

The Kelvin–Helmholtz instability and smoothed particle hydrodynamics

Terrence S. Tricco  

Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St George Street, Toronto, ON M5S 3H8, Canada

Accepted 2019 July 7. Received 2019 June 23; in original form 2018 August 3

ABSTRACT

We perform simulations of the Kelvin–Helmholtz instability using smoothed particle hydrodynamics (SPH). The instability is studied both in the linear and strongly non-linear regimes. The smooth, well-posed initial conditions of Lecoanet et al. (2016) are used, along with an explicit Navier–Stokes viscosity and thermal conductivity to enforce the evolution in the non-linear regime. We demonstrate convergence to the reference solution using SPH. The evolution of the vortex structures and the degree of mixing, as measured by a passive scalar ‘colour’ field, match the reference solution. Tests with an initial density contrast produce the correct qualitative behaviour. The \mathcal{L}_2 error of the SPH calculations decreases as the resolution is increased. The primary source of error is numerical dissipation arising from artificial viscosity, and tests with reduced artificial viscosity have reduced \mathcal{L}_2 error. A high-order smoothing kernel is needed in order to resolve the initial velocity amplitude of the seeded mode and inhibit excitation of spurious modes. We find that standard SPH with an artificial viscosity has no difficulty in correctly modelling the Kelvin–Helmholtz instability and yields convergent solutions.

Key words: hydrodynamics – instabilities – methods: numerical.

1 INTRODUCTION

There exists doubt about the ability of smoothed particle hydrodynamics (SPH) to correctly model the Kelvin–Helmholtz instability and hydrodynamical mixing instabilities in general (e.g. Heß & Springel 2010; Springel 2010b; Bauer & Springel 2012; Sijacki et al. 2012; Hopkins 2013, 2015). This contention should be taken seriously, for if true, it would imply that SPH is unable to solve the fundamental equations of hydrodynamics. This disbelief is borne primarily from the work of Agertz et al. (2007), who, after comparing results between grid-based methods and SPH, concluded that SPH has a fundamental shortcoming in the mixing of multiphase fluids.

Wadsley, Veeravalli & Couchman (2008) and Price (2008) were quick to point out that the results of Agertz et al. (2007) were not due to a fundamental flaw of SPH, but are explainable by the combined effect of discontinuous initial conditions with the lack of numerical treatment for contact discontinuities. These are solvable issues.

The primary issue in Agertz et al. (2007)’s interpretation of their results arises from the presence of a contact discontinuity in the initial conditions. Discontinuities of any form lead to a multivalued fluid, and obtaining a solution numerically may lead to an unphysical solution if the fluid is assumed to be differentiable. What is required is either to make no assumption about differentiability (e.g. by solving the equations in integral form), or to include a numerical method to keep the fluid single valued. In the case of

shocks, for example, the fluid is kept single valued by the addition of a Riemann solver (for grid codes) or artificial viscosity (SPH). Contact discontinuities require similar treatment. For grid methods, some measure of numerical dissipation is inherently present due to the truncation error involved in advection of fluid between grid cells (e.g. Robertson et al. 2010). For SPH, there is no such corresponding numerical dissipation since the equations of motion for the particles are derived from the Lagrangian. Advection of entropy is exact (Springel & Hernquist 2002; Price 2008). What results is a spurious surface tension at contact discontinuities that keeps the two phases apart (the ‘pressure gap’; Okamoto et al. 2003; Agertz et al. 2007; Read, Hayfield & Agertz 2010; Springel 2010a; Abel 2011). In some sense, the problem is that SPH is too good, with the only dissipation present that which is explicitly added.

Solutions exist to handle contact discontinuities in SPH. One option is to explicitly add dissipation in the form of an artificial conductivity, similar to artificial viscosity, which keeps the fluid single valued by smoothing the contact discontinuity to the order of the resolution length (Price 2008; Wadsley et al. 2008). This has been shown to be effective in capturing contact discontinuities, removing the pressure gap, and promoting mixing (Valcke et al. 2010; Kawata et al. 2013). One challenge lies in the detection of contact discontinuities so that the artificial dissipation does not unnecessarily diffuse pressure gradients, for example, when pressure gradients are balanced by external forces (such as gravity). Options have been suggested for such cases (Valdarnini 2012; Price et al. 2017).

Other solutions have focused on changing the base SPH formulation. If a kernel summation is used for the internal energy

* E-mail: tricco@cita.utoronto.ca

or entropy, then the assumption of differentiability is not implied in the resulting equations. These methods, often referred to as ‘density-independent’ or ‘pressure–entropy’ formulations (Ritchie & Thomas 2001; Hopkins 2013; Saitoh & Makino 2013), perform well in promoting mixing, though no formal convergence study of the Kelvin–Helmholtz instability has been conducted. That they can be used for general astrophysical simulations without introducing other issues remains to be determined. Hopkins (2013) found that ‘pressure–entropy’ formulations do not provide as good an estimate of the mass density, recommending that a hybrid scheme be used for the density summation and entropy summation when the mass density is required (i.e. additional physics), and Read & Hayfield (2012) found that this formulation performs poorly when dealing with shocks. Alternative approaches to handle contact discontinuities include kernel correction techniques (e.g. García-Senz, Cabezón & Escartín 2012; Frontiere, Raskin & Owen 2017).

Beyond the context of SPH, it is a problem in general to use discontinuous initial conditions to study the Kelvin–Helmholtz instability. Discontinuities present in the initial conditions, such as in the velocity or density, prohibit formal convergence since the initial conditions are ill defined, remaining unresolved irrespective of the spatial resolution (Robertson et al. 2010; McNally, Lyra & Passy 2012b). Convergence cannot be achieved because increasing the resolution of the calculation only permits access to higher wavenumber modes. Even if a chosen mode is seeded initially, higher wavenumber modes can be excited by the discontinuity and become dominant since they are the fastest growing modes (Chandrasekhar 1961).

This has not prevented authors from using discontinuous initial conditions to assert the accuracy of their simulations. Springel (2010b) and Hopkins (2015), for example, argued that since their voronoi tessellation and meshless finite mass (MFM) methods produce more small-scale structure than grid-based methods, they are closer to a ‘true’ physical solution. This may be untrue. Since these small-scale structures are generated by high wavenumber modes seeded by the numerics, there is no guarantee that they represent physical structures (McNally et al. 2012b; Duffell & MacFadyen 2015). It is inappropriate to attribute structure generated from numerical noise as corresponding to resolved physical structure.

Smooth, well-defined initial conditions of the Kelvin–Helmholtz instability produce converged solutions, avoiding the issues inherent to discontinuous initial conditions. Robertson et al. (2010), McNally et al. (2012b), and Lecoanet et al. (2016) have put forward well-defined initial conditions. They each demonstrate convergence in the linear regime, and, for Lecoanet et al. (2016), additionally in the non-linear regime by the inclusion of a physical Navier–Stokes viscosity and thermal conductivity. Convergence is demonstrated by these authors for a variety of codes and methods. Collectively, they use the Eulerian grid-based codes ART (Kravtsov, Klypin & Khokhlov 1997; Kravtsov, Klypin & Hoffman 2002), ENZO (O’Shea et al. 2005; Bryan et al. 2014), ATHENA (Stone et al. 2008), the PENCIL CODE (Brandenburg & Dobler 2002), the pseudo-spectral code DEDALUS (Burns et al. 2016, 2019), and the Lagrangian mesh-free code PHURBAS (Maron, McNally & Mac Low 2012; McNally, Maron & Mac Low 2012a). Convergence of the McNally et al. (2012b) test has also been demonstrated for the implicit large eddy simulation code MUSIC (Goffrey et al. 2017).

