

A General Class of Lagrangian Smoothed Particle Hydrodynamics Methods and Implications for Fluid Mixing Problems

Philip F. Hopkins¹*

¹*Department of Astronomy, University of California Berkeley, Berkeley, CA 94720*

Submitted to MNRAS, June, 2012

ABSTRACT

Various formulations of smooth-particle hydrodynamics (SPH) have been proposed, intended to resolve certain difficulties in the treatment of fluid mixing instabilities. Most have involved changes to the algorithm which either introduce “artificial” correction terms or violate what is arguably the greatest advantage of SPH over other methods: manifest conservation of energy, entropy, momentum, and angular momentum. Here, we show how a class of alternative SPH equations of motion (EOM) can be derived self-consistently from a discrete particle Lagrangian – guaranteeing manifest conservation – in a manner which tremendously improves treatment of these instabilities and contact discontinuities. Saitoh & Makino recently noted that the volume element used to discretize the EOM does not need to explicitly invoke the mass density (as in the “standard” approach); we show how this insight, and the resulting degree of freedom, can be incorporated into the rigorous Lagrangian formulation that retains ideal conservation properties and includes the “ ∇h ” terms that account for variable smoothing lengths. We derive a general EOM for any choice of volume element (particle “weights”) and method of determining smoothing lengths. We then specify this to a “pressure-entropy formulation” which resolves problems in the traditional treatment of fluid interfaces. Implementing this in a new version of the GADGET code, we show it leads to good performance in mixing experiments (e.g. Kelvin-Helmholtz & “blob” tests). And conservation is maintained even in strong shock/blastwave tests, where formulations without manifest conservation produce large errors. This also improves the treatment of sub-sonic turbulence, and lessens the need for large kernel particle numbers. The code changes are trivial and entail no additional numerical expense. This provides a general framework for self-consistent derivation of different “flavors” of SPH.

Key words: methods: numerical — hydrodynamics — instabilities — turbulence — cosmology: theory

1 INTRODUCTION

Smoothed particle hydrodynamics (SPH) is a method for solving the equations of hydrodynamics (in which pseudo-Lagrangian discretized mass elements are followed; Lucy 1977; Gingold & Monaghan 1977) which has found widespread application in astrophysical simulations and a range of other fields as well (for recent reviews, see Rosswog 2009; Springel 2010; Price 2012b).

The popularity of SPH owes to a number of properties: compared to many other methods, it is computationally inexpensive (fast), numerically very robust (stable), trivially allows the tracing of individual fluid elements (pseudo-Lagrangian), automatically produces improved resolution in high-density regions without the need for any ad-hoc pre-specified “refinement” criteria (inherently adaptive), is Galilean-invariant, couples properly and conservatively to N-body gravity schemes, exactly solves the particle continuity equation,¹ and has excellent conservation properties. The latter character stems from the fact that – unlike Eulerian grid methods – the SPH equations of motion (EOM) can be rigorously and exactly derived from a discretized particle Lagrangian, in a manner that guarantees manifest and simultaneous conservation of energy, entropy, linear momentum, and angular momentum (Springel & Hernquist 2002, henceforth S02).

However, there has been considerable discussion in the liter-

ature regarding the accuracy with which the most common SPH algorithms capture certain fluid mixing processes (particularly the Kelvin-Helmholtz instability; see e.g. Morris 1996; Dils 1999; Ritchie & Thomas 2001; Marri & White 2003; Okamoto et al. 2003; Agertz et al. 2007). Comparison between SPH and Euler (grid) methods shows that while agreement is quite good for supersonic flows, strong shock problems, and regimes with external forcing (e.g. gravity); “standard” SPH appears to suppress mixing in sub-sonic, thermal pressure-dominated regimes associated with contact discontinuities (Kitsionas et al. 2009; Price & Federrath 2010; Bauer & Springel 2011; Sijacki et al. 2011).² The reason is, in part, that in standard SPH the kernel-smoothed density enters the EOM, and so behaves incorrectly near contact discontinuities (introducing an artificial “surface tension”-like term) where the density is not differentiable.

A variety of “flavors” (alternative formulations of the EOM or kernel estimators) of SPH have been proposed which remedy this (see above and Monaghan 1997; Ritchie & Thomas 2001; Price 2008; Wadsley et al. 2008; Read et al. 2010; Read & Hayfield 2012; Abel 2011; García-Senz et al. 2012). These approaches share an essential common principle, namely recognizing that the pressure at contact discontinuities must be single-valued (effectively removing

* E-mail: phopkins@astro.berkeley.edu

¹ This is the continuity equation for a *discretized* particle field. Exactly solving the continuity equation for a continuous fluid, of course, requires infinite resolution or infinite ability to distort the Lagrangian particle “shape.”

² In fairness, we should emphasize that it has long been well-known that Eulerian grid codes, on the other hand, err on the side of *over-mixing* (especially when resolution is limited), and in fact this problem actually motivated some of the SPH work discussed above. This may, however, be remedied in moving-mesh approaches (though further study is needed; see e.g. Springel 2010).

the surface tension term). Some of these show great promise. However, many (though not all) of these formulations either introduce additional (potentially unphysical) dissipation terms and/or explicitly violate the manifest conservation and continuity solutions described above – perhaps the greatest advantages of SPH. This can lead to severe errors in problems with strong shocks or high-Mach number flows, limited resolution, or much larger gradients between phase boundaries (J. Read, private communication; see also the discussion in Price 2012b; Read & Hayfield 2012; Abel 2011). All of these regimes are inevitable in most astrophysically interesting problems.

Recently however, Saitoh & Makino (2012) (henceforth SM12) pointed out that the essential results of most of these flavors can be derived self-consistently in a manner that does properly conserve energy. The key insight is that the “problematic” inclusion of the density in the EOM (as opposed to some continuous property near contact discontinuities) arises because of the ultimately arbitrary choice of how to discretize the SPH volume element (typically chosen to be $\sim m_i/\rho_i$). Beginning with an alternative choice of volume element, one can in fact consistently derive a conservative EOM. They propose a specific form of the volume element involving internal energy and pressure, and show that this eliminates the surface tension term and resolves many problems of mixing near contact discontinuities.

In this paper, we develop this approach to provide a rigorous, conservative, Lagrangian basis for the formulation of alternative “flavors” of SPH, and show that this can robustly resolve certain issues in mixing. Although the EOM derived in SM12 conserves energy, it was derived from an ad-hoc discretization of the hydrodynamic equations, not the discrete particle Lagrangian. As such it cannot guarantee *simultaneous* conservation of energy and entropy (as well as momentum and angular momentum). And the EOM they derive is conservative only for constant SPH smoothing lengths (in time and space); to allow for adaptive smoothing (another major motivation for SPH), it is necessary to derive the “ ∇h ” terms which account for their variations. This links the volume elements used for smoothing in a manner that necessitates a Lagrangian derivation. And their derivation depends on explicitly evolving the particle internal energy; there are a number of advantages to adopting entropy-based formulations of the SPH equations instead.

We show here that – allowing for a different initial choice of which thermodynamic volume variable is discretized – an entire extensible class of SPH algorithms can be derived from the discrete particle Lagrangian, and write a general EOM for these methods (Eq. 12, our key result). We derive specific “pressure-energy” (Eq. 18) and “pressure-entropy” (Eq. 21) formulations of the EOM, motivated by the approaches above that endeavor to enforce single-valued SPH pressures near contact discontinuities. We consider these methods in a wide range of idealized and more complex test problems, and show that they *simultaneously* maintain manifest conservation while tremendously improving the treatment of contact discontinuities and fluid mixing processes.

2 THE SPH LAGRANGIAN & EQUATIONS OF MOTION

2.1 A Fully General Derivation

Following S02, note that the SPH equations can be derived self-consistently from the discrete particle Lagrangian

$$L(\mathbf{q}, \dot{\mathbf{q}}) = \frac{1}{2} \sum_{i=1}^N m_i \dot{\mathbf{r}}_i^2 - \sum_{i=1}^N m_i u_i \quad (1)$$

in the independent variables $\mathbf{q} = (\mathbf{r}_1, \dots, \mathbf{r}_N, h_1, \dots, h_N)$, namely the positions and smoothing lengths of each volume element/particle, and the internal energy per unit mass u . If the smoothing lengths h are constant, then the only independent variables are the \mathbf{r}_i and the equations of motion follow from $d(\partial L/\partial \dot{q}_i)/dt = \partial L/\partial q_i$. We require the derivatives of the u_i ; recalling that this is at constant entropy, so $du = -(P/m)dV$, we have:

$$\left. \frac{\partial u_i}{\partial q_i} \right|_A = -\frac{P_i}{m_i} \frac{\partial \Delta V_i}{\partial q_i} \quad (2)$$

where P_i is the pressure and ΔV_i is *some* estimator of the particle “volume.”

We clearly require some thermodynamic variable to determine P ; we can choose “which” to follow, for example internal energy u or entropy A . If we follow particle-carried u_i , then self-consistency with the thermodynamic equation above requires that we define the pressure in the above equation as $P_i = (\gamma - 1)u_i(m_i/\Delta V_i)$; if we follow the entropy A_i we must define $P_i = A_i(m_i/\Delta V_i)^\gamma$.

If h is allowed to vary, then we require some relation which makes it differentiable in order to make progress. This usually amounts to enforcing some condition on the effective “neighbor number” or “mass inside a kernel.” We stress that this language is somewhat misleading: it is not actually the case that there is exactly a certain neighbor number or mass inside the kernel, but rather that a continuous relation between h and some local volumetric quantity is enforced, for example $(4\pi/3)h_i^3\rho_i = M_{\text{kernel}} = m_i N_{\text{ngb}}$ (so that $h_i \propto \rho_i^{-1/3}$). Any such constraint (if continuous) is equally valid: motivated by the “effective neighbor number” approach we can define the constraint equation

$$\phi_i(\mathbf{q}) \equiv \frac{4\pi}{3} h_i^3 \frac{1}{\Delta \tilde{V}} - N_{\text{ngb}} = 0 \quad (3)$$

where $\Delta \tilde{V}$ is some continuous estimator of the “particle volume”; e.g. for $\Delta \tilde{V} = m_i/\rho_i$, we recover the approximate “mass in kernel” constraint.

