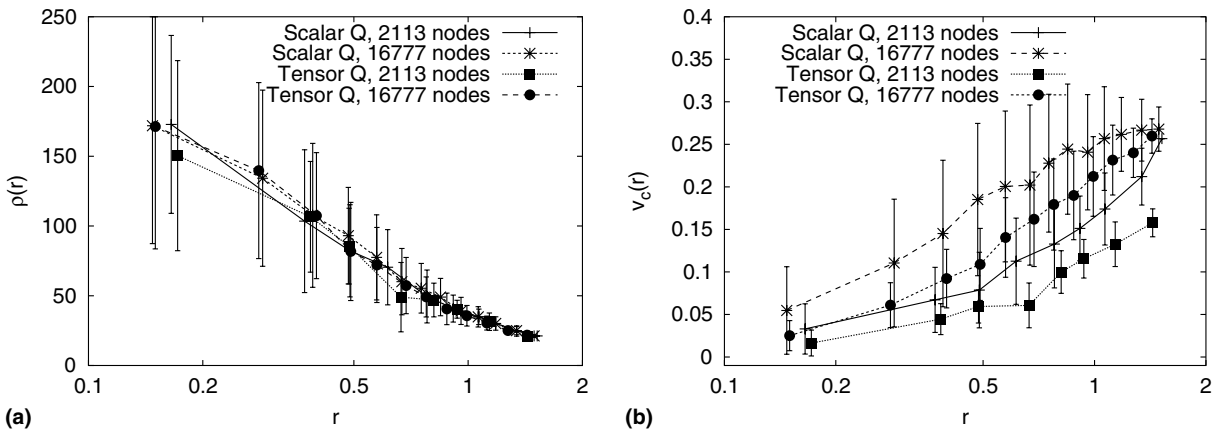


Fig. 12. Radial profiles of the mass density and rotational velocity in the three-dimensional collapsing disk simulations at $t = 100$, where the disk is selected as any nodes with mass density $\rho \geq 20$. Nodes are binned radially and the profiles averaged – the errorbars show the 1σ distribution of values in each radial bin. (a) shows the mass density and (b) the rotational velocity.

mass evidently maintains axisymmetry very well with $F_s(t) \gtrsim 0.95$ throughout the runs. The total mass shows more elongation in the two low-res runs, though both high resolution simulations maintain axisymmetry for the total mass very accurately. Therefore, just as in the two-dimensional collapsing disk tests we again find that these simulations maintain axisymmetry well throughout the runs, so the angular momentum transport we see in Fig. 12(b) does not seem to be due to non-axisymmetry in the disks.

So what do these collapsing gas cloud simulations tell us about the tensor viscosity and angular momentum? All of the tensor viscosity tests here demonstrate good conservation of the total angular momentum, which is the necessary lowest bar we must clear in order for the results to be believable. This argues that the fact that the tensor viscosity is not guaranteed to conserve angular momentum should not rule out the tensor viscosity as a useful tool. Additionally, the fact that once the disk has formed the tensor viscosity does not show a greater rate of loss in angular momentum from the disk than the shear corrected scalar viscosity