What is missing is a comprehensive study demonstrating convergence with SPH. A number of studies have been conducted using discontinuous initial conditions (Price 2008; Cha, Inutsuka & Nayakshin 2010; Junk et al. 2010; Valcke et al. 2010; Shadloo & Yildiz 2011; Hubber, Falle & Goodwin 2013). Studies with smooth

initial conditions have been limited. While McNally et al. (2012b) included results from the SPH code NDSPMHD (Price 2012) in their work, their analyses were not focused on SPH, with a total of five codes compared simultaneously. Other studies have either been limited in scope (Price et al. 2017), for incompressible fluids (Yue et al. 2015), or have used higher order corrections (Cabezón, García-Senz & Figueira 2017; Frontiere et al. 2017).

In this paper, a convergence study of the Kelvin–Helmholtz instability in SPH is undertaken using the smooth, well-defined initial conditions of Lecoanet et al. (2016). Importantly, a physical Navier–Stokes viscosity and thermal conductivity are included to set the dissipation. The aim is to demonstrate convergence in both the linear and non-linear regimes.

The standard formulation of SPH is used, rather than alternative ‘pressure–entropy’ formulations or higher order corrections. This is motivated by the many public astrophysical SPH codes implemented using standard formulations of SPH with artificial dissipation terms, e.g. GASOLINE2 (Wadsley, Keller & Quinn 2017), GANDALF (Hubber, Rosotti & Booth 2018), and PHANTOM (Price et al. 2017).

The paper is organized as follows. The set-up, initial conditions and numerical method used for the Kelvin–Helmholtz test, is described Section 2. Results are presented and analysed in Section 3. A summary and conclusion are given in Section 4.

2 PROBLEM SET-UP

2.1 Equations

We solve the Navier–Stokes equations of hydrodynamics that include a thermal conductivity, given by

$$\frac{d\rho}{dt} = -\rho \nabla \cdot \mathbf{v}, \quad (1)$$

$$\frac{d\mathbf{v}}{dt} = -\frac{\nabla P}{\rho} - \frac{1}{\rho} \nabla \cdot \mathbf{\Pi}, \quad (2)$$

$$\frac{du}{dt} = -\frac{P}{\rho} \nabla \cdot \mathbf{v} + \nabla \cdot (\mathbf{v} \cdot \mathbf{\Pi}) + \frac{1}{\rho} \nabla \cdot (\chi \rho \nabla T), \quad (3)$$

where ρ is the density, \mathbf{v} is the velocity, P is the thermal pressure, u is the internal energy, and T is the temperature. The equations are written in Lagrangian form, using the material derivative $d/dt \equiv \partial/\partial t + \mathbf{v} \cdot \nabla$. The Navier–Stokes stress tensor is

$$\Pi^{ij} = \nu \rho \left(\frac{\partial v^i}{\partial x^j} + \frac{\partial v^j}{\partial x^i} - \frac{2}{3} \frac{\partial v^k}{\partial x^k} \delta^{ij} \right), \quad (4)$$

where ν is the shear viscosity. Heating from the viscosity is included as the second term in equation (3). The third term in equation (3) is a physical thermal conductivity, with χ the thermal diffusivity. An ideal gas equation of state is used, $P = \rho T$, with the ratio of specific heats $\gamma = 5/3$. The temperature is thus related to the internal energy according to $T = (\gamma - 1)u$.

A scalar field is added to the simulation to quantify mixing. This colour field, or dye, is evolved according to

$$\frac{dc}{dt} = \frac{1}{\rho} \nabla \cdot (\nu_c \rho \nabla c). \quad (5)$$

The colour field is passively advected with the flow and does not influence the dynamics. A ‘physical’ diffusion term is present to set the diffusion rate. This is important to achieve convergence in grid codes since numerical diffusion is introduced due to advection. For

SPH, advection is dissipationless, but the explicitly added diffusion is retained for these calculations in order to converge to the results of Lecoanet et al. (2016).

2.2 Initial conditions

The initial conditions are given by

$$\rho = 1 + \frac{\Delta\rho}{\rho_0} \left[\tanh\left(\frac{y-y_1}{a}\right) - \tanh\left(\frac{y-y_2}{a}\right) \right], \quad (6)$$

$$v_x = v_0 \left[\tanh\left(\frac{y-y_1}{a}\right) - \tanh\left(\frac{y-y_2}{a}\right) - 1 \right], \quad (7)$$

$$v_y = A \sin(2\pi x) \left[\exp\left(-\frac{(y-y_1)^2}{\sigma^2}\right) + \exp\left(-\frac{(y-y_2)^2}{\sigma^2}\right) \right], \quad (8)$$

$$P = 10, \quad (9)$$

$$c = \frac{1}{2} \left[\tanh\left(\frac{y-y_1}{a}\right) - \tanh\left(\frac{y-y_2}{a}\right) + 2 \right], \quad (10)$$

where $a = 0.05$ and $\sigma = 0.2$ dictate the width of the interface regions. The simulations are run in a periodic box of size $x \in [0, L]$ and $y \in [0, 2L]$ with $L = 1$. The interfaces are at $y_1 = 0.5$ and $y_2 = 1.5$. For the chosen initial conditions, the top half mirrors the bottom half.

We take $v_0 = 1$, with x -velocity that smoothly transitions between the two regions. The flow remains subsonic with Mach number $\mathcal{M} \approx 0.25$. The amplitude of the v_y velocity perturbation used to seed the instability is $A = 0.01$. These initial conditions are well resolved even at modest resolutions.

The Reynolds number is $\text{Re}=10^5$, where the Reynolds number is defined according to

$$\text{Re} = \frac{L\Delta v}{\nu}. \quad (11)$$

The box size is $L = 1$ and the difference in fluid velocity is $\Delta v = 2v_0$. This yields $\nu = 2 \times 10^{-5}$, with $\chi = \nu_c = \nu$.

Lecoanet et al. (2016) considered cases where $\Delta\rho/\rho_0 = 0$ and 1. That is, for uniform density with $\rho = 1$ throughout, and with a density contrast between densities of $\rho = 1$ and 2. The stratified test introduces secondary instabilities that produce complex, small-scale structure in the flow. This a far more computationally demanding test case, with Lecoanet et al. (2016) only achieving convergence with the grid-code ATHENA in the non-linear regime when $16\,384 \times 32\,768$ grid cells are used, requiring ~ 1 million cpu hours. Since we do not have 1 million spare cpu hours available, our tests focus on the unstratified, uniform density test case. This is sufficient to quantify the convergence properties of the Kelvin–Helmholtz instability in SPH. Tests including a density contrast are studied in Section 3.9.

2.3 SPH

The equations are solved using SPH. The continuum equations are discretized into a set of particles that are smoothed over their local volume to recover continuum-like behaviour. The implementation of SPH used is similar to that of the PHANTOM SPH code (Price et al.