We stress that $\Delta \tilde{V}$ does *not* need to be the same as ΔV ; one is the effective volume used to evolve the thermodynamics, the other is simply *any* continuous function used to define the h_i , so as to make them differentiable and thus allow us to include the appropriate “ ∇h ” terms in the EOM.

The equations of motion can then be determined from

$$\frac{d}{dt} \frac{\partial L}{\partial \dot{q}_i} - \frac{\partial L}{\partial q_i} = \sum_{j=1}^N \lambda_j \frac{\partial \phi_j}{\partial q_i} \quad (4)$$

where λ_i are the Lagrange multipliers. The second half of these equations ($q_i = h_i$) lead to the Lagrange multipliers

$$\lambda_i = -\frac{3P_i \Delta \tilde{V}^2}{4\pi h_i^3} \psi_i \quad (5)$$

$$\psi_i \equiv \frac{h_i}{3\Delta \tilde{V}} \frac{\partial \Delta V}{\partial h_i} \left[1 - \frac{h_i}{3\Delta \tilde{V}} \frac{\partial \Delta \tilde{V}}{\partial h_i} \right]^{-1} \quad (6)$$

Inserting this into the first half of the equations gives the EOM

$$m_i \frac{d\mathbf{v}_i}{dt} = \sum_{j=1}^N P_j \left[\nabla_i \Delta V_j + \psi_j \nabla_i \Delta \tilde{V}_j \right] \quad (7)$$

Now we require some way of defining “volumes.” In SPH this is done with respect to the kernel sum: for any particle-carried scalar value x_i , the x -weighted volume average $y_i = \bar{y}$ can be constructed as

$$y_i \equiv \sum_{j=1}^N x_j W_{ij}(h_i) \quad (8)$$

where $W_{ij}(h_i) = W_{ij}(|\mathbf{r}_{ij}|/h_i)$ is the smoothing kernel (discussed further below) as a function of $|\mathbf{r}_{ij}| \equiv |\mathbf{r}_i - \mathbf{r}_j|$. The corresponding x -weighted volume element can then be defined as $\Delta V_i \equiv x_i/y_i$. Note that for the choice $x_i = m_i$, we recover the familiar “standard” SPH choices of $y_i = \rho_i$ and $\Delta V_i = m_i/\rho_i$; but any other choice of x_i is in principle equally valid. For any well-behaved kernel, this makes the ΔV_i fully differentiable functions of \mathbf{r} and h :

$$\frac{\partial \Delta V_i}{\partial h_i} = -\frac{x_i}{y_i^2} \frac{\partial y_i}{\partial h_i}, \quad \nabla_i \Delta V_j = -\frac{x_j}{y_j^2} \nabla_i y_j \quad (9)$$

with

$$\nabla_i y_j = x_i \nabla_i W_{ij}(h_j) + \delta_{ij} \sum_{k=1}^N x_k \nabla_i W_{ik}(h_i) \quad (10)$$

$$\frac{\partial y_i}{\partial h_i} = -\sum_{j=1}^N \frac{x_j}{h_j} \left(3 W_{ij}(h_i) + \frac{|\mathbf{r}_{ij}|}{h_i} \frac{\partial W(|\mathbf{r}|/h)}{\partial (|\mathbf{r}|/h)} \Big|_{|\mathbf{r}_{ij}|/h_i} \right) \quad (11)$$

Putting all of this together, we obtain

$$m_i \frac{d\mathbf{v}_i}{dt} = -\sum_{j=1}^N x_i x_j \left[\frac{P_i}{y_i^2} f_{ij} \nabla_i W_{ij}(h_i) + \frac{P_j}{y_j^2} f_{ji} \nabla_i W_{ij}(h_j) \right] \quad (12)$$

$$f_{ij} \equiv 1 - \frac{\tilde{x}_j}{x_j} \left(\frac{h_i}{3 y_i} \frac{\partial y_i}{\partial h_i} \right) \left[1 + \frac{h_i}{3 y_i} \frac{\partial y_i}{\partial h_i} \right]^{-1} \quad (13)$$

Because the “potential” (thermal energy) in the Lagrangian here depends only on coordinate differences and is rotationally symmetric, the pair-wise force in Eq. 12 is automatically anti-symmetric (i.e. obeys Newton’s third law); energy, entropy, momentum, and angular momentum are all *manifestly* conserved, provided that smoothing lengths are adjusted to ensure the appropriate ϕ constraint (it is straightforward to verify this explicitly). The so-called “ ∇h ” terms, which depend on derivatives in time and space of the smoothing lengths, are implicitly included to all orders, via the f terms. The final EOM for any choice of x_i and \tilde{x}_i involves essentially identical information, constraints, and cost. In other words, because all of the formulations we consider are simply replacing the choice of particle-carried scalar in Eq. 12, they involve identical computational expense for otherwise equal gas conditions.

2.2 Formulations of the SPH Equations

2.2.1 Density-Entropy Formulation

If we take $x_i = \tilde{x}_i = m_i$ (giving $y_i = \tilde{y}_i = \bar{\rho}_i$, the kernel-averaged mass density estimator, and volume estimator $\Delta V_i = m_i/\rho_i$), and follow the entropy A_i (giving $P_i = \bar{P}_i = A_i \bar{\rho}_i^\gamma$), we obtain the EOM from S02:

$$\begin{aligned} \frac{d\mathbf{v}_i}{dt} &= -\sum_{j=1}^N m_j \left[\frac{f_i \bar{P}_i}{\bar{\rho}_i^2} \nabla_i W_{ij}(h_i) + \frac{f_j \bar{P}_j}{\bar{\rho}_j^2} \nabla_i W_{ij}(h_j) \right] \\ &= -\sum_{j=1}^N m_j \left[f_i A_i \bar{\rho}_i^{\gamma-2} \nabla_i W_{ij}(h_i) + f_j A_j \bar{\rho}_j^{\gamma-2} \nabla_i W_{ij}(h_j) \right] \\ f_i &= \left[1 + \frac{h_i}{3 \bar{\rho}_i} \frac{\partial \bar{\rho}_i}{\partial h_i} \right]^{-1}, \quad \bar{\rho}_i \equiv \sum_{j=1}^N m_j W_{ij}(h_i) \end{aligned} \quad (14)$$

Note that for adiabatic evolution, we require no energy equation, since entropy is followed; for a specific energy *defined* as $u \equiv \bar{P}_i/[(\gamma-1)\bar{\rho}_i] = (\gamma-1)^{-1} A_i \bar{\rho}_i^{\gamma-1}$, energy conservation is manifest from the EOM above.

As discussed in § 1, this formulation is known to have trouble treating certain contact discontinuities. Because the only

volumetric quantity that enters is the density, this fails when the densities are no longer differentiable, even when pressure is smooth/constant. Specifically, consider a contact discontinuity, $\rho_1 c_1^2 = \rho_2 c_2^2$ (where quantities ‘1’ and ‘2’ are on either side of the discontinuity). As we approach the discontinuity, the kernel-estimated density must trend to some average because it spherically averages over both “sides,” $\rho \rightarrow \langle \rho \rangle$, but the particle-carried sound speeds c remain distinct, so the pressure is now multi-valued, with a ‘pressure blip’ of magnitude $\sim (\rho_{\max}/\rho_{\min}) P_{\text{true}}$ appearing. This has a gradient across the central smoothing length, causing an artificial, repulsive “surface tension” force that suppresses interpenetration across the discontinuity.

2.2.2 Pressure-Energy Formulation

Instead consider $x_i = \tilde{x}_i = (\gamma-1)U_i \equiv (\gamma-1)m_i u_i$, proportional to the particle internal energy. Now, $y_i = \tilde{y}_i$ is by definition a direct kernel-averaged pressure estimator $y_i = \bar{P}_i$, and $\Delta V_i = (\gamma-1)m_i u_i/P_i$. This means that the pressure itself is now the directly kernel-averaged quantity entering the EOM and is therefore always single-valued. So long as the pressure is smooth/differentiable (regardless of how the density varies), the EOM should be well-behaved. For this choice, we obtain:³

$$\begin{aligned} \frac{d\mathbf{v}_i}{dt} &= -\sum_{j=1}^N (\gamma-1)^2 m_j u_i u_j \left[\frac{f_i}{\bar{P}_i} \nabla_i W_{ij}(h_i) + \frac{f_j}{\bar{P}_j} \nabla_i W_{ij}(h_j) \right] \\ f_i &= \left[1 + \frac{h_i}{3 \bar{P}_i} \frac{\partial \bar{P}_i}{\partial h_i} \right]^{-1}, \quad \bar{P}_i \equiv \sum_{j=1}^N (\gamma-1) m_j u_j W_{ij}(h_i) \end{aligned} \quad (15)$$

Recall, we now need to evolve the energy explicitly. From Eq. 2, $du/dt = -(P/m)(d\Delta V/dt)$, and the evolution of the volume element ΔV just follows from the Lagrangian continuity equation, so we obtain

$$\frac{du_i}{dt} = \sum_{j=1}^N (\gamma-1)^2 m_j u_i u_j \frac{f_i}{\bar{P}_i} (\mathbf{v}_i - \mathbf{v}_j) \cdot \nabla_i W_{ij}(h_i) \quad (16)$$

It is straightforward to see that this guarantees explicit energy conservation; entropy conservation is implicit and also easily verified.

There are, however, significant drawbacks to the choice of $\tilde{x}_i = x_i$ ($\Delta \tilde{V} = \Delta V$), in which case we are implicitly defining h such that $h_i^3 \propto u_i/\bar{P}_i$. This is, by itself, perfectly valid and is easily solved by the same bisector method as in the “standard” ($\Delta \tilde{V} = m/\rho$) formulation. However, in practice, the particle u_i values vary much more widely than the m_i . This leads to some potential problems. First, if there is large variation in u_i , convergence in h_i can become quite expensive. Second, again under circumstances with large disorder, the required h_i can become very large, leading to an effective loss of resolution. Third and most problematic, under some circumstances the constraint ϕ can have multiple solutions; in this case if h_i “jumps,” it is no longer continuously differentiable, and so exact energy conservation is broken.

³ Eq. 15 is similar to the EOM derived in SM12, itself identical to that derived earlier in Ritchie & Thomas (2001) from purely heuristic arguments. These do, after all, motivate the derivation here. However there are two key differences. First, we derive and include the ∇h terms ($f_i = 1$ in SM12), necessary for conservation if h varies. Second, the ∇W terms enter differently (with different multipliers and indices). This stems from the Lagrangian derivation and is necessary – even for constant h – for the EOM to *simultaneously* conserve energy and entropy (i.e. to properly advect the thermodynamic volume element).