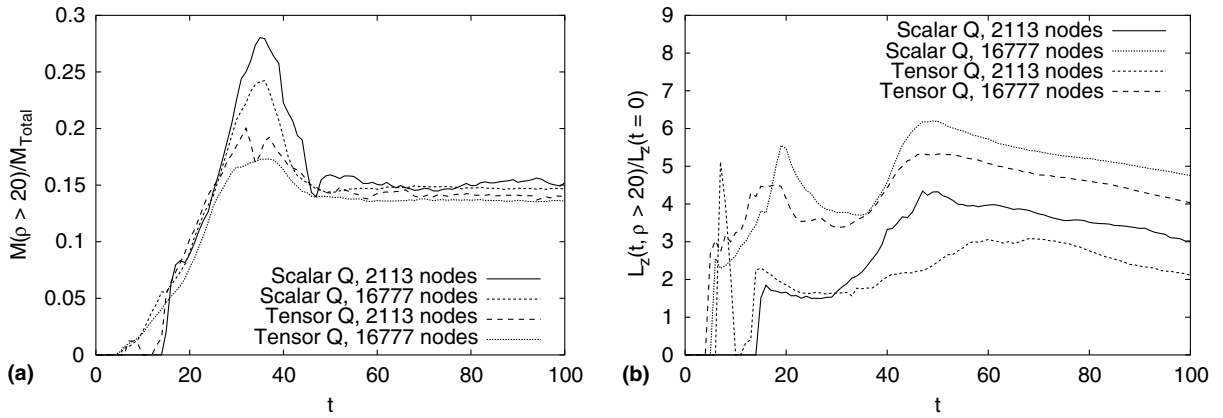


Fig. 13. Time evolution of the mass and angular momentum ratio in the three-dimensional collapsing disk simulations, where the disk is selected as any nodes with mass density $\rho \geq 20$. (a) shows the evolution of the disk mass and (b) the evolution of $L_z(t)/L_z(t=0)$.

is also a very good sign. Previous investigators such as Navarro and Steinmetz [12] have verified that the shear corrected scalar viscosity formalism does an adequate job of maintaining the angular momentum of a rotating disk once it is formed, whereas without the use of the shear correction the scalar viscosity fails such tests badly. The fact that the tensor viscosity is able to match the angular momentum transport loss rate of the shear corrected scalar viscosity for the rotating disk after it is formed without resorting to an explicit shear correction is therefore quite impressive. However, there are yet unanswered questions about how both of these viscosities are affecting the local transport of angular momentum during the assembly of a rotating structure such as this disk problem. The angular momentum and rotational velocity of the simulations employing the scalar viscosity never converge, whereas while the simulations using the tensor viscosity do converge (at least in the two-dimensional examples shown here), they characteristically transport less angular momentum into the disk than the corresponding scalar viscosity runs and therefore have a lower rotational




Fig. 14. Time evolution of the shape parameter (Eq. (39)) of the mass distribution in the three-dimensional collapsing disk simulations. (a) shows the results using all the mass in the simulations, while (b) shows the results in the various simulations for the mass in the disk only.

velocity. Despite the convergence of the tensor viscosity runs, ultimately it is not yet clear which methodology is doing a better job of modeling the angular momentum transport in this system.

5.5. Computational expense of the tensor viscosity

An obvious concern with the implementation of the tensor viscosity is the computational expense it will incur relative to the simpler scalar viscosity. The new tensor formalism laid out in Section 4 clearly involves much more work than the scalar prescription described in Section 3, and since the tensor viscosity is a pair-wise operation much of this work is required in the innermost loops of the derivative evaluation in an SPH scheme. Additionally, in a straightforward implementation we require an extra loop over the SPH nodes to evaluate the background gradient of the velocity $(\partial v^\alpha / \partial x^\beta)_i$ before the standard fluid derivative evaluations can begin. This implies an extra pass through the whole neighbor node selection process, which depending on the implementation can incur a large overhead. These sorts of details make comparing the expense of a given technique difficult for different implementations. The final bottom line cost of a new algorithm such as this depends strongly on how the time spent evaluating the viscosity compares to the other costs in the code: details of how SPH node neighbor selections are done, the exact choice of time integration scheme, etc. will all have a strong influence on the total run-time. Therefore, we present two different measures of the relative expense of the tensor viscosity in our current implementation (as compared with the scalar viscosity): the ratio of the time spent evaluating just the viscous component of the fluid derivatives, and the total run-time ratio. Table 3 summarizes these comparisons for the two-dimensional and three-dimensional collapsing disk problems shown in Section 5.4.