2017). No modifications have been made to the SPH formalism in order to simulate the Kelvin–Helmholtz instability. The calculations use the standard form of SPH with an artificial viscosity, of the type that has been employed for decades (Monaghan 1992, 2005). The discrete set of equations for Euler fluid flow, equivalent to equations (1)–(3) absent the Navier–Stokes viscous terms and thermal conductivity, is given by

$$\rho_a = \sum_b m_b W_{ab}(h_a), \quad (12)$$

$$\frac{d\mathbf{v}_a}{dt} = - \sum_b m_b \left[\frac{P_a}{\Omega_a \rho_a^2} \nabla_a W_{ab}(h_a) + \frac{P_b}{\Omega_b \rho_b^2} \nabla_a W_{ab}(h_b) \right] + \sum_b \frac{m_b}{\bar{\rho}_{ab}} v_{ab}^{\text{sig}} \mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab} \nabla_a \bar{W}_{ab}, \quad (13)$$

$$\frac{du_a}{dt} = \frac{P_a}{\Omega_a \rho_a^2} \sum_b m_b \mathbf{v}_{ab} \cdot \nabla_a W_{ab}(h_a) - \sum_b \frac{m_b}{\bar{\rho}_{ab}} v_{ab}^{\text{sig}} (\mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab})^2 \hat{\mathbf{r}}_{ab} \cdot \nabla_a \bar{W}_{ab}. \quad (14)$$

Here, a and b refer to particle indices and the summations are over neighbouring particles. The shorthand notation, $\mathbf{v}_{ab} \equiv \mathbf{v}_a - \mathbf{v}_b$, is used, along with $\hat{\mathbf{r}}_{ab} \equiv (\mathbf{r}_a - \mathbf{r}_b)/|\mathbf{r}_a - \mathbf{r}_b|$, which is the unit vector along the line of the sight between particles. The smoothing kernel is $W_{ab}(h_a) \equiv W(|\mathbf{r}_a - \mathbf{r}_b|, h_a)$, with $\bar{W}_{ab} \equiv \frac{1}{2}[W_{ab}(h_a)/\Omega_a + W_{ab}(h_b)/\Omega_b]$ the averaged smoothing kernel. The particle mass is m , and these sets of calculations use particles of equal mass.

The smoothing length, h , is obtained through iteration with the density, ρ , according to the relation

$$h_a = \eta \left(\frac{m_a}{\rho_a} \right)^p, \quad (15)$$

where p is the number of dimensions (two for these calculations). We use $\eta = 1.2$ for the ratio of smoothing length to interparticle spacing, which is ideal for the bell-shaped spline kernels (Dehnen & Aly 2012). Gradients of the smoothing length are accounted for by the factor

$$\Omega_a = 1 + \frac{h_a}{p\rho_a} \sum_b m_b \frac{\partial W_{ab}(h_a)}{\partial h_a}, \quad (16)$$

which is important to ensure exact energy conservation in the spatial discretization (Monaghan 2002; Springel & Hernquist 2002). The discretized momentum equation exactly conserves linear and angular momentum.

Artificial viscosity is included by the second terms of the momentum and energy equations (equations 13 and 14). This implementation closely matches that of Monaghan (1997). The signal velocity used is

$$v_{ab}^{\text{sig}} = \frac{1}{2} \bar{\alpha}_{ab} (c_{s,a} + c_{s,b}) - \beta \mathbf{v}_{ab} \cdot \hat{\mathbf{r}}_{ab}, \quad (17)$$

where c_s is the sound speed, α and β are dimensionless quantities, and $\bar{\alpha}_{ab} = \frac{1}{2}(\alpha_a + \alpha_b)$. The Morris & Monaghan (1997) switch is used to reduce numerical dissipation. Each particle has an individual α , which is time integrated according to

$$\frac{d\alpha_a}{dt} = \max(\nabla \cdot \mathbf{v}_a, 0) - \frac{\alpha_a - \alpha_0}{\tau_a}, \quad (18)$$

such that α is increased in regions of compression. The range of values is enforced to be $\alpha \in [\alpha_0, 1]$, with $\alpha_0 = 0.1$ the minimum

value used in these calculations. The time-scale over which α decays is $\tau_\alpha = 0.1h_a / \max(v_{ab}^{\text{sig}})$, where $\max(v_{ab}^{\text{sig}})$ is the maximum signal velocity as calculated over all neighbouring particles. This decay time-scale roughly corresponds to the time it takes for a sound wave to cross five smoothing lengths (Morris & Monaghan 1997). The β term represents a von Neumann–Richtmyer-type viscosity and is used to prevent particle interpenetration (Monaghan 1989). In following of the PHANTOM and GASOLINE2 codes, this term is not multiplied by α . In general, the kinetic energy dissipated by the β term tends to be lower than the α term since it scales as αh^2 .

The one important ingredient required to achieve convergence in these simulations is a high-order smoothing kernel. The amplitude of the initial velocity perturbation for the seeded Kelvin–Helmholtz mode has Mach number of the order of $\mathcal{M} \approx 0.0025$, notably smaller than in the initial conditions used by Springel (2010b), Robertson et al. (2010), McNally et al. (2012b), and Hopkins (2015). It is required that velocity noise caused by errors in the pressure gradient be below this amplitude in order to resolve the linear growth phase. Additionally, the excitation of other non-seeded modes by velocity noise must be inhibited in order to correctly model the non-linear phase. We found that the cubic spline is unable to do this for this problem, and that while the quintic spline produces all of the correct features of the reference solution, the kernel bias prevents convergence. The calculations presented here use the septic spline, which exhibits convergence, as will be demonstrated. Though the septic spline has a smoothing volume greater than the cubic or quintic splines, thus includes more neighbour particles, the resolution length of the calculation is the same. The ratio of smoothing length to interparticle spacing, as obtained via equation (15), is unchanged. The bias of the smoothing kernel on our calculations is discussed further in Section 3.7, and a list of the spline kernels, up to the nonic spline, is given in Appendix A.

2.4 Physical dissipation in SPH

Navier–Stokes viscosity is implemented using two first derivatives, in a manner similar to Flebbe et al. (1994), Watkins et al. (1996), and Sijacki & Springel (2006). Each component of the stress tensor (equation 4) is calculated using a difference derivative estimate (Price 2012) according to

$$\frac{\partial v_a^i}{\partial x_a^j} = -\frac{1}{\Omega_a \rho_a} \sum_b m_b (v_a^i - v_b^i) \frac{\partial W_{ab}(h_a)}{\partial x_a^j}. \quad (19)$$

The corresponding term in the momentum equation can be derived from the Lagrangian to be

$$\frac{dv_a^i}{dt} = -\sum_b m_b \left[\frac{\Pi_a^{ij}}{\Omega_a \rho_a^2} \nabla_a^j W_{ab}(h_a) + \frac{\Pi_b^{ij}}{\Omega_b \rho_b^2} \nabla_b^j W_{ab}(h_b) \right]. \quad (20)$$

The derivative of the stress tensor is a symmetric derivative operator. The conjugate nature of difference and symmetric derivative operators has been noted in other contexts (Cummins & Rudman 1999; Price 2010; Tricco & Price 2012).

Heating from the Navier–Stokes viscosity is given by

$$\frac{du_a}{dt} = \nu \rho_a \left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} \right)^2 - \frac{2}{3} \nu \rho_a \left(\frac{\partial v_a^k}{\partial x_a^k} \right)^2, \quad (21)$$

where by the square of the first term we mean the tensor summation $\sum_i \sum_j T_{ij} T_{ij}$ with $T_{ij} = \partial v_a^i / \partial x_a^j + \partial v_a^j / \partial x_a^i$. The heating term is positive definite provided the terms of the stress tensor are calculated using equation (19). The derivation of the viscous heating term is

given in Appendix B. This spatial discretization of the Navier–Stokes viscosity exactly conserves energy and linear momentum.