An obvious alternative is to use $\tilde{x} \equiv 1$, i.e. $\Delta\tilde{V} \equiv 1/\bar{n}_i$, where \bar{n}_i is the “particle number density”

$$\bar{n}_i = \bar{y}_i(\tilde{x} = 1) \equiv \sum_{j=1}^N W_{ij}(h_i) \quad (17)$$

This restores the effective “number of neighbors” criterion for h_i , and is always well-behaved since all particles are weighted equally. If we do this, the EOM become:

$$\begin{aligned} \frac{d\mathbf{v}_i}{dt} &= - \sum_{j=1}^N (\gamma - 1)^2 m_j u_i u_j \left[\frac{f_{ij}}{\bar{P}_i} \nabla_i W_{ij}(h_i) + \frac{f_{ji}}{\bar{P}_j} \nabla_i W_{ij}(h_j) \right] \\ f_{ij} &= 1 - \left(\frac{h_i}{3(\gamma - 1)\bar{n}_i m_j u_j} \frac{\partial \bar{P}_i}{\partial h_i} \right) \left[1 + \frac{h_i}{3\bar{n}_i} \frac{\partial \bar{n}_i}{\partial h_i} \right]^{-1} \end{aligned} \quad (18)$$

Note that this is just the previous equation with $f_i \rightarrow f_{ij}$; in other words, the EOM are identical up to the “ ∇h ” corrections, which is what we expect, since the only function of the $\Delta\tilde{V}$ term is to determine how the h_i evolve. Trivially, then, the energy equation is also the same as above but with $f_i \rightarrow f_{ij}$.

As discussed in SM12, because the volumetric quantity used in the EOM here is now directly the kernel-estimated pressure (instead of the density), this formulation automatically *gaurantees* that pressure is single-valued at contact discontinuities, and so removes the pressure “blip” and surface tension force. The equations will now be well-behaved so long as pressure is smooth. This is true by definition in contact discontinuities; it is of course not true at shocks, but neither (typically) is the density constant there – so we do not lose any desirable behaviors of the density-entropy formulation. In either case, we require some artificial viscosity to treat shocks.

2.2.3 Pressure-Entropy Formulation

If we wish to retain a direct kernel-estimate of the pressure entering the EOM, but formulate this in terms of entropy, we must instead consider $x_i = m_i A_i^{1/\gamma}$. In this case,⁴ we obtain by definition (from the consistency requirement for P_i)

$$\bar{P}_i = y_i^\gamma = \left[\sum_{j=1}^N m_j A_j^{1/\gamma} W_{ij}(h_i) \right]^\gamma \quad (19)$$

If we also assume $\tilde{x}_i = x_i$, the EOM become

$$\begin{aligned} \frac{d\mathbf{v}_i}{dt} &= - \sum_{j=1}^N m_j (A_i A_j)^{1/\gamma} \left[\frac{f_{ij} \bar{P}_i}{\bar{P}_i^{2/\gamma}} \nabla_i W_{ij}(h_i) + \frac{f_{ji} \bar{P}_j}{\bar{P}_j^{2/\gamma}} \nabla_i W_{ij}(h_j) \right] \\ f_i &= \left[1 + \frac{h_i}{3\bar{P}_i^{1/\gamma}} \frac{\partial \bar{P}_i^{1/\gamma}}{\partial h_i} \right]^{-1} \end{aligned} \quad (20)$$

Note however that using $\tilde{x}_i = x_i$ implies $\Delta\tilde{V} = m_i (A_i/\bar{P}_i)^{1/\gamma}$; as in the previous section, this can introduce problems of convergence and diffusion. Therefore, instead consider as before $\tilde{x}_i = 1$

⁴ This choice of x_i may seem a bit strange, but in fact this is the only self-consistent “entropy formulation” which directly evaluates the pressure. If we simply substituted $u_j = A_j \bar{P}_j^{\gamma-1}/(\gamma-1)$ in $x_i = (\gamma-1)m_i u_i$, we would re-introduce the density $\bar{\rho}$ (which we are trying to avoid in this formulation of the equations); we could instead define $u_j = A_j (\bar{P}_j/A_j)^{\gamma-1}/(\gamma-1)$, but this involves \bar{P}_j in its own definition and would require a prohibitively expensive iterative solution over all particles every timestep.

($\Delta\tilde{V} = 1/\bar{n}_i$). This gives:

$$\begin{aligned} \frac{d\mathbf{v}_i}{dt} &= - \sum_{j=1}^N m_j \left\{ \left[\left(\frac{A_j}{A_i} \right)^{1/\gamma} + f_i \right] \left(\frac{A_j}{\bar{P}_i} \right)^{2/\gamma} \bar{P}_i \nabla_i W_{ij}(h_i) \right. \\ &\quad \left. + \left[\left(\frac{A_i}{A_j} \right)^{1/\gamma} + f_j \right] \left(\frac{A_j}{\bar{P}_j} \right)^{2/\gamma} \bar{P}_j \nabla_i W_{ij}(h_j) \right\} \end{aligned} \quad (21)$$

with

$$f_i = \frac{h_i}{3A_i^{1/\gamma} m_i \bar{n}_i} \frac{\partial \bar{P}_i^{1/\gamma}}{\partial h_i} \left[1 + \frac{h_i}{3\bar{n}_i} \frac{\partial \bar{n}_i}{\partial h_i} \right]^{-1} \quad (22)$$

As for the density-entropy formulation, we explicitly evolve the entropy so for adiabatic evolution require no additional evolution equation.

This formulation is very similar to the pressure-energy formulation, (and has the identical advantages of good behavior at contact discontinuities). The only difference is the free choice of thermodynamic variable. This formulation trivially conserves entropy, and manifestly conserves energy to machine differencing accuracy if constant timesteps are used (the choice of pressure-energy or pressure-entropy formulation can lead to some differences when adaptive timesteps are used, but we show these are generally small). It is largely a matter of convenience and minor computational expense which method is preferred.

When the pressure is smooth and there is good particle order, the $f_i \approx 0$ here, which means our choice of how to regularize h is unimportant, and no spurious “surface tension” force is introduced. For the choice $\tilde{x} = 1$, the correction terms remain well-behaved even if there is large particle disorder in A_i , critical to stability in simulations when heating/cooling are included and entropy is no longer conserved. Another useful feature here is the following: imagine the case where there is large particle disorder so $P_j \gg P_i$ and $A_j \gg A_i$. Since the A terms enter as multiplicative pre-factors (we can re-write the latter formulation with $\Delta\tilde{V} = 1/\bar{n}_i$ in the same manner), their difference does not introduce errors into the sum; gradient errors will arise from differencing the P_i terms, but for $\gamma = 5/3$, these enter only as $P_i^{-1/5}$, so differencing errors are greatly suppressed.

3 ADDITIONAL SIMULATION PROPERTIES

3.1 Entropy & Artificial Viscosity

As is standard in SPH, the algorithm is inherently inviscid and some artificial viscosity must be included to properly capture shocks. However, we include a more sophisticated treatment of the artificial viscosity term (as compared to S02 and “standard” GADGET; which follow Gingold & Monaghan 1983), following Morris & Monaghan (1997) with a Balsara (1989) switch. This includes a particle-by-particle artificial viscosity that grows rapidly in strong shocks and rapidly decays away from shocks (to a minimum $\alpha = 0.05$), reducing numerical dissipation by more than an order of magnitude away from shocks compared to the previous constant artificial viscosity prescription. For detailed comparison of the viscosity algorithms, we refer to Cullen & Dehnen (2010).

3.2 Thermodynamic Evolution & Timestep Criteria

For all problems discussed here, we employ the GADGET adaptive timestep algorithm, which dramatically reduces the computational expense for almost all interesting problems (relative to using a constant simulation timestep). However, as pointed out in Saitoh & Makino (2009) and developed further in Durier & Dalla Vecchia (2012), in problems with very high Mach number shocks/bulk

flows, the standard adaptive timestepping can lead to problems if particles with long timesteps interact suddenly mid-timestep with material evolving on much shorter timesteps. Fortunately this is easily remedied, and we do so by implementing a timestep limiter identical to that in Durier & Dalla Vecchia (2012). At all times, any active particle informs its neighbors of its timesteps and none are allowed to have a timestep > 4 times that of a neighbor; and whenever a timestep is shortened (or energy is injected in feed-back) particles are “activated” and forced to return to the timestep calculation as soon as possible.

3.3 Smoothing Kernel

Our derivation of the EOM allows for an arbitrary choice of SPH smoothing kernel W , so long as it is differentiable. This choice can have a significant effect on some test problems via its effect on the resolution pressure gradient errors (see e.g. Morris 1996; Dils 1999; Read et al. 2010, and discussion below). We have experimented with a wide range of possible kernel shapes, following the ~ 10 discussed in Fulk & Quinn (1996) and Hongbin & Xin (2005), the triangular kernels proposed in Read et al. (2010), and the variant Wendland kernels proposed in Dehnen & Aly (2012). However our intention here is not to study SPH kernels; for simplicity we therefore adopt a standard quintic spline kernel with $N_{\text{NGB}} = 128$ neighbors in all tests shown here (unless otherwise noted). This is the ‘optimal’ spline kernel suggested in both Hongbin & Xin (2005) and Dehnen & Aly (2012) and has effective resolution equal to a cubic spline with 34 neighbors (but significantly higher accuracy). In most cases, we obtain similar results using a lower-order cubic spline⁵ with $N_{\text{NGB}} = 32$; but we discuss where this is not the case. We do not see qualitative improvement in the specific tests here using yet higher-order kernels and/or increasing neighbor number as high as $N_{\text{NGB}} \approx 500$.