In the two-dimensional case we see that calculating the viscosity forces alone is about 1.7 times more expensive with the tensor viscosity, while the total run-time increases by a factor of 2.5. The reason the total run-time increases more than the derivative penalty alone is the extra pass over the nodes evaluating $(\partial v^\alpha / \partial x^\beta)_i$. In the three-dimensional case we see that the individual expense of calculating the viscous acceleration increases by almost a factor of 8, while the total run-time increases by a factor of 4. The percentage of time spent in various parts of the code changes going from two to three dimensions, which is why in three dimensions the overall run-time does not increase by as much as the viscosity evaluation itself would suggest. It is also likely that much of the reason for the increase in the expense of evaluating the tensor viscosity term in three dimensions vs. two dimensions is due to the fact that in two dimensions we analytically solve for the rotational transformation R_S which diagonalizes the symmetric component of the velocity gradient (see Eq. (21)), whereas in three dimensions we are using the iterative Jacobi scheme to diagonalize this term (as described in Appendix C.) Since this diagonalization is occurring in the innermost derivative loop for every interacting pair of nodes, it would probably pay to similarly solve this problem analytically in the three-dimensional case. In a similar vein it is also worth emphasizing that the performance numbers here represent a worst case scenario, since the tensor viscosity implementation used here was written with flexibility in mind in order to facilitate experimentation with different forms for the viscosity. There has not yet been an effort to optimize this implementation (with improvements like avoiding the pair-wise Jacobi solve in three dimensions), so presumably this performance can be improved. Ultimately of course the tensor

Table 3
The tensor viscosity/scalar viscosity run-time ratios for the collapsing disk simulations in Section 5.4

	Total	Viscosity only
Two-dimensional	2.5	1.7
Three-dimensional	3.9	7.8

The first column shows the ratio of the total time taken by the tensor viscosity simulations over the scalar, while the second shows the ratio of time spent just evaluating the viscous forces in the fluid derivative calculations.

viscosity will always be more expensive than the simpler scalar form, and it is up to the numerical practitioner to decide if the gains in the solution obtained are worth the expense.

6. Summary and discussion

We have developed a tensor artificial viscosity for use with SPH, which reduces to the well-known Monaghan and Gingold [10] form in one dimension. We have also derived a directional limiter designed to limit the acceleration due to the artificial viscosity to the direction most directly aligned with any local shocks. This limiter is well suited for use with a tensor viscosity because the tensor formalism allows for directionality in the acceleration, which can then be effectively masked or allowed by the limiter. Our motivation in this work has been to better address multi-dimensional problems such as the shearing shock problem in Section 5.2. We believe these sorts of truly multi-dimensional problems are more representative of the physical systems which we use tools such as SPH to study than the traditional one-dimensional test cases we typically tune hydrodynamic algorithms against. The shearing shock problem is manifestly two-dimensional, as opposed to the more typically utilized one-dimensional test problems such as the standard Noh shock problem. Although the cylindrically and spherically convergent Noh problems are conceptually two- and three-dimensional tests, in fact these are really one-dimensional tests in that they can be entirely specified as one-dimensional problems in the appropriate coordinate system. Unfortunately most of the tests we use to validate our numerical hydrodynamic methods are intrinsically one-dimensional in this sense, and therefore do not necessarily prove that a given technique is truly appropriate for multi-dimensional physics. Such one-dimensional tests are a necessary first bar any hydrodynamic technique must pass, but we need more multi-dimensional tests if we are truly to trust the application of these tools for the study of complex multi-dimensional fluid flows. In this sense we believe the shearing shock problem presented here is as important as the tensor viscosity itself.