Thermal conductivity is implemented in SPH according to Cleary & Monaghan (1999), given by

$$\frac{du_a}{dt} = \chi \sum_b m_b (\rho_a + \rho_b) (u_a - u_b) \frac{\mathbf{r}_{ab} \cdot \nabla_a \overline{W}_{ab}(h_a)}{r_{ab}^2}. \quad (22)$$

This uses a direct second derivative in the manner of Brookshaw (1985). Cleary & Monaghan (1999) proposed that the harmonic mean of the thermal diffusivity coefficients, $4\chi \rho_a \rho_b / (\rho_a + \rho_b)$, is better for discontinuous thermal conductivities. For these calculations, the arithmetic mean, $\frac{1}{2}\chi(\rho_a + \rho_b)$, is used instead since the thermal conductivity for this problem varies smoothly.

Colour diffusion is implemented in an analogous fashion to the thermal conductivity, given by

$$\frac{dc_a}{dt} = \nu_c \sum_b m_b (\rho_a + \rho_b) (c_a - c_b) \frac{\mathbf{r}_{ab} \cdot \nabla_a \overline{W}_{ab}}{r_{ab}^2}. \quad (23)$$

This formulation exactly conserves the volume integral of colour, that is $\sum_a m_a c_a$, since it is equal and opposite for each particle pair.

2.5 Calculations

Calculations are performed using 256×592 , 512×1184 , 1024×2364 , and 2048×4728 particles in the x - and y -directions, respectively, labelled as $n_x = 256, 512, 1024$, and 2048 . The particles are arranged on triangular lattices, with the largest calculation consisting of ~ 10 million particles ($n_x = 2048$), requiring $\sim 70\,000$ cpu hours of computational time.

2.6 Reference solution

There is no analytic solution to the Lecoanet et al. (2016) Kelvin–Helmholtz test. In its absence, we turn to a solution obtained using high-resolution calculations, which acts as an approximation to the true solution. The SPH calculations performed in this work are compared to the data from the D2048 calculation of Lecoanet et al. (2016). This solution is obtained using the pseudo-spectral code DEDALUS (Burns et al. 2016, 2019),¹ on a grid of 2048×4096 points. The data consist of the density and colour fields at times of $t = 2, 4, 6$, and 8 .

3 RESULTS

The SPH calculations are qualitatively and quantitatively compared with the solution of the D2048 DEDALUS calculation from Lecoanet et al. (2016). Our focus is to study the convergence properties of SPH with respect to this solution, particularly in the strongly non-linear regime.

3.1 Linear regime; $t < 2$

Fig. 1 shows the colour field at $t = 2$ for the SPH calculations at resolutions of $n_x = 512, 1024$, and 2048 particles, along with the D2048 reference solution (far right-hand panel). Only the bottom half ($y < 1$) of the calculations are shown since the top and bottom mirror each other. At this time, the calculation is at the end of the linear evolution phase of the Kelvin–Helmholtz instability.

¹<http://dedalus-project.org>




Figure 1. The colour field in the region $x, y \in [0, 1]$ for the $n_x = 512, 1024,$ and 2048 particle SPH calculations (three leftmost panels) with the D2048 DEDALUS calculation (far right-hand panel) at $t = 2$. In all cases, the SPH calculations have grown the seeded mode to the correct shape. For the $n_x = 512$ particle calculation, there is some noise in the inner region of the curl arising from particle motion. This noise is absent from the $n_x = 2048$ particle calculation, which reproduces the D2048 reference solution in all aspects.

In all cases, the SPH calculations form a single large curl corresponding to the seeded $k = 2\pi$ mode, in agreement with the reference solution. For the $n_x = 512$ particle calculation, the inner tip of the curl is not wound as tightly as the reference solution. There is some additional noise in the inner region, due to particles breaking off lattice. These differences disappear for the $n_x = 2048$ particle calculation. Indeed, the visual appearance of the curl for the highest resolution SPH calculation ($n_x = 2048$) closely reproduces the reference solution.

The initial velocity perturbation grows exponentially in the linear regime of the Kelvin–Helmholtz instability, occurring at a rate $\propto \exp(\gamma t)$. The linear growth rate for an incompressible, infinite domain fluid is given by (Chandrasekhar 1961)

$$\gamma = \frac{(\rho_1 \rho_2)^{1/2}}{\rho_1 + \rho_2} |v_1 - v_2| k. \quad (24)$$

For the initial conditions in Section 2.2, the corresponding quantities are $k = 2\pi$, $\rho_1 = \rho_2 = 1$, $v_1 = 1$, and $v_2 = -1$. This yields a growth rate of $\gamma = 2\pi$.

However, the growth rate in these calculations is expected to differ since the interfaces are smoothed. Wang, Ye & Li (2010) derive analytic estimates of the growth rate for incompressible fluids with smoothed density and/or velocity interfaces. Their key conclusion is that smoothing the velocity interface slows the growth rate. Using their analytic estimate (cf. their equation 18) with the initial conditions employed here yields $\gamma \approx 5.09$, about a 20 per cent reduction from the estimate using equation (24).

Our calculations additionally smooth the velocity perturbation, and are for a compressible fluid in a periodic domain with a peak density fluctuation of ≈ 10 per cent. Therefore, it is expected that the growth rate will be further reduced. It is likely that an estimate using the approach of Wang et al. (2010) merely represents an upper bound.

Berlok & Pfrommer (2019) derive the linear theory for the Kelvin–Helmholtz instability with smooth initial velocity and density profiles (as well as the magnetized Kelvin–Helmholtz instability). Using their code, PSECAS, a numerical solution was obtained for the linear growth rate for the initial conditions of Lecoanet et al. (2016), with the growth rate of the $k = 2\pi$ mode found to be $\gamma \approx 3.227$.

To measure the growth rate of the $k = 2\pi$ mode in our calculations, we use a calculation based on a discrete convolution,

as described by McNally et al. (2012b). A weighting term for each particle, i , is defined according to

$$d_i = \begin{cases} h_i^2 \exp(-|y_i - \frac{1}{2}|/\sigma^2) & y < 1, \\ h_i^2 \exp(-|(2 - y_i) - \frac{1}{2}|/\sigma^2) & y \geq 1, \end{cases} \quad (25)$$

where h is the smoothing length and the weighting has characteristic size of the v_y perturbation as given in the initial conditions (equation 8). The wave components for the $k = 2\pi$ mode are given by

$$s_i = v_{y,i} \sin(2\pi x_i), \quad (26)$$

$$c_i = v_{y,i} \cos(2\pi x_i). \quad (27)$$

The mode amplitude, M , may then be computed by the weighted summations of the components of the mode according to

$$M = 2 \left[\left(\frac{\sum_{i=1}^N s_i d_i}{\sum_{i=1}^N d_i} \right)^2 + \left(\frac{\sum_{i=1}^N c_i d_i}{\sum_{i=1}^N d_i} \right)^2 \right]^{1/2}, \quad (28)$$

where the summation of d_i in the denominator acts as a normalization condition. The summations are over all particles. This approach avoids the need to interpolate to a fixed grid since the v_y values are taken directly from the SPH particles.

Fig. 2 shows the amplitude of the $k = 2\pi$ mode in the linear regime. For all calculations, it undergoes exponential amplification of the y -velocity, as expected for the linear regime of the Kelvin–Helmholtz instability. The mode amplitude reaches a maximum at $t \approx 2$, establishing that the non-linear regime begins at this time.

The measured growth rates per resolution are summarized in Table 1. The $n_x = 2048$ calculation has a growth rate $\gamma \approx 3.139$, which is in close agreement to the expected value of $\gamma \approx 3.227$ obtained using the PSECAS code (Berlok & Pfrommer 2019), with the growth rate converging linearly in resolution to the expected value.