3.4 Density Estimation in Density-Independent SPH

An estimate of the density is often required for other calculations (e.g. cooling), even if it is not needed for the EOM. This can still be directly estimated from the standard SPH kernel, as ρ_i above (i.e. a kernel sum with $x_i = m_i$). But in principle a density could also be inferred or estimated as $\rho_{th} \approx m_i / \Delta V_i$ for the thermodynamic ΔV “pressure formulations” (i.e. from the combination of the variables \bar{P}_i and A_i or u_i). However, in “mixed” regions near contact discontinuities, this will lead to multi-valued densities in neighboring particles (since entropies are still conserved at the particle level). This relates to the fact that the two densities represent physically distinct quantities. The former, direct kernel ρ_i ($x_i = m_i$) is simply the volume-average mean mass density at the particle location. The latter (thermodynamically-inferred ρ_{th}) is the energy or entropy-weighted mean mass density (in the pressure-energy or pressure-entropy formulations, respectively), averaged over all neighbors within the kernel. Because the cooling is typically calculated on a particle-by-particle basis, we find the former estimator is more stable and appropriate in most applications. However, there may be situations where alternative weightings for the density estimator are optimal, and a more complete treatment of mixing may

require a mechanism for equilibrating entropies inside the kernel and generating mixing entropy.

4 TEST PROBLEMS

4.1 Sod Shock Tube

If we repeat the standard one-dimensional Sod shock tube test using the parameter sets in SM12, our results are indistinguishable from their Figs. 2-4. As shown therein (also their Figs. 1 & 10), for the “pressure-energy” formulation (and identically for the “pressure-entropy” formulation here), the “pressure blip” which appears at the contact discontinuity in the standard “density-entropy” formulation is removed. To the extent that there are differences in the results with different formulations, they are much more pronounced in the Sedov blastwave test below; we therefore focus our attention on this test.

4.2 Strong Sedov-Taylor Blastwaves

Here we consider an extreme Sedov-Taylor blastwave, with very large mach number, designed to be a powerful test of conservation. A box of side-length 6kpc is initially filled with 128³ equal-mass particles at constant density $n = 0.5 \text{ cm}^{-3}$ and temperature 10K; $6.78 \times 10^{46} \text{ J}$ of energy is added to the central 64 particles in a top-hat distribution. This triggers a blastwave with initial Mach number ~ 1000 , which we compare at 20Myr, where the shock front should be at $r \approx 1.19 \text{ kpc}$. In this test and below, unless otherwise specified, we assume a $\gamma = 5/3$ gas equation of state.

First we compare the “fully conservative” algorithms we derive above. In “standard” SPH (density-entropy formulation; Eq. 14), the analytic solution is reproduced very well up to the SPH smoothing/resolution limit. The narrow shock jump is not perfectly resolved, for example, but averaged over the same smoothing, the agreement is excellent between analytic solution and mean particle values. There is also some scatter in particle properties (most notably velocity), but it is generally narrow. The largest deviation from the analytic solution is in the post-shock “ringing” in the velocity field, a well-known effect. Note that the behavior at small radii $\lesssim 0.5 \text{ kpc}$ is noisy simply because the extremely low density means there are no particles here. As expected, the conservation properties are excellent as well; energy is conserved to a part in $\sim 10^4$; similar conservation obtains for entropy, energy, momentum, and angular momentum. In fact the conservation errors are overwhelmingly dominated by the adaptive time-stepping scheme (which violates conservation by not always evolving mutually interacting particles at the same timesteps), not the formulation of the EOM. If we simply force constant, global timesteps, the conservation errors are reduced to machine accuracy (however this comes at great numerical expense).⁶

To compare, the pressure-energy formulation ($x_i = (\gamma - 1)U_i$ as Eq. 18, using the neighbor number density volume element $\bar{x}_i = 1$ to define h) agrees very well, and also gives excellent conservation. The particle noise/scatter is slightly larger, most noticeably in the post-shock temperature. This occurs because of some “mixing” of thermal properties in the kernel average from higher-energy particles. However the post-shock ringing in the velocity solution (although still present) is reduced.

The pressure-entropy formulation ($x_i = m_i A_i^{1/\gamma}$ with $\bar{x}_i = 1$,

⁵ It is well-known that the “standard” Schoenberg (1946) spline kernels become unstable for arbitrarily high neighbor number; the “correct” way to increase N_{NGB} to better sample the kernel and reduce gradient errors is to move to higher-order kernels of $\mathcal{O}(n) \approx -1 + (1.0 - 1.2)N_{\text{NGB}}^{1/3}$. Done properly, this allows increasing N_{NGB} and accuracy without decreasing the effective spatial resolution (albeit at additional computational expense).

⁶ We do confirm the point discussed in § 3.2, that motivates our timestepping algorithm: without care in “signaling” when adaptive timesteps are taken, conservation errors can be severe.

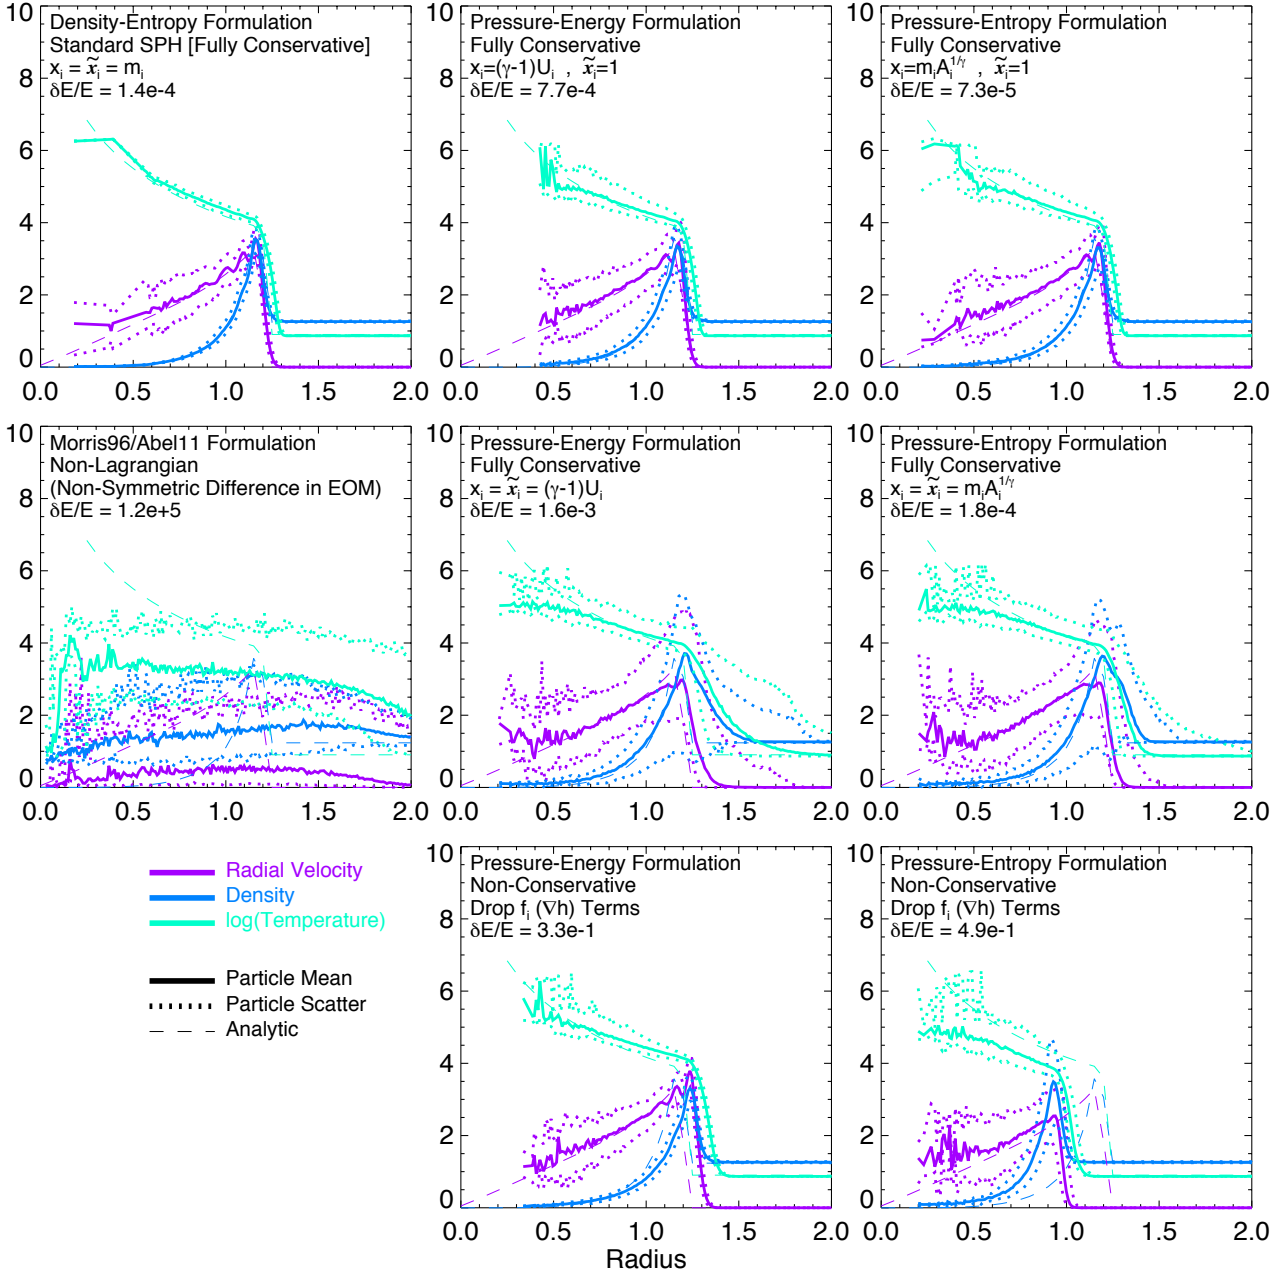


Figure 1. Strong Sedov-Taylor blastwave test. We compare the solution at fixed time and resolution, with different SPH algorithms discussed in the text. We plot radial velocity, density, and temperature (for clarity, we plot the median and 1% – 99% interval of particle values binned in radial intervals $\Delta r = 0.01$), with the analytic solution. We measure the accuracy of energy conservation $\delta E/E \equiv |E(t) - E(t=0)|/E(t=0)$. *Top Left:* Standard SPH (density-entropy; $x_i = \tilde{x}_i = m_i$); Eq. 14. *Top Center:* Lagrangian Pressure-Energy formulation ($x_i = (\gamma - 1)U_i$) with particle-number density based smoothing lengths ($\tilde{x}_i = 1$); Eq. 18. *Top Right:* Lagrangian Pressure-Entropy formulation ($x_i = m_i A_i^{1/3}$) with particle-number density based smoothing lengths ($\tilde{x}_i = 1$); Eq. 21. Conservation and accuracy is excellent up to the resolution limits (lack of particles dominates the noise at small radii), in all three algorithms. The pressure formulations slightly increase the particle scatter in temperature, but reduce the post-shock “ringing” of the velocity solution. Conservation errors in all three cases are overwhelmingly dominated by the adaptive timesteps, not EOM. *Middle Left:* Non-Lagrangian SPH as considered in Morris (1996); Abel (2011); this algorithm minimizes linear pressure errors and removes the surface tension term, but violates conservation of energy and momentum. Conservation errors grow exponentially and dominate the solution. *Middle Center:* Lagrangian Pressure-Energy formulation with $\tilde{x}_i = x_i$; Eq. 15. Although conservation is maintained, this choice of \tilde{x}_i (how to regularize the h_i) leads to enormous numerical diffusion and particle disorder (“spreading” the shock location over a large radius and affecting the pre-shock gas). *Middle Right:* Lagrangian Pressure-Entropy formulation with $\tilde{x}_i = x_i$; Eq. 20. Again, this leads to severe diffusion. *Bottom Center:* Pressure-Energy formulation without the ∇h terms ($f_i = 1$ in Eq. 15). Variations in h now cause order-unity energy conservation errors, leading to the shock being in the wrong position. *Bottom Right:* Pressure-Entropy formulation without the ∇h terms. Again order-unity conservation errors appear and the shock evolves incorrectly.