The question of angular momentum conservation is one of the most important issues with a tensor artificial viscosity such as this. Standard SPH utilizing one of the typical scalar viscosity formalisms is somewhat unique among numerical hydrodynamic methods in that it simultaneously conserves both linear and angular momentum, a property we are loathe to surrender. However, the advantages of the tensor viscosity formalism in being able to discriminate different velocity signals in complex flows (and align the local acceleration appropriately) make it often more appropriate for multi-dimensional problems, so long as the angular momentum conservation is not grievously violated. We have investigated the angular momentum conservation of the tensor viscosity extensively in this paper, and found that even in the most challenging rotational problems total angular momentum is always conserved to a few percent. This places the conservation of total angular momentum on par with typical conservation of total energy in such simulations, so we believe the tensor viscosity passes the angular momentum test.

However, clearly a great deal of work clearly remains to be done on angular momentum transport issues in SPH simulations. Although we demonstrate here that the tensor viscosity conserves total angular momentum acceptably, there are still many questions about the transport of angular momentum in both the scalar and tensor viscosity formalisms. It is worrisome that in the two-dimensional scalar viscosity simulations of the collapsing disk in Section 5.4 neither the rotation curves nor the angular momentum transport for the disk converge even with very highly resolved simulations. Although the tensor viscosity simulations do converge, they demonstrate a much lower rotation rate than the scalar viscosity simulations, the angular momentum history looks quite different, and ultimately we do not know that either case is correct. This sort of problem is very relevant for physical systems investigators are trying to model today, such as the formation of disk galaxies in a cosmological context. Cosmological simulations have long been successful at creating simulated galaxy populations with about the expected mass distribution, but almost invariably the final population of simulated disk galaxies lose far too much of their angular momentum

during the formation process so that the final disk is too centrally concentrated and rotating too slowly. This error can be quite severe, with angular momentum errors by as much as one or two orders of magnitude [4,11–13]. A variety of explanations for this problem have been proposed supposing that the problem is related to missing physics in the models, but the lack of convergence we see in the idealized collapsing rotational problems here suggests that at least some of the error could be numerical in origin. This problem is one example of a situation wherein correctly determining the local transport of angular momentum is more relevant than achieving perfect global conservation.

We would also like to point out that the prescription for defining the tensor viscosity in Section 4 consists of a number of distinct pieces that are not necessarily unique, and therefore there is much room to experiment. The most obvious step that can be changed is the method for determining $\overline{\partial v^\alpha / \partial x^\beta}$ (as well as $\overline{\nabla(\nabla \cdot \mathbf{v})}$ for the limiter), which is part of why we separate out those sections into the appendices. A simple choice that can be made is to determine these quantities using the SPH formalism: $\overline{\partial v^\alpha / \partial x^\beta} = \langle \partial v^\alpha / \partial x^\beta \rangle = -\rho_i^{-1} \sum_j m_j (v_i^\alpha - v_j^\alpha) \partial W_{ij} / \partial x^\beta$. We have experimented with this choice and found the tensor viscosity results improve for the convergent Noh problems (in particular the spherical result is less noisy in the post-shock regime) as well as the shearing Noh problem (the post shock solution in the $v_0 = 5$ case in Fig. 4 shows less of a dip in the mass density curve), which is probably all related to having less noise in the direction of the gradient term. However, we also found that using the SPH definition for $\overline{\partial v^\alpha / \partial x^\beta}$ results in catastrophic angular momentum conservation errors in the collapsing disk problem in Section 5.4, with errors of order $\Delta L_z / L_z \sim 50\%$. We believe this error arises because the SPH estimate of the gradient fails in a systematic way for systems with very large gradients in the mass density/SPH node distribution. This is related to the boundary problem in SPH, where the SPH definition for interpolated field values and gradients characteristically falls off as one approaches a material boundary where the SPH nodes stop. The simple pair-wise estimate for $\overline{\partial v^\alpha / \partial x^\beta}$ in Appendix A is specifically designed to be robust in such boundary situations, and results in the much improved angular momentum properties we present in this paper. Nonetheless, there is still significant room to experiment with this and other choices in the recipe defining the tensor viscosity.