3.2 Non-linear regime; $t > 2$

Fig. 3 shows the colour field at $t = 4, 6,$ and 8 for the $n_x = 512, 1024,$ and 2048 particle SPH calculations with comparison to the D2048 DEDALUS calculation. Only the bottom half ($y < 1$) of the calculation is shown since the top and bottom mirror each other.




Figure 2. Amplitude of the $k = 2\pi$ seeded mode for the $n_x = 256, 512, 1024,$ and 2048 SPH calculations. The growth is exponential in the linear regime, as expected for the Kelvin–Helmholtz instability. The $n_x = 2048$ calculation has a growth rate $\Gamma \propto \exp(3.139t)$.

Table 1. Measured growth rates, Γ , of the linear phase.

n_x	Γ
256	2.92
512	3.055
1024	3.113
2048	3.139

At these times, the Kelvin–Helmholtz instability is in the strongly non-linear regime. The evolution of the instability in the non-linear regime is dictated by the dissipation – both numerical and physical. The curl produced in the linear regime ($t = 2$; Fig. 1) continues to wind tighter as time progresses ($t = 4$), producing multiple features as the curl turns into itself ($t > 6$). At $t = 8$, the physical dissipation applied to the colour field has blended many of the smaller features together, producing a well-mixed colour field along the interface.

The SPH calculations qualitatively reproduce the evolution and morphology of the reference solution. For the $n_x = 512$ particle calculation, there is noise in the colour field due to particle motion, evident by the ‘jaggedness’ of the filaments comprising the curl. The origin of the noise is attributed to the high Reynolds number ($\text{Re} = 10^5$) of the calculations, which is challenging to meet at low resolutions. This noise is absent in the $n_x = 2048$ particle calculation.

Differences between the solutions, as seen in multiple local features, are reducing as the resolution improves. For all time snapshots, the innermost tip of the red curl is more tightly wound with increasing resolution. At $t = 4$, the outer features of the curl become more sharply defined, and, at $t = 8$, the top and bottom spurs are lengthening with increasing resolution. In all features, it is clear that the SPH calculations are converging to the reference solution.

3.3 Colour entropy

The volume-integrated colour entropy is measured to quantify the mixing of the two regions. Following Lecoanet et al. (2016), the specific colour entropy is defined to be

$$s \equiv -c \ln(c), \quad (29)$$

with the total colour entropy given by the volume integral

$$S = \int \rho s dV. \quad (30)$$

In SPH, this integral can be computed by summation according to

$$S = \sum_a m_a s_a, \quad (31)$$

such that ρdV is equivalent to the mass element dm . This quantity has the property that it monotonically increases in time when $v_c > 0$.

Fig. 4 shows the total colour entropy as a function of time for the SPH calculations. The black dots at $t = 2, 4, 6,$ and 8 are the values from the reference solution. The total colour entropy of the SPH calculations reproduces the behaviour of the reference solution (compare also to fig. 3 in Lecoanet et al. 2016). In the non-linear phase ($t > 2$), the difference in total colour entropy between the SPH calculations and the reference data points decreases as the resolution increases. That is, the SPH calculations are converging to the reference solution. At late times, the lower resolution SPH calculations have lower total colour entropy than the reference solution, in contrast to $\text{Re} = 10^6$ ATHENA results by Lecoanet et al. (2016) that had increased colour entropy for lower resolutions. This difference in behaviour is caused by advection in grid-based codes introducing numerical diffusion, thereby artificially increasing the rate of mixing. Advection is exact in SPH.

Deviations of the total colour entropy from the reference solution are minor even for our lowest resolution of $n_x = 256$ particles, and the shape and magnitude of the total colour entropy curve are well represented. From a practical perspective, SPH, even at modest resolutions, not only reproduces the correct qualitative features of the Kelvin–Helmholtz instability (cf. Fig. 3), but also the correct amount of mixing.

3.4 \mathcal{L}_2 error

The \mathcal{L}_2 error is computed to quantify the convergence of the SPH calculations towards the D2048 calculation. To compute the error between the two calculations, the SPH particles are interpolated to a 2048×4096 grid so that the two solutions can be directly compared. The colour at each grid point is computed according to

$$c(\mathbf{r}) = \frac{\sum_b m_b c_b W_{ab}(h_b)}{\sum_b m_b W_{ab}(h_b)}, \quad (32)$$

where the summation in the denominator is the normalization condition. The \mathcal{L}_2 error between the SPH interpolated grid and D2048 grid is computed as

$$\mathcal{L}_2 = \left[dV \sum_a |c_a^{\text{SPH}} - c_a^{\text{D2048}}|^2 \right]^{1/2}, \quad (33)$$

where dV is the volume of each grid cell (i.e. $1/2048^2$), and c_a^{SPH} and c_a^{D2048} are the colour fields of the SPH interpolated grid and the D2048 grid, respectively.

Fig. 5 shows the \mathcal{L}_2 error between the SPH calculations and the D2048 DEDALUS calculation. The error is computed at $t = 2, 4, 6,$ and 8 , with the lines between points present only for visual aid. The convergence rates at each time are listed in Table 2. The convergence rate at $t = 2$ is close to linear, with each doubling of resolution reducing the \mathcal{L}_2 error by half. The solution at $t = 4$ is converging towards the reference solution, though at a rate slower than linear. In the late non-linear regime of the instability ($t = 6$ and




Figure 3. The colour field in the region $x, y \in [0, 1]$ in the strongly non-linear regime for the $n_x = 512$, 1024, and 2048 SPH calculations (three leftmost columns) with the D2048 DEDALUS calculation (rightmost column) at $t = 4, 6$, and 8 (top, middle, and bottom rows, respectively). The SPH calculations reproduce the large-scale morphology of the DEDALUS calculation at all times. Differences in small-scale features are seen to be converging away as the resolution increases.




Figure 4. Total colour entropy for the $n_x = 256, 512, 1024$, and 2048 particle SPH calculations, with the black dots the values from the D2048 solution. The shape and magnitude of the total colour entropy curve agree with the reference data points for all resolutions.

8), the SPH solution only begins to demonstrate convergence for resolutions above $n_x = 512$ particles. The error between the $n_x = 256$ and 512 is similar at these times, a consequence of the artificial viscosity still dominant over the Navier–Stokes viscosity. The long-term behaviour, for the resolutions presented, is governed still by the numerical dissipation. Only when the physical dissipation begins to be resolved, that is in the $n_x = 1024$ and 2048 particle calculations, does the \mathcal{L}_2 error decrease with resolution, though at a rate slower than linear. The effect of numerical dissipation is investigated in detail in Section 3.5.

Overall, the SPH calculations are converging towards the D2048 reference solution. At all times, the \mathcal{L}_2 errors of the $n_x = 2048$ particle calculation are lower than for the $n_x = 1024$ particle calculation, which are lower than the $n_x = 512$ particle calculation.

3.5 Convergence of the dissipation rate

One question is whether the dissipation of kinetic energy is dominated by the physical Navier–Stokes viscosity or the numerical




Figure 5. \mathcal{L}_2 error of the SPH calculations with respect to the D2048 solution at $t = 2, 4, 6,$ and 8 . The connecting lines are merely for visual aid. The convergence rate is linear at $t = 2$. In the late non-linear regime ($t > 6$), the error only decreases once the Navier–Stokes viscosity begins to be resolved. The SPH calculations are converging to the reference solution at all times, with the \mathcal{L}_2 error of the $n_x = 2048$ particle calculation lower than the $n_x = 1024$ particle calculation, which is lower than the $n_x = 512$ particle calculation.