Eq. 21) is essentially identical to the pressure-energy formulation here (the slightly better conservation owes to how the sound speeds enter the timestep signaling scheme). This is not surprising – the two represent essentially identical information but just choose to explicitly follow different thermodynamic variables.

As discussed in § 2.1, the EOM for the pressure-energy and pressure-entropy formulations take a simpler form if we assume $\tilde{x}_i = x_i$, as in standard SPH, but this can lead to problems. We show that here, by considering these implementations (Eq. 15 & 20 for the pressure-energy and pressure-entropy formulations, respectively). Recall, these are still fully conservative Lagrangian formulations; as a result we still see good energy conservation. However, forcing the smoothing lengths to evolve not just with particle number density but with local thermodynamic quantities leads to occasional enormous “super-smoothing” that is both computationally expensive and introduces enormous numerical diffusion. We see that here, where the upper envelope of particles and even the mean are biased in the pre-shock medium (information from the post-shock region has clearly been over-smoothed into the regions here). Although some mean and post-shock quantities are well-behaved, this algorithm has introduced far too much numerical mixing.

It is also instructive to see what happens if we drop the ∇h terms in these formulations (returning to $\tilde{x}_i = 1$). The EOM will now violate manifest energy conservation to the level that the smoothing lengths vary over an individual smoothing length (in a smooth medium, this correction should vanish at infinite resolution, but it will never vanish in shocks if variable h are allowed). Indeed we now see order-unity energy errors. Although the qualitative solution appears similar, the pressure-energy and pressure-entropy solutions gain and lose energy, respectively, and so the shock is simply in the wrong position. This occurs in other, similar non-conservative formulations as well, for example see the discussion of problems with Ritchie & Thomas (2001) & ‘OSPH’ methods in Read & Hayfield (2012, Appendix A).

Finally, we consider an algorithm which is manifestly non-conservative. Specifically, we consider the Morris (1996) formulation of the pressure derivative, where the kernel sum for the EOM is over $(P_j - P_i)/(\rho_i \rho_j) \nabla W_{ij}$. This is very similar (although not identical) to the EOM proposed in Abel (2011) as well. It is possible to show that this formulation actually eliminates the leading-order gradient errors associated with the “standard” SPH EOM (see e.g. Price 2012b). However, the cost is manifest resolution-level violation of conservation (of energy, momentum, and entropy). If we adopt our standard initial conditions with this algorithm, we see that the conservation errors grow exponentially and quickly swamp the real solution. The problem, as described in Price (2012b) is in the non-linear terms that maintain particle order in SPH.⁷

4.3 Kelvin-Helmholtz Instabilities

We next consider a (three-dimensional) Kelvin-Helmholtz (KH) test. The initial conditions are taken from the Wengen multiphase test suite⁸ and described in Agertz et al. (2007); Read et al. (2010).

⁷ Abel (2011) do show that the behavior of Sedov blastwaves in their formulation is tremendously improved if the initial conditions are modified so that the blastwave does not start from a point injection or top-hat particle distribution, but from an *already-developed* smaller blastwave with pre-initialized, resolved pressure gradients. However that is not the particular test here, and – as they caution – is not often the case in astrophysical systems where e.g. early-stage SNe explosions are unresolved.

⁸ Available at <http://www.astrosim.net/code/doku.php>




Figure 2. Specific entropy map of Kelvin-Helmholtz instabilities at time 2 times the linear K-H growth timescale τ_{KH} (top), and at $8\tau_{KH}$ (bottom). Left: Standard (density-entropy) SPH. Note the “surface layer” surrounding the curls and the breakup into oil-and-water style blobs at later times. Right: Pressure-entropy formulation ($\tilde{x}_i = 1$).

Briefly, in a periodic box with size 256, 256, 16 kpc in the x, y, z directions (centered on 0, 0, 0), respectively, $\approx 10^6$ equal-mass particles are initialized in a cubic lattice, with density, temperature, and x -velocity $= \rho_1, T_1, v_1$ for $|y| < 4$ and $= \rho_2, T_2, v_2$ for $|y| > 4$, with $\rho_2 = 0.5\rho_1$, $T_2 = 2.0T_1$, $v_2 = -v_1 = 40 \text{ km s}^{-1}$. The values for T_1 are chosen so the sound speed $c_{s,2} \approx 8|v_2|$; the system has constant initial pressure. To trigger instabilities, a sinusoidal velocity perturbation is applied to v_y near the boundary, with amplitude $\delta v_y = 4 \text{ km s}^{-1}$ and wavelength $\lambda = 128 \text{ kpc}$.

Fig. 2 compares the resulting behavior in the “standard” density-entropy formulation of SPH (Eq. 14), and the pressure-entropy formulation (Eq. 21) proposed here. Because of the severe diffusion seen in the Sedov test associated with the choice of $x_i = \tilde{x}_i$ for the pressure-entropy formulation, we focus on the formulation with the “neighbor number density” volume element used to define h , i.e. $\tilde{x}_i = 1$. However, for the sub-sonic, pressure-equilibrium systems simulated in this test, the choice of \tilde{x}_i makes little difference (highlighting the importance of our strong shock tests). We also do not further separately consider the pressure-energy formulation, as it gives very similar results in our subsequent tests. As expected from the discussion in § 1, the density-entropy formulation poorly captures the KH instability. The surface tension term forms a sharp boundary layer between the two phases, which suppresses all the small-wavelength mixing and leads at late times to the breakup of the “rolls.” This ultimately produces blobs that resemble an oil-and-water morphology (a system with physical surface tension).

In the pressure-entropy formulation, however, the behavior is dramatically improved. The long-wavelength mode grows on the correct (linear) KH timescale, but we can also plainly see the growth of modes at all wavelengths (along the “edge” of the density surface) down to the kernel scale. At later times, we resolve $\sim 3 - 5$ full “wraps” of the rolls while the standard SPH solution breaks up. Most important, there is obviously actual mixing occur-




Figure 3. Comparison of the Kelvin-Helmholtz behavior at $t = 12 \tau_{\text{KH}}$ on artificial viscosity in the Pressure-Entropy formulation. *Left:* Our standard, by-particle time-dependent artificial viscosity following Morris & Monaghan 1997. *Right:* The original GADGET simplified constant artificial viscosity prescription (as Gingold & Monaghan 1983 with a Balsara 1989 switch). Although very important for some behaviors, it has little effect on KH instabilities within this formulation.




Figure 4. Comparison of the Kelvin-Helmholtz behavior at $t = 3 \tau_{\text{KH}}$ on the smoothing kernel in the Pressure-Entropy formulation. *Top Left:* Quartic core-triangle kernel with neighbor number $N_{\text{NGB}} = 442$ from Read et al. (2010). *Top Right:* Cubic core-triangle with $N_{\text{NGB}} = 128$. *Bottom Left:* Quintic spline with $N_{\text{NGB}} = 128$. *Bottom Right:* Cubic spline with $N_{\text{NGB}} = 32$.

ring “inside” the rolls. Direct comparison to high-resolution results from grid codes (both fixed-grid and adaptive mesh solutions) as well as Godunov SPH and moving-mesh methods, using identical initial conditions, show that the result here is quite similar (compare e.g. Read et al. 2010 Fig. 6, or Murante et al. 2011 Figs. 5-9).

In Fig. 3, we repeat this experiment with the pressure-entropy formulation, but replace our standard treatment of artificial viscosity (see § 3.1) with a more simplified treatment (using a constant artificial viscosity for all particles and times); the latter is known to produce significantly more numerical dissipation away from shocks. Although it is well-known that this can produce significant differences in some situations (e.g. sub-sonic turbulence and rotating shear flows; see Cullen & Dehnen 2010; Price & Federrath 2010; Bauer & Springel 2011), it has little effect on this particular test. This owes in part to the implementation of the Balsara (1989) switch in both cases, which reduces the artificial shear viscosity; without this, we see significantly greater damping of shear motions.

In Fig. 4, we again repeat this experiment with the pressure-




Figure 5. Comparison of Kelvin-Helmholtz behavior with different initial conditions in the Pressure-Entropy formulation at $t \approx 1 \tau_{\text{KH}}$ (top) and $t \approx 3 \tau_{\text{KH}}$ (bottom) on smoothing kernel (*left:* quintic spline with $N_{\text{NGB}} = 128$; *right:* cubic spline with $N_{\text{NGB}} = 32$). Compared to Fig. 2-4, we modify the initial conditions by multiplying the particle number, density, sound speed, pressure, shear velocity, and initial perturbation amplitude by factors of $= 0.5, 2.0, 2.0, 8.0, 0.5, 0.5$, respectively. All of these changes increase the ratio of particle noise and pressure gradient errors relative to the KH growth. In this limit, the cubic spline with $N_{\text{NGB}} = 32$, which involves larger gradient errors, is barely able to capture the instability; however the quintic spline with $N_{\text{NGB}} = 128$ does well. “Standard” SPH completely fails to develop any “curls” in this limit.

entropy formulation, but vary the smoothing kernel. We compare two spline kernels, for the neighbor numbers advocated in Price (2012b); Dehnen & Aly (2012), and the “core-triangle” kernels with sharp central peaks, proposed in Read et al. (2010). For the standard initial conditions here, we see only subtle differences over a wide range from even the simplest kernel with 32 neighbors up through more than an order of magnitude higher neighbor number.⁹ By explicitly eliminating the “surface tension” term, while maintaining exact conservation, the pressure-entropy formulation appears to significantly reduce the importance of kernel-level pressure gradient errors in obtaining the correct KH solution (compare e.g. the more significant kernel dependence found for standard but non-conservative density-entropy SPH in Read et al. 2010, Figs. 4 & 6). We have also tested the Wendland kernels proposed in Dehnen & Aly (2012) and find (as they do) essentially identical performance to our spline results for the neighbor numbers here (although the kernels proposed there show greatly improved stability properties at higher neighbor number).