In closing, all the work presented here has been performed in the context of an open source meshless hydrodynamics project, and instructions for obtaining the set of code used in this paper as well as the test setups for all the problems presented here are available at <http://spherical.sf.net/TensorQ.html>. We invite anyone interested to download this software and try out reproducing either our results or try constructing additional test problems of your own.

Acknowledgment

This work was performed under the auspices of the US Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48.

Appendix A. Estimating $\overline{\partial v^\alpha / \partial x^\beta}$

It is critical that the method used to estimate $\overline{\partial v^\alpha / \partial x^\beta}$ for use with the tensor viscosity be robust. After experimenting with several options, we settled on a low order pair-wise summed definition given by

$$\frac{\overline{\left(\frac{\partial v^\alpha}{\partial x^\beta}\right)}_i}{\sum_j \frac{m_j}{\rho_j} W_{ij} |(R_\sigma^{\gamma\alpha})_{ij} (1^\gamma \delta^{\psi 1}) (R_\sigma^{\psi\beta})_{ij}|} = \frac{\sum_j \frac{m_j}{\rho_j} W_{ij} (R_\sigma^{\gamma\alpha})_{ij} \left(\frac{(R_\sigma^{\gamma\alpha})_{ij} v_j^\alpha}{|r_{ij}|} \delta^{\psi 1} \right) (R_\sigma^{\psi\beta})_{ij}}{\sum_j \frac{m_j}{\rho_j} W_{ij} |(R_\sigma^{\gamma\alpha})_{ij} (1^\gamma \delta^{\psi 1}) (R_\sigma^{\psi\beta})_{ij}|}, \tag{A.1}$$

where $\mathbf{1}^\alpha$ is a vector with unity for each element, $\delta^{\alpha\beta}$ is the Kronecker delta function, and $(R_\sigma^{\gamma\alpha})_{ij}$ is the rotational transformation which rotates to the frame where the x' axis is aligned with the x_{ji}^α vector. This expression is simply a summed average of the individual pair-wise estimates for the gradient between each node pair defined in Eq. (19). Eq. (A.1) deviates from a standard SPH estimate in that we explicitly normalize each element by the total weighting which has gone into the pair-wise contributions to that element. This is necessary because each pair-wise contribution only increments the first column of the gradient tensor in the rotated ij frame, and therefore the sum of the total weighting applied to each element in the full tensor is unique and does not sum to unity. This explicit normalization also helps maintain accuracy in the gradient estimate near edges and over strong density gradients.

Appendix B. Estimating $\overline{\nabla(\nabla \cdot \mathbf{v})}$

We also use a low order (but robust) pair-wise formalism to estimate the second derivative $\nabla(\nabla \cdot \mathbf{v}) = \partial^2 v^\beta / (\partial x^\alpha \partial x^\beta)$ as

$$\overline{\left(\frac{\partial^2 v^\beta}{\partial x^\alpha \partial x^\beta} \right)_i} = \frac{\sum_j \frac{m_j}{\rho_j} \frac{\partial W_{ij}}{\partial x^\alpha} \frac{v_{ij}^\beta}{|r_{ij}|^2 + \epsilon^2 h_i^2}}{\sum_j \frac{m_j}{\rho_j} W_{ij}}, \quad (\text{B.1})$$

where ϵ is a small number to prevent division by zero errors. Just as in Eq. (A.1) we explicitly renormalize this expression by $\sum_j m_j / \rho_j W_{ij}$ in order to improve the accuracy of the result near strong density gradients and edges.