Table 2. Convergence rates of the \mathcal{L}_2 error fit to $\propto n_x^{-\sigma}$.

t	σ (artificial viscosity switch)	σ ($\alpha=0$)
2.0	0.875	1.108
4.0	0.623	0.754
6.0	0.257	0.381
8.0	0.278	0.376

dissipation of the artificial viscosity. This may be estimated by simple analytic arguments and through measurements directly from the calculations.

As discussed by Artymowicz & Lubow (1994), Murray (1996), and Lodato & Price (2010), the α term of the artificial viscosity can be equated to a shear and bulk Navier–Stokes viscosity, with shear component

$$v^{\text{AV}} = \frac{1}{10} \alpha v_{\text{sig}} h, \quad (34)$$

and bulk component that is $\zeta^{\text{AV}} = 5/3 \times v^{\text{AV}}$. Since the fluid is subsonic, the Morris & Monaghan (1997) artificial viscosity switch is effective at reducing α for all particles in all calculations to $\alpha \approx \alpha_0$, where $\alpha_0 = 0.1$. For the $n_x = 256$ particle calculation, this would imply that $v^{\text{AV}} \approx 18v$, that is, the numerical dissipation is 18 times greater than the Navier–Stokes viscosity. For our highest resolution calculation ($n_x = 2048$), it would thus be expected that v^{AV} be roughly twice as large as v , with the Navier–Stokes viscosity becoming dominant over the numerical dissipation only for resolutions $n_x > 4096$ particles.

The rate at which kinetic energy is dissipated by each source of dissipation may be directly measured from the calculations. Both the artificial viscosity and Navier–Stokes viscosity are explicitly added sources of dissipation. The rate of change of kinetic energy can be measured by computing

$$\frac{dE_k}{dt} = \sum_a m_a v_a \cdot \frac{dv_a}{dt}, \quad (35)$$




Figure 6. Rate of change of kinetic energy (dE_k/dt) from sources of dissipation, as directly measured from the $n_x = 2048$ particle calculation. The grey dotted line represents twice the average Navier–Stokes dissipation. In agreement with analytic expectation, the artificial viscosity is about twice as dissipative as the explicit Navier–Stokes viscosity at this resolution. The increase of dissipation by the artificial viscosity at $t \approx 1$ corresponds to when the particles first break off the initial lattice arrangement.

and substituting in dv/dt from the artificial viscosity (from equation 13) and Navier–Stokes viscosity (equation 20).

Fig. 6 shows dE_k/dt for the artificial viscosity, Navier–Stokes viscosity, and the sum of the two for the $n_x = 2048$ calculation. This measurement confirms the analytic estimate that the rate at which kinetic energy is dissipated by the artificial viscosity is approximately twice that of the Navier–Stokes viscosity. The β viscosity must be an insignificant contribution to the overall numerical dissipation in this case. The sharp increase of dissipation by the artificial viscosity at $t \approx 1$ occurs from the particles breaking off the initial lattice arrangement.

These analyses suggest that the non-linear evolution of the SPH calculations is influenced considerably by numerical dissipation, with a ratio of 2:1 of numerical-to-physical dissipation for our highest resolution calculation. The ratio of kinetic energy dissipation is expected to be equal for $n_x = 4096$ particles, with the explicit Navier–Stokes viscosity becoming the dominant source of dissipation only at resolutions greater than that.

It is worth noting that the calculations in this work have used the Morris & Monaghan (1997) switch with $\alpha_0 = 0.1$. More sophisticated artificial viscosity limiters have been developed in recent years that are better at reducing numerical dissipation (Cullen & Dehnen 2010; Read & Hayfield 2012; Wadsley et al. 2017), in particular permitting the use of $\alpha_0 = 0$ in some instances. Using one of these more advanced switches should improve the convergence properties of SPH on this Kelvin–Helmholtz test. This is explored further in Section 3.6.

3.6 Reducing the artificial viscosity ($\alpha = 0$)

To investigate the convergence properties when the Navier–Stokes viscosity is the dominant source of kinetic energy dissipation, we have repeated our highest resolution calculation ($n_x = 2048$) using a fixed $\alpha = 0$ for all particles. By reducing the numerical dissipation, this result should be indicative of higher resolutions calculations that use typical artificial viscosity parameters and switches.

It is important to be careful about setting $\alpha = 0$ since the primary role of artificial viscosity in these calculations is to maintain particle regularity. As discussed in Section 3.5, the shear viscosity




Figure 7. The colour field in the region $x, y \in [0, 1]$ at $t = 6$ (left column) and $t = 8$ (right column) for the standard $n_x = 2048$ particle SPH calculation using the Morris & Monaghan (1997) viscosity switch (first row), using a fixed $\alpha = 0$ to reduce artificial viscosity (middle row), and the D2048 reference solution for comparison (bottom row). At this resolution, the explicit Navier–Stokes viscosity is sufficient to keep the particles regular. By reducing the numerical dissipation ($\alpha = 0$), the qualitative agreement between the SPH and D2048 solutions is dramatically improved in comparison to calculations that employ the Morris & Monaghan (1997) artificial viscosity switch.

introduced by the artificial viscosity for the $n_x = 2048$ particle calculation is about twice that of the explicitly added Navier–Stokes viscosity. It is reasonable to expect that the Navier–Stokes viscosity may be sufficient by itself to keep the particles regular for this resolution. Note that the β term of the artificial viscosity is retained for these calculations since it represents a von Neumann–Richtmyer-type viscosity. Its contribution to the overall dissipation of kinetic energy should be small.

Fig. 7 shows the colour field at $t = 6$ and 8 for fixed $\alpha = 0$. The qualitative agreement between the SPH calculation and the reference solution is significantly improved in comparison to the standard calculation (cf. Fig. 3). Most noticeable is the inner tip of the curl, which is wound as tightly as the reference solution in both time snapshots. The thick red filament in the upper region of the curl at $t = 8$ is concentrated towards the upper right, as in the reference solution, whereas this feature in the standard calculation is diffuse across the top of the curl. Overall, numerous features




Figure 8. \mathcal{L}_2 error of SPH calculations using fixed $\alpha = 0$ in the artificial viscosity with respect to the D2048 solution at $t = 2, 4, 6,$ and 8 . The connecting lines are for visual aid. The solution of the $n_x = 2048$ particle calculation with reduced numerical dissipation is in better agreement to the D2048 reference solution than the standard SPH calculation (cf. Fig. 5).

can be identified by eye as resembling the reference solution more closely in this calculation than for the standard calculation.

Fig. 8 shows the \mathcal{L}_2 error for fixed $\alpha = 0$ calculations for resolutions of $n_x = 256$ – 2048 particles. The convergence rates are given in Table 2. At low resolutions, substantial noise occurs in the particle arrangement due to the absence of sufficient viscosity (either numerical or physical). The results at low resolution are not significantly improved in the non-linear regime, except for at $t = 2$ since the particles have yet to rearrange significantly. However, at high resolution ($n_x = 2048$), the \mathcal{L}_2 error is uniformly lower than the standard calculation (cf. Fig. 5), by approximately 20 per cent at $t = 2$ and 30 per cent in the late non-linear regime ($t \geq 4$). The solution obtained at this resolution is more accurate than the standard calculation, suggesting that SPH with typical artificial viscosity parameters will indeed converge to the reference solution as the resolution increases due to decreasing numerical dissipation. This is substantiated by the improved rate of convergence when the artificial viscosity is reduced (cf. Table 2), although the rate of convergence still remains sublinear in the non-linear regime of the instability.

It is expected that the results of this calculation would be similar to a calculation that used a higher order artificial viscosity limiter, such as the approaches by Cullen & Dehnen (2010), Read & Hayfield (2012), or Wadsley et al. (2017). The Cullen & Dehnen (2010) artificial viscosity switch, for example, uses the time derivative of the divergence of the velocity, $d(\nabla \cdot \mathbf{v})/dt$, promising more accurate shock detection than the Morris & Monaghan (1997) switch. In practice, this permits the use of $\alpha_0 = 0$ (e.g. equation 18), granting significant reduction in numerical dissipation. Therefore, it is reasonable to expect that more accurate switches would allow for α to be near zero for this problem, and that the results obtained here are indicative of results that use a higher order artificial viscosity switch.

3.7 Kernel bias

The calculations performed in this work used the M8 septic spline. It is important to note that the use of a high-order kernel is not an intrinsic requirement to capture the Kelvin–Helmholtz instability. Rather, a high-order B-spline has been used to minimize velocity noise caused by errors in the pressure gradient. This is examined in greater detail below.




Figure 9. The colour field in the region $x, y \in [0, 1]$ for $n_x = 1024$ particle SPH calculations at $t = 4$ using the cubic, quartic, quintic, septic, and nonic splines (left to right). All produce a dominant curl of the correct mode frequency, but the cubic spline, and to a lesser degree the quartic spline, evolves incorrectly in comparison to the reference solution (cf. Fig. 3). The quintic spline shows substantial noise along the spiral arm edges, arising due to particle motion. The colour fields of the septic and nonic spline results are smooth and virtually indistinguishable, demonstrating that the kernel bias does not affect the evolution for these high-order splines.

Fig. 9 shows the colour field at $t = 4$ for $n_x = 1024$ particle calculations using the B-splines between cubic (M4) and nonic (M10). In all cases, a single dominant curl forms, mixing the two fluids along the interface. However, the shape of this vortex changes distinctly with respect to the quality of the smoothing kernel. The cubic spline does not resemble the reference solution except that the correct mode has developed. It lacks the interior structure present in the reference solution. The calculation with the quartic spline improves upon this. It resembles the reference solution better than the cubic spline result, though the curl is not wound as tightly as it should, and the filaments are broader and stunted. It is also narrower in the y -direction. The quintic spline improves upon this further. The overall shape and structure of the vortex is in close proximity to the reference solution. However, the interface along the spirals shows substantial noise, noticeable by the ‘jaggedness’ of the spiral edges.

The best results are obtained when the septic or nonic spline is used. In both cases, a smooth spiral structure develops that closely resembles the reference solution, as discussed in Section 3.2. The noise present in the quintic spline result is gone. The nonic spline solution is virtually identical to the septic spline result, demonstrating that the kernel bias is no longer the dominant source of error.

To better understand the effect of the kernel bias on the growth of the Kelvin–Helmholtz instability, we consider the maximum kinetic energy in the y -direction. Fig. 10 shows the maximum of $\frac{1}{2}\rho v_y^2$ over all particles. This maximum is initially $\sim 10^{-4}$, as defined by the initial conditions.

For the septic and nonic spline calculations, the maximum y -kinetic energy undergoes steady exponential amplification over four orders of magnitude between $0 < t < 2$. This is the expected behaviour in the linear regime of the Kelvin–Helmholtz instability, and further establishes that the non-linear regime begins at $t > 2$. Additionally, the results between the septic and nonic splines are virtually indistinguishable at all times, showcasing further that the kernel bias has no effect on the results obtained in the standard calculations.

On the other hand, the cubic and quartic spline results show an immediate discontinuous jump of two to three orders of magnitude in y -kinetic energy. The maximum y -velocities are $\mathcal{M} \approx 0.05$ and $\mathcal{M} \approx 0.03$, respectively, significantly larger than the $\mathcal{M} \approx 0.0025$




Figure 10. The maximum y -kinetic energy, $\frac{1}{2}\rho v_y^2$, over all particles for calculations using the cubic to nonic spline kernels. The results using the cubic and quartic splines undergo sharp initial jumps, occurring due to errors in the pressure gradient leading to particle rearrangement. The quintic spline result shows similar behaviour. In contrast, calculations using the septic spline show steady exponential growth, as expected of the Kelvin–Helmholtz instability. The nonic spline result matches the septic spline result, showing that the kernel bias no longer affects the evolution of the instability.

amplitude of the initial velocity perturbation. Similar behaviour was noted by McNally et al. (2012b). Despite the differing growth of y -kinetic energy, the peak energy for the cubic and quartic spline calculations is of similar magnitude to the high-order splines. The peaks are delayed, however, occurring at later times. Overall, the initial jump and delayed peak of y -kinetic energy demonstrate that the evolution of the instability is strongly affected by the kernel bias for these low-order kernels (see also Fig. 9).

The quintic spline result is between these two extremes. It exhibits a sharp jump in y -kinetic energy, as with the low-order splines, but at $t \approx 0.25$ and not as large in magnitude, with a corresponding maximum y -velocity of $\mathcal{M} \approx 0.01$. However, between $1 < t < 2$, the y -kinetic energy undergoes steady exponential amplification along the track of the high-order splines. This is reflected in the colour field (Fig. 9), which qualitatively contains the features of the reference solution, but with noise introduced by the early sharp increase in energy.

The sharp increase in y -kinetic energy for the low-order splines is due to errors in the pressure gradient, which generates spurious velocity noise. The error in the discretization of the pressure gradient scales as $\mathcal{O}(1)$ with respect to resolution. That is, the errors are not reduced by increasing the resolution of the calculation. However, the pressure gradient in the standard formulation of SPH is derived from the discretized Lagrangian using the density summation (Price 2012), such that the SPH equations of motion are exactly the Lagrangian equations of motion for the particle system. The errors in the pressure gradient act to regularize the particles, pushing them towards a better arrangement, which in turn leads to better interpolation properties. This means that maintaining particle regularity is inherently built into SPH, though at the cost of a small degree of velocity noise. Using a more accurate discretization for the pressure gradient is problematic because it requires the introduction of a particle regularization scheme. Instead, the primary approach to reduce pressure gradient errors is to use a better smoothing kernel. This has led to investigation of new kernels, such as the Wendland family of kernels (Dehnen & Aly 2012), but the same effect can be obtained by switching to high-order splines.

Fig. 10 demonstrates that the magnitude of the velocity noise decreases as the quality of the smoothing kernel improves. It is clear that the amplitude of the velocity perturbation used to seed the Kelvin–Helmholtz instability must be greater than the velocity noise introduced by particle arrangement in order to correctly model the linear regime. For these calculations, the lowest order spline kernel that meets this criterion is the septic spline. This may not be the sole requirement. It would additionally be expected that other modes not seeded in the initial conditions will be excited by any introduced velocity noise, as if they had initial amplitudes similar to the threshold specific to each kernel. The growth of these spurious modes would have the potential to affect the solution.

This conjecture was tested by restarting the $n_x=2048$ calculation at $t = 2$ with the smoothing kernel switched from the septic spline to the cubic spline. In this way, the linear growth phase will be accurately modelled using a high-order kernel before switching to a low-order kernel. Fig. 11 shows the colour field of the cubic spline restarted calculation at $t = 4$, alongside the septic spline calculation for comparison. A substantial degree of noise is present in the vortex structure of the cubic spline calculation, occurring due to velocity noise introduced by errors in the pressure gradient. This demonstrates that a high-quality smoothing kernel is necessary for both the linear and non-linear phases.

3.8 Calculations without physical dissipation

Astrophysical simulations often rely solely on numerical dissipation, that is, without explicit physical dissipation. Here we present the results of calculations without Navier–Stokes viscosity or thermal conductivity. The artificial viscosity is as described in Section 2.3. An artificial thermal conductivity is applied according to

$$\frac{du_a}{dt} = \frac{1}{10} \sum_b \frac{m_b}{\bar{\rho}_{ab}} \sqrt{\frac{|P_a - P_b|}{\bar{\rho}_{ab}}} (u_a - u_b) \hat{\mathbf{r}}_{ab} \cdot \nabla_a W_{ab}(h_a), \quad (36)$$

making use of the Price (2008) switch. Since advection is dissipationless in SPH, there is no numerical dissipation associated with the colour field in the absence of explicitly prescribed terms. In such a case, the total colour entropy would remain constant over time as the colour of each particle would remain fixed. Rather than creating a contrived ‘artificial colour diffusion’ term for the colour field for




Figure 11. The $n_x=2048$ calculation at $t = 4$ (left-hand panel) alongside a calculation where the smoothing kernel was changed from the septic spline to the cubic spline at $t = 2$ (right-hand panel). Switching to the low-order kernel partway through the calculation introduces substantial noise into the solution.

the purposes of diffusing colour between particles, we simply retain the ‘physical’ dissipation term given by equation (23).

Fig. 12 shows the colour field for these calculations at $t = 4$ and 6, alongside the standard calculations with physical dissipation terms. Slight differences are evident to the calculation with physical dissipation; however, as the artificial dissipation is the dominant source of kinetic dissipation for both calculations (cf. Section 3.6), the differences are minimal. Since the artificial viscosity can be equated to a Navier–Stokes shear and bulk viscosity, it is expected that a more complex vortex structure would develop only when the artificial dissipation is lower than the physical dissipation. That is, when the Reynolds number exceeds 10^5 , occurring for resolutions $n_x > 4096$ particles.

3.9 Calculations with a density contrast

Lecoanet et al. (2016) studied the Kelvin–Helmholtz for both initial conditions that have uniform density and a difference in density between the two flow regions. Including a density contrast introduces a significant degree of complexity to the structure and evolution of the instability. In this case, the computational demand to achieve convergence is strict, with Lecoanet et al. (2016) only achieving convergence between ATHENA and DEDALUS when $16\,384 \times 32\,768$ grid cells were used.

We study the behaviour of SPH on the Kelvin–Helmholtz instability problem of Lecoanet et al. (2016) that has an initial density contrast. The initial density profile is specified as in equation (6) with $\Delta\rho/\rho_0 = 1$. The particles are arranged using the stretch mapping technique (Price et al. 2017) to achieve the correct density profile. The calculations are performed for $n_x = 512, 1024,$ and 2048 particles. The D4096 solution from Lecoanet et al. (2016) is used as the reference solution.

Fig. 13 shows the colour field at $t = 2, 4,$ and 6 for the SPH calculations for resolutions up to $n_x = 2048$ particles. The instability is triggered in all cases, with a curled vortex forming by $t = 2$ that is similar in shape to the reference solution. The similarity improves as the resolution increases.

At $t = 4$, the physical thermal conductivity is insufficient to properly mix the two flow regions for the $n_x = 512$ particle calculation, resulting in a clumpy structure. This issue is alleviated at higher resolution ($n_x = 2048$). Including an artificial thermal

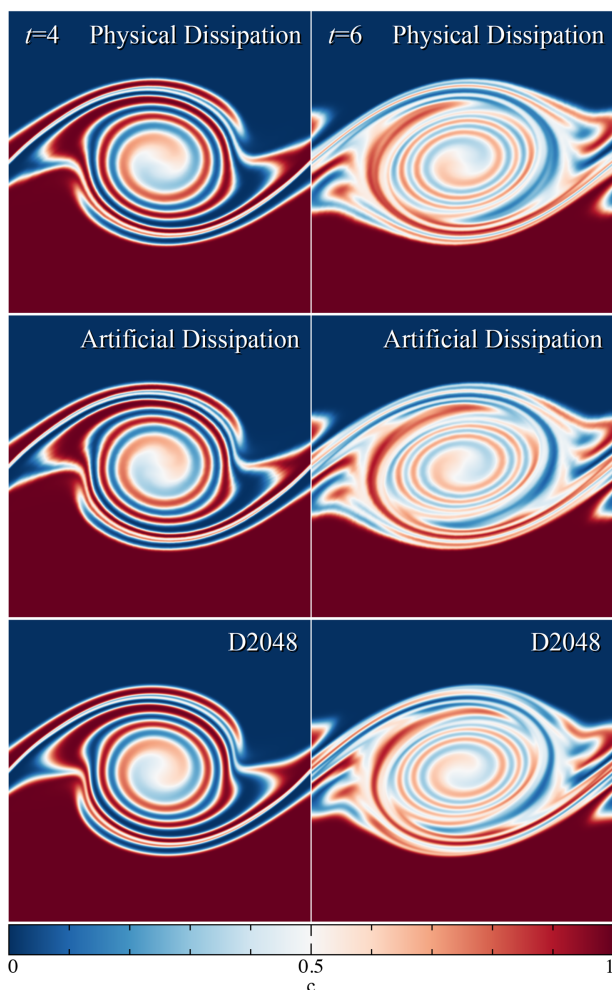


Figure 12. Calculations relying solely on artificial dissipation terms for the velocity and internal energy fields, contrasted with the standard SPH calculations with physical dissipation terms, and the DEDALUS reference solution. At this resolution, $n_x = 2048$ particles, there is no significant difference in the vortex structure as the artificial dissipation is the dominant source of dissipation even when physical dissipation is present.

conductivity would help prevent this issue from occurring at low resolution (Price 2008).

The non-linear evolution of the vortex structure for the $n_x = 1024$ and 2048 particle SPH calculations exhibits a greater degree of mixing in comparison to the reference solution. Similar behaviour was found for ATHENA calculations (Lecoanet et al. 2016) at lower resolution, thus this outcome is not surprising given that the highest resolution ATHENA calculation has length resolution one-eighth that of the $n_x = 2048$ particle SPH calculation. Overall, the qualitative behaviour of the SPH calculations is as expected for the resolutions simulated.

4 CONCLUSION

We have performed calculations of the Kelvin–Helmholtz instability using SPH. Our tests have used the unstratified, uniform density test of Lecoanet et al. (2016) with a Reynolds number of $\text{Re}=10^5$. The calculations have included a passive scalar ‘colour’ field to quantify the degree of mixing by the instability. Physical dissipation terms

have been included in order to control the evolution of the instability in the non-linear regime.

We have compared our results qualitatively and quantitatively to the D2048 results obtained using the pseudo-spectral code DEDALUS, as presented by Lecoanet et al. (2016). Results have been examined in both the linear and non-linear regimes of the instability. Our key observations are:

(i) The SPH calculations qualitatively agreed with the reference solution on the linear and non-linear evolution of the Kelvin–Helmholtz instability. A single dominant curl formed in the linear regime, which continued to wind producing substantial substructure in the non-linear regime. At late times ($t > 8$), the instability, in combination with the physical dissipation terms, produced a well-mixed interface. The internal structure of the vortices resembled the reference solution for the times compared, with the visual agreement improving as the resolution was increased.

(ii) No modifications to SPH were required to activate or correctly evolve the instability. The Kelvin–Helmholtz instability can be captured using the standard density formulation of SPH with artificial dissipation terms, of the kind that has been used for decades.

(iii) The exponential