However, in Fig. 5 we again examine the effects of different kernels (with the pressure-entropy formulation), but with different initial conditions. We reduce the particle number, shear velocity, and initial perturbation amplitude by factors of 2 and double the initial sound speed; these changes all reduce the magnitude of the initial KH instability and its growth rate, relative to the particle noise and error terms stemming from the kernel sum in the SPH pressure gradients that enter the EOM. This is designed to be challenging (even for grid codes). With this much weaker seeding and noiser

⁹ Note that, at the same error level, the core-triangle kernels require significantly higher neighbor number, because they introduce a sharp bias towards the central particles.




Figure 6. Comparison of KH instabilities at $t \approx 2\tau_{\text{KH}}$ (top) and $t \approx 10\tau_{\text{KH}}$ (bottom), for the “standard” (density-entropy) SPH (left) and pressure-entropy (right; with $\tilde{x}_i = 1$) SPH formulations. The initial conditions are as Fig. 2, but the initial density contrast is increased from a factor of 2 to a factor of 20, with the initial particles no longer being equal mass but 4 times more massive in the initial high-density region. Sharp boundary layers that lead to “gloopy” morphology are evident in standard SPH. In Pressure-Entropy SPH it remains well-behaved, although hints of “gloopiness” appear in the transition between linear growth and fully non-linear instability.

particle distribution, we begin to see a dependence on the smoothing kernel. The simplest $N_{\text{NGB}} = 32$ cubic spline, with the smallest neighbor number, leads to particle noise in the pressure gradients comparable to the actual signal. While the instability is still (barely) captured, the rolls are “too thin” and end up being sheared before they reach the proper height and can wrap appropriately. However, our standard quintic spline with $N_{\text{NGB}} = 128$ performs well, recovering all the key behaviors (note that the non-linear behavior is different here than in the previous test, as it should be for the different pressure and shear velocities). Going to still higher resolution or varying the kernel at higher N_{NGB} gives well-converged results at this point.

In Fig. 6, we repeat our KH test again but multiply the initial density contrast by a factor of 10, and instead of using constant-mass particles we use particles with masses a factor of 4 larger in the high-density region. As discussed in Read & Hayfield (2012), many proposed alternative formulations of SPH, designed to improve fluid mixing, fail in this regime (see e.g. Fig 6 in Read et al. 2010 and Figs. A1 & E1 in Read & Hayfield 2012). And we see an even more pronounced “boundary layer” separating the phases in the standard density-entropy SPH formulation. This occurs because the higher density contrast exacerbates any (even small residual) surface tension term, and the multi-mass particles increase particle disorder and leading-order errors in the pressure gradient estimator. Multi-mass particles also make it critical to have a well-behaved criterion for smoothing lengths, and increase the errors from neglecting ∇h terms. But we see that the pressure-entropy formulation remains well-behaved in this case. There is some increased hint of “gloopiness” as the instability transitions between linear and non-linear growth, but at least some of this is because of the (correct) slower growth of the small-wavelength modes.

Finally, it is worth noting (though perhaps as more of a cu-

riosity) that the “poor” solution of the density-entropy formulation in the cases above looks very similar to the “correct” solution for a modestly magnetized medium (with some field component parallel to the shear or tangled; see e.g. Frank et al. 1996 and references therein). This occurs because, if we consider the linear perturbation analysis of the KH instability, the “incorrect” surface tension force here is (for the initial linear stage) almost mathematically identical to a “correct” magnetic tension term for parallel fields with strength $\beta \sim 1$.

4.4 The “Blob” Test

We next consider the “blob” test, which is designed to test processes of astrophysical interest (e.g. ram-pressure stripping, mixing, KH and Rayleigh-Taylor instabilities in a multiphase medium). Again our initial conditions come from the Wengen test suite and are described in Agertz et al. (2007). Briefly, the initial conditions include a spherical cloud of gas with uniform density in pressure equilibrium with an ambient medium, in a wind-tunnel with periodic boundary conditions. The imposed wind has Mach number $\mathcal{M} = 2.7$ and the initial density/temperature ratios are $= 10$. The box is a periodic rectangle with dimensions $x, y, z = 2000, 2000, 6000$ kpc, with the cloud centered on $0, 0, -2000$ kpc; the 9.6×10^6 particles are equal-mass and placed on a lattice.

Fig. 7 shows the resulting morphology of the cloud as a function of time, in the pressure-entropy formulation. Fig. 8 attempts to quantify the rate of cloud disruption following Agertz et al. (2007), who defined a useful standard criterion for measuring the degree of mixing of initially dense blob gas: at each time, we simply measure the total mass in gas with $\rho > 0.64\rho_c$ and $T < 0.9T_a$ (where ρ_c and T_a are the initial cloud density and ambient temperature). We show the standard SPH (density-entropy) result from that paper, as compared to the prediction from the pressure-entropy formulation here. We also compare the result from the identical test in SM12, using the pressure-energy formulation but without including the ∇h terms in the EOM, as well as the results from a high-resolution run with Enzo (O’Shea et al. 2004), a grid code.

The wind-cloud interaction generates a bow shock and immediately begins disrupting the cloud via a combination of KH and RT instabilities at its surface. By the middle panel in Fig. 7, there is no visible large concentration of dense material. Compare this to the identical initial conditions run with standard SPH and grid codes in Agertz et al. (2007) (their Figs. 4 & 7). In standard SPH, the cloud is compressed to a “pancake,” but the tension term prevents mixing at the surface and so a sizeable fraction survives disruption for large timescales. In contrast, the predicted morphology of the cloud here agrees very well with that in adaptive mesh codes therein (as well as moving-mesh methods in Sijacki et al. 2011). And quantitatively, we see this in Fig. 8.

However, once again we should note that the solution derived from the density-entropy formulation is remarkably similar to the “correct” solution with magnetic fields with $\beta \gtrsim 1$ (compare e.g. Mac Low et al. 1994; Shin et al. 2008), because the artificial surface tension term acts similarly to magnetic tension.

4.5 Sub-Sonic Turbulence

Bauer & Springel (2011) present a detailed study of the properties of idealized driven isothermal turbulence in SPH – specifically using GADGET-3 with the density-entropy formulation – as compared to moving mesh and grid methods (in the code AREPO; Springel 2010). Consistent with earlier results, they found that different methods agree well when turbulence is super-sonic (e.g. Kitsonas et al. 2009; Price & Federrath 2010). But when the turbu-




Figure 7. Central density slices in the “blob test” (a high-density, pressure-equilibrium cloud hit by a wind) with the pressure-entropy formulation. Time (in units of the KH growth time) increases from top to bottom. The high-density (red/orange) cloud gas is efficiently mixed by instabilities within a couple cloud crossing times; the morphology and density distribution agree well with grid codes.

lence is sub-sonic, they found that density-entropy SPH tended to reproduce a smaller inertial range. However as discussed there and in Price (2012a), this can depend quite sensitively on the artificial viscosity prescriptions and other numerical details. We therefore reproduce the sub-sonic, driven turbulence experiment from Bauer & Springel (2011) but adopt the pressure-entropy formulation, to test whether it also differs significantly from the results in the density-entropy formulation. We adopt the identical setup and resolution (A. Bauer, private communication), with the initial conditions and driving algorithm described therein. Briefly, we initialize a box of unit length, density, and sound speed with 128^3 particles, and drive the turbulence in a narrow range of large- k modes with characteristic Mach number $\mathcal{M} = 0.3$ on the largest scales; unlike our other experiments, the gas is isothermal ($\gamma = 1$). We run the experiment to $t = 25$ (more than sufficient to reach steady-state).

We compare the turbulent velocity spectrum measured following Bauer & Springel (2011) (we linearly interpolate the particle-centered velocity values to a uniform grid, fourier transform each




Figure 8. Fraction of the initial cloud in Fig. 7 which remains both cold and dense (i.e. avoids mixing) as a function of time, relative to the KH growth time at the cloud surface. We compare the standard SPH (density-entropy) and pressure-entropy formulations here, as well as the pressure-energy formulation in SM12 (which does not include the ∇h terms) and the results from a high-resolution grid code method (Enzo). The grid code and pressure formulations (independent of the ∇h terms) agree reasonably well. Density formulations show slower mixing/stripping.




Figure 9. Velocity power spectrum in sub-sonic ($\mathcal{M} = 0.3$) driven, isothermal turbulence. Each simulation uses 128^3 particles and identical driving. We compare the analytic Kolmogorov inertial-range model, to that calculated from our standard density-entropy (Eq. 14) and pressure-entropy (Eq. 21) formulations. The mean softening length h is shown as the vertical dotted line. Both do well down to a few softening lengths, where they first fall below the analytic result (excess dissipation) then rise above (kernel-scale noise). The EOM choice has a weak effect on the results. We compare different SPH algorithms and codes (description in text); agreement is good where the same methods are used. The “PHANTOM” and “GADGET-3 with Morris97 AV” use a density-entropy EOM with similar artificial viscosity treatment to our calculations (which dominates the inertial range), but different SPH kernels and power spectrum calculation methods (which dominate the noise at $\lesssim 3h$). The “standard” GADGET-3 calculation uses a constant artificial viscosity and different kernel, and produces almost no inertial range. The AREPO calculation considers a grid/moving mesh at the same resolution: deviations from Kolmogorov occur around the resolution limit, but are of a different character.

velocity component, and average in bins of fixed $|\mathbf{k}|$ to obtain the kinetic energy power spectrum). The resulting power spectrum is plotted from $k \approx 2\pi$ (though $k \lesssim 10$ are affected by the turbulent forcing as well as the finite box size) to $k \approx 500$ (a smoothing length h). We compare the pressure-entropy formulation to the results of standard (density-entropy formulation) SPH. The results are nearly identical (although there is slightly more power at intermediate scales in the density-entropy formulation). The results follow the expected Kolmogorov inertial range, until $\sim 4h$, where they drop below owing to artificial numerical dissipation; they then rise at $\sim 2h$, owing to kernel-scale noise (largely from pressure gradient errors).

We can compare this to other SPH results using the identical initial conditions and driving routine. First, consider the results from Bauer & Springel (2011), using “standard” GADGET-3 (density-entropy EOM, without the additional algorithm improvements discussed in § 3) but with an artificial viscosity scheme following Morris & Monaghan (1997) similar (though slightly different) to that adopted here, and the noisier $N_{\text{NGB}} = 32$ cubic spline kernel. We also compare the results from Price (2012a) using a different SPH code (PHANTOM), with the same density-entropy EOM, and an artificial viscosity treatment identical to that here, but using an intermediate kernel ($N_{\text{NGB}} = 58$ cubic spline) and different power spectrum calculation method. As shown there, the artificial viscosity treatment dominates large/intermediate scales. So with identical artificial viscosity, the Price (2012a) result is identical to ours over the inertial range (including the initial “drop off”). The smaller-scale ($\lesssim 2 - 3h$) behavior is dominated by the kernel choice and method of power spectrum calculation. If we compare “standard” GADGET-3 with constant artificial viscosity (Gingold & Monaghan 1983) from Bauer & Springel (2011), we see much more severe excess dissipation, and almost no inertial range.

Comparing to the results from AREPO in Bauer & Springel (2011) (run either in fixed-grid or moving-mesh mode, they are nearly identical), we see that the deviations from Kolmogorov there are opposite in sign and quite distinct. In SPH, artificial viscosity produces excess dissipation on larger (resolved) scales, while kernel gradient errors lead to particle noise that boosts the power at the smallest scales. The former (larger-scale deviations), being dominated by artificial viscosity, are only indirectly tied to resolution; considering higher-resolution runs we confirm the results in Bauer & Springel (2011) regarding the relatively slow convergence the high- k spectrum in SPH relative to grid codes. The latter (smaller-scale deviations) are related to the conservative nature of the code and its ability to dissipate particle noise (their “return” to and overshoot of the Kolmogorov power-law are numerical, rather than physical consequences of a turbulent cascade). In Eulerian approaches, the lack of any but numerical viscosity concentrates the dissipation at the resolution/cell scale, which produces a “bottleneck” of excess power that cannot be dissipated at the larger (marginally resolved) scales; but this scales directly with the resolution.

4.6 Stellar Feedback in Isolated Galaxies

Finally we test our modified algorithm on a fully non-linear system of particular astrophysical interest. Specifically, we re-run one of the star-forming galaxy disk models described in Hopkins et al. (2011b) and a subsequent series of papers (Hopkins et al. 2012, 2011c,a). These simulations include gravity, collisionless stars and




Figure 10. Morphology of the gas in a simulation of a star-forming disk galaxy (an isolated, SMC-mass dwarf) from Hopkins et al. (2012), following the resolved physics of the ISM and stellar feedback explicitly. Brightness encodes projected gas density (logarithmically scaled with ≈ 6 dex stretch); color encodes gas temperature with the blue material being $T \lesssim 1000$ K molecular and atomic gas, pink $\sim 10^4 - 10^5$ K warm ionized gas, and yellow $\gtrsim 10^6$ K hot gas. We show the galaxy face-on *top* and edge-on *bottom* after several orbital periods of evolution (during which the galaxy is in quasi-steady state). We compare the identical simulation using the density-entropy (Eq. 14) and pressure-entropy (Eq. 21) formulations of SPH. The two are largely similar. There is slightly sharper delineation between phases (e.g. molecular clouds and hot bubbles) in the density-entropy formulation; this includes the sharper transition between the star-forming and outer disks (physically, where the cooling time becomes longer than the dynamical time).

dark matter, and gas with a wide range of cooling processes¹⁰ from $\sim 10 - 10^9$ K, super-sonic turbulence, shocks, star formation in resolved giant molecular clouds, and explicit treatment of stellar feedback from SNe Types I & II, stellar winds, photoionization and photoelectric heating, and radiation pressure. Our purpose here is to test how robust these algorithms are (since many problems with e.g. non-linear error terms or conservation errors often do not manifest in simple test problems such as those above), and to examine how significant the pure numerical issues of fluid mixing are, relative to e.g. the previously-studied effects of adding/removing different physics in these simulations.

We specifically re-run the “SMC” model (a model of an Small Magellanic Cloud-mass, gas-rich isolated dwarf galaxy), at “high” resolution (1 pc softening length) as defined in the papers above, with all the above physical processes included exactly as in Hopkins et al. (2012), but in one case adopting the density-entropy formulation and in another the pressure-entropy formulation. We choose this model because, of the galaxy types studied therein, it has the largest mass fraction in hot gas and is most strongly affected by thermal feedback mechanisms (e.g. SNe), and as a result is most

¹⁰ Cooling requires a density estimate, independent of whether one is needed for the EOM. For this, we use the standard SPH kernel density estimator, for the reasons in § 3.4.




Figure 11. Star formation rate and stellar wind mass-loading (ratio of total unbound gas mass to total new stellar mass formed since the start of the simulation) as a function of time, for the galaxy simulations in Fig. 10. Although stochastic variations in the SFR differ by factors $\sim 2 - 3$, the time-averaged SFR is only $\sim 20\%$ lower in the density-entropy formulation. The wind mass-loading is systematically higher in this case by a factor $\sim 1.5 - 2.0$. This stems from increased phase mixing in the pressure-entropy formulation introducing more cold, dense gas into “hot” gas, enhancing its cooling before it can vent out of the disk.

sensitive to fluid mixing and phase structure. Figs. 10-11 compare the visual morphologies and star formation rates from these simulations. Taking into account that the systems are turbulent and non-linear, the morphologies are similar. They differ as one might expect: the pressure-entropy formulation increases mixing along phase boundaries, leading to less-sharp divisions between molecular clouds and/or hot, under-dense bubbles and the surrounding medium. This also makes the division between the star-forming and “extended” disk less sharp, as it is largely determined by the radius where the cooling rate becomes fast relative to the dynamical time.

Quantitatively, the time-averaged star formation rates differ remarkably little, just $\approx 20\%$ (with higher SFRs in the pressure-entropy case); this owes to the fact that they are globally set by equilibrium between momentum injection from stellar feedback and dissipation of turbulent energy/momentum on a dynamical time (Hopkins et al. 2012). The instantaneous rates vary more rapidly and can differ by factors of $\sim 2 - 3$. We also compare the mass-loading of galactic winds in each simulation, defined as in Hopkins et al. (2011c) as the ratio of total “wind mass” (material with positive bernoulli parameter, i.e. which will be unbound in the absence of external pressure forces) to new stellar mass formed since the beginning of the simulation. Here the qualitative behavior is the same but the wind mass-loading is systematically smaller in the pressure-entropy case, by a factor of $\sim 1.5 - 2$. This is directly related to the higher SFR; both owe to the increased mixing reducing the cooling time of the hot gas bubbles (which provide, for this small galaxy, much of the wind mass). In more massive galaxies, the winds in these simulations are less dominated by hot gas, so the effect will be smaller. These differences (while not negligible) are smaller than those typically caused by adding or removing different feedback mechanisms, or even changing details of their implementation (shown in Hopkins et al. 2011c). For example, removing stellar feedback entirely in this simulation changes the SFR by a factor of ~ 50 ! Integrated over cosmological times, or allowing for different galaxy conditions because of the integrated effects of different heating/cooling efficiencies as gas actually accretes into ha-

los, however, these small differences can easily build up into more significant divergence (see e.g. Vogelsberger et al. 2011).

5 DISCUSSION

With inspiration from SM12, we re-derive a fully self-consistent set of SPH equations of motion which are independent of the kernel-calculated density, and therefore remove the well-known “surface tension” terms that can suppress fluid mixing. The equations still depend on the medium having a differentiable pressure (hence still require some artificial viscosity to capture shocks), but unlike the traditional SPH EOM, remain valid through contact discontinuities.

Our derivation of the EOM relies on the key conceptual point in SM12, that the SPH volume element does not have to explicitly involve the mass density. However, our derivation and resulting EOM adds four key improvements: **(1)** We rigorously derive the equations from the discretized particle Lagrangian. This guarantees one of the most powerful features of SPH, namely *manifest* simultaneous conservation of energy, entropy, momentum, and angular momentum, and an exact solution to the particle continuity equation. **(2)** We similarly derive the “ ∇h ” terms, which are *required* for manifest conservation if the SPH smoothing lengths h_i are not everywhere constant. **(3)** We derive an “entropy formulation” of the equations that allows for the direct evolution of the entropy, avoiding the need to construct/evolve an energy equation, and gives better entropy conservation properties as in S02; this also happens to minimize the correction terms involved in using a “particle neighbor number” definition to define h , as compared to the “energy formulation.” **(4)** We show how the Lagrangian derivation can be generalized to *separate* definitions of the thermodynamic volume element (relating e.g. P and u) and that used to define the smoothing lengths. This resolves problems of numerical stability and excess diffusion in strong shocks and/or large density contrasts, and automatically allows for varying particle masses.

In fact, we derive a completely general, Lagrangian form of the EOM, including the ∇h terms, for *any* definition of the SPH thermodynamic volume element. Essentially *any* particle-carried quantity can be used in the kernel sum entering the EOM, and any (not necessarily the same) differentiable function used to define how the smoothing lengths h_i are scaled. In some ways this replaces the long-known “free weighting functions” used to define the SPH EOM in their original “discretized volume element” formulation. However, in that approach, the choice of different functions generically violates conservation and continuity; here, we demonstrate that a similar physical degree of freedom can be utilized in the discretization of the equations of motion without such violations.

Based on this degree of freedom, it is easy to see how different discretizations of the EOM might be optimized for some problems. By choosing the required kernel-evaluated element to directly represent a very smooth/stable property in the system, one not only removes spurious “tension” terms associated with discontinuities in other system variables, but also minimizes the inevitable discretization error from representing these quantities with a kernel sum. For the constant-pressure (but mixed density) fluid mixing tests we show here, the optimal choice is the “pressure formulation.” In MHD applications this kernel sum could trivially be altered to include the magnetic pressure. However if simulating an incompressible or weakly compressible fluid, the “density formulation” may well be superior. Direct kernel sums of nominally “higher order” properties such as the vorticity or vortensity are also valid and may represent useful formulations for some problems. It is even possible (in principle) to generalize our derivation to one in which

different particle subsets have differently-defined volume elements; although we caution that such an approach requires great care.

For the test problems here, we show that the “pressure-entropy” and “pressure-energy” formulations dramatically improve the treatment of fluid interface instabilities including the Kelvin-Helmholtz instability and the “blob test” (a mix of Kelvin-Helmholtz and Raleigh-Taylor instabilities as well including non-linear evolution and shock capturing); giving results very similar to grid methods. SM12 also demonstrate that a “pressure formulation” dramatically improves the behavior in Rayleigh-Taylor instabilities (see their Figs. 10-11), and that it removes the “deforming” effect of the surface tension term (allowing, for example, the long-term evolution of an irregular shape of gas at constant pressure but high density contrast; their Figs. 5-6); deformation is difficult to avoid even in grid codes (unless the chosen geometry matches the grid), and would otherwise require moving-mesh approaches to follow. However, unlike some of the modifications in the literature proposed to improve the fluid mixing in SPH (which violate conservation), the manifest conservation properties of our derivation mean that it remains well-behaved even in very strong shocks and does not encounter problems of either energy conservation or particle order in e.g. extremely strong blastwave problems.

With these changes in place, we find weaker (albeit still significant) residual effects from improvement in the artificial viscosity scheme. Comparisons of such schemes are well-studied and what we implement here is still not the most sophisticated possible treatment, although it still considerably reduces artificial viscosity away from shocks (for more detailed studies, see e.g. Cullen & Dehnen 2010). We find similar effects from changes to the SPH smoothing kernel. Our favored kernel is taken from more detailed kernel comparison studies in Hongbin & Xin (2005); Dehnen & Aly (2012); however, unlike some other SPH formulations, we find that even the “simplest” kernel possible ($N_{\text{NGB}} = 32$ cubic spline) reproduces good results in several tests, except in the expected regime where we wish to resolve kernel-scale growing instabilities that rely on sub-sonic motions at the level of $\mathcal{M} \lesssim N_{\text{NGB}}^{-1}$ and so summation errors dominate. This relates to the manifest conservation and maintenance of good particle order implicit in the EOM (see Price 2012b).

We test the algorithm not just in the “standard” set of test problems but also an example of direct astrophysical interest, simulating the evolution of galaxies with a multi-phase ISM. This is useful because it makes clear that for this problem, at least, the differences arising from the treatment of different physics (e.g. how cooling, star formation, stellar feedback, and AGN feedback are implemented) makes, on average, larger differences than the numerical scheme (also shown in other code comparisons; e.g. Scannapieco et al. 2012). This is not surprising: those choices lead to orders-of-magnitude differences as opposed to the (still significant) factor \sim couple effects of numerical choices. Moreover the differences we are concerned with here largely pertain to mixing in sub-sonic, non-radiative flows dominated by thermal pressure; in contrast many astrophysical problems of interest involve highly super-sonic, radiative, gravity-dominated flows. In that limit, the differences owing to the algorithm are often – though certainly not always – minimized (see references in § 1). But there are important regimes with transonic flows where the numerical approach can make larger differences (e.g. cosmological inflows & outflows; see Vogelsberger et al. 2011; Keres et al. 2011; Torrey et al. 2011). Even in idealized test problems, we caution that simple physical differences can produce larger distinctions than the numerical method. For example, for several fluid mixing problems considered here, the “correct” MHD solution in the presence of an equipartition magnetic field

can resemble the “standard” (density-entropy) SPH solution without a magnetic field, as opposed to the results from our pressure-entropy formulation or grid codes without such fields. The reason is that the real magnetic tension suppresses mixing, similar to the (purely numerical) “surface tension” term discussed in the text (so one might obtain a more “realistic” solution, but for entirely wrong reasons).

Ultimately, the numerical formulations derived here should provide the basis for a more rigorous approach to the “flavors” of SPH, and a means to compare the consequences of the fundamental choice of how to discretize any SPH approach. This change to the algorithm is not a panacea! Fortunately, the modified equations of motion proposed here can be trivially incorporated with many other methods that improve on other numerical aspects, for example the inviscid algorithm in Cullen & Dehnen (2010), the higher-order dissipation switches in Price (2008) and Read & Hayfield (2012), and/or the gradient error reducing integral formulation of the kernel equations in García-Senz et al. (2012). We wish to stress that – although SPH certainly has some disadvantages which we have not attempted to address here – poor fluid mixing in contact discontinuities is not necessarily an “inherent” property of SPH. This problem can be improved without requiring additional dissipation terms (and without additional computational expense) while retaining what is probably the greatest advantage of SPH algorithms, namely their excellent conservation properties.

ACKNOWLEDGMENTS

We thank Dusan Keres and Lars Hernquist for helpful discussions, and Oscar Agertz, Andrey Kravtsov, Robert Feldmann, and Nick Gnedin for contributions motivating this work. We also thank Justin Read, Volker Springel, and Eliot Quataert for useful comments, and thank Andreas Bauer for sharing the implementation of the turbulent driving routine. Support for PFH was provided by NASA through Einstein Postdoctoral Fellowship Award Number PF1-120083 issued by the Chandra X-ray Observatory Center, which is operated by the Smithsonian Astrophysical Observatory for and on behalf of the NASA under contract NAS8-03060.

REFERENCES

- Abel, T. 2011, MNRAS, 413, 271
- Agertz, O., et al. 2007, MNRAS, 380, 963
- Balsara, D. S. 1989, Ph.D. thesis, Univ. Illinois at Urbana-Champaign
- Bauer, A., & Springel, V. 2011, MNRAS, in press, arXiv:1109.4413
- Cullen, L., & Dehnen, W. 2010, MNRAS, 408, 669
- Dehnen, W., & Aly, H. 2012, MNRAS, in press arXiv:1204.2471
- Dilts, G. A. 1999, International Journal for Numerical Methods in Engineering, 44, 1115
- Durier, F., & Dalla Vecchia, C. 2012, MNRAS, 419, 465
- Frank, A., Jones, T. W., Ryu, D., & Gaalaas, J. B. 1996, ApJ, 460, 777
- Fulk, D. A., & Quinn, D. W. 1996, Journal of Computational Physics, 126, 165
- García-Senz, D., Cabezón, R. M., & Escartín, J. A. 2012, A&A, 538, A9
- Gingold, R. A., & Monaghan, J. J. 1977, MNRAS, 181, 375
- . 1983, MNRAS, 204, 715
- Hongbin, J., & Xin, D. 2005, Journal of Computational Physics, 202, 699
- Hopkins, P. F., Keres, D., Murray, N., Quataert, E., & Hernquist, L. 2011a, MNRAS, in press [arXiv:1111.6591]

- Hopkins, P. F., Quataert, E., & Murray, N. 2011b, MNRAS, 417, 950
- . 2011c, MNRAS, in press, arXiv:1110.4638 [astro-ph]
- . 2012, MNRAS, 421, 3488
- Keres, D., Vogelsberger, M., Sijacki, D., Springel, V., & Hernquist, L. 2011, MNRAS, in press, arXiv:1109.4638
- Kitsionas, S., et al. 2009, A&A, 508, 541
- Lucy, L. B. 1977, AJ, 82, 1013
- Mac Low, M.-M., McKee, C. F., Klein, R. I., Stone, J. M., & Norman, M. L. 1994, ApJ, 433, 757
- Marri, S., & White, S. D. M. 2003, MNRAS, 345, 561
- Monaghan, J. 1997, Journal of Computational Physics, 138, 801
- Morris, J. P. 1996, Publications of the Astronomical Society of Australia, 13, 97
- Morris, J. P., & Monaghan, J. J. 1997, Journal of Computational Physics, 136, 41
- Murante, G., Borgani, S., Brunino, R., & Cha, S.-H. 2011, MNRAS, in press arXiv:1105.1344
- Okamoto, T., Jenkins, A., Eke, V. R., Quilis, V., & Frenk, C. S. 2003, MNRAS, 345, 429
- O'Shea, B. W., Bryan, G., Bordner, J., Norman, M. L., Abel, T., Harkness, R., & Kritsuk, A. 2004, eprint, arXiv:astro-ph/0403044
- Price, D. J. 2008, Journal of Computational Physics, 227, 10040
- . 2012a, MNRAS, 420, L33
- . 2012b, Journal of Computational Physics, 231, 759
- Price, D. J., & Federrath, C. 2010, MNRAS, 406, 1659
- Read, J. I., & Hayfield, T. 2012, MNRAS, 422, 3037
- Read, J. I., Hayfield, T., & Agertz, O. 2010, MNRAS, 405, 1513
- Ritchie, B. W., & Thomas, P. A. 2001, MNRAS, 323, 743
- Rosswog, S. 2009, New Astronomy Reviews, 53, 78
- Saitoh, T. R., & Makino, J. 2009, ApJL, 697, L99
- . 2012, ApJ, in press, arXiv:1202.4277
- Scannapieco, C., et al. 2012, MNRAS, 2970
- Schoenberg, I. J. 1946, Q. Appl. Math, IV, 45
- Shin, M.-S., Stone, J. M., & Snyder, G. F. 2008, ApJ, 680, 336
- Sijacki, D., Vogelsberger, M., Keres, D., Springel, V., & Hernquist, L. 2011, MNRAS, in press, arXiv:1109.3468
- Springel, V. 2010, MNRAS, 401, 791
- Springel, V. 2010, ARA&A, 48, 391
- Springel, V., & Hernquist, L. 2002, MNRAS, 333, 649
- Torrey, P., Vogelsberger, M., Sijacki, D., Springel, V., & Hernquist, L. 2011, ApJ, in press, arXiv:1110.5635
- Vogelsberger, M., Sijacki, D., Keres, D., Springel, V., & Hernquist, L. 2011, MNRAS, in press arXiv:1109.1281
- Wadsley, J. W., Veeravalli, G., & Couchman, H. M. P. 2008, MNRAS, 387, 427