Appendix C. Rotational transformations

C.1. $R_\sigma^{\alpha\beta}$

In Section 4 we discuss the rotational transformation $R_\sigma^{\alpha\beta}$, which aligns the rotated x' axis with the $x_{ji}^\alpha = x_j^\alpha - x_i^\alpha$ vector between nodes i and j . In terms of the unit vector $\hat{x}_{ji}^\alpha = x_{ji}^\alpha / (x_{ji}^\beta x_{ji}^\beta)^{1/2}$, $R_\sigma^{\alpha\beta}$ in two dimensions is

$$R_\sigma = \begin{pmatrix} \hat{x}_{ji} & \hat{y}_{ji} \\ -\hat{y}_{ji} & \hat{x}_{ji} \end{pmatrix}, \quad (\text{C.1})$$

while in three dimensions

$$R_\sigma = \begin{pmatrix} \hat{x}_{ji} & \hat{y}_{ji} & \hat{z}_{ji} \\ -\hat{y}_{ji} / (\hat{x}_{ji}^2 + \hat{y}_{ji}^2)^{1/2} & \hat{x}_{ji} / (\hat{x}_{ji}^2 + \hat{y}_{ji}^2)^{1/2} & 0 \\ -\hat{x}_{ji} \hat{z}_{ji} / (\hat{x}_{ji}^2 + \hat{y}_{ji}^2)^{1/2} & -\hat{y}_{ji} \hat{z}_{ji} / (\hat{x}_{ji}^2 + \hat{y}_{ji}^2)^{1/2} & (\hat{x}_{ji}^2 + \hat{y}_{ji}^2)^{1/2} \end{pmatrix}. \quad (\text{C.2})$$

C.2. $R_S^{\alpha\beta}$

$R_S^{\alpha\beta}$ is the rotational transformation which diagonalizes the symmetric component of the velocity gradient (Eq. (21)). In two dimensions we define the rotation angle

$$\theta_S = \frac{1}{2} \tan^{-1} \frac{S_{ij}^{21}}{S_{ij}^{11} + S_{ij}^{22}}, \quad (\text{C.3})$$

and the rotational transformation is

$$R_S = \begin{pmatrix} \cos \theta_S & \sin \theta_S \\ -\sin \theta_S & \cos \theta_S \end{pmatrix}. \quad (\text{C.4})$$

In three dimensions we could in principle similarly determine R_S analytically, but in practice we simply use the Jacobi iterative diagonalization scheme to numerically determine R_S .

References

- [1] D.S. Balsara, *J. Comput. Phys.* 121 (1995) 357.
- [2] W. Benz, *The Numerical Modeling of Nonlinear Stellar Pulsations*, Kluwer Academic Publishers, Boston, 1990 p. 269.
- [3] J.C. Campbell, M.J. Shashkov, *J. Comput. Phys.* 172 (2001) 739.
- [4] A.E. Evrard, F.J. Summers, M. Davis, *Astrophys. J.* 503 (1994) 569.
- [5] R.A. Gingold, J.J. Monaghan, *MNRAS* 181 (1977) 375.
- [6] L.B. Lucy, *Astron. J.* 82 (1977) 1013.
- [7] L.G. Margolin, Technical Report UCRL-53882, Lawrence Livermore National Laboratory, 1988.
- [8] D.S. Miller, Ph.D. Thesis, UCRL-LR-115467, Lawrence Livermore National Laboratory, 1993.
- [9] J.J. Monaghan, *Annu. Rev. Astron. Astrophys.* 30 (1992) 543.
- [10] J.J. Monaghan, R.A. Gingold, *J. Comput. Phys.* 52 (1983) 374.
- [11] J.F. Navarro, W. Benz, *Astrophys. J.* 380 (1991) 320.
- [12] J.F. Navarro, M. Steinmetz, *Astrophys. J.* 478 (1997) 13.
- [13] J.F. Navarro, M. Steinmetz, *Astrophys. J.* 538 (2000) 477.
- [14] W.F. Noh, *J. Comput. Phys.* 72 (1987) 78.
- [15] J.M. Owen, J.V. Villumsen, P.R. Shapiro, H. Martel, *Astrophys. J. Suppl.* 116 (1998) 155.
- [16] I.J. Shchoenberg, *J. Approx. Theory* 2 (1969) 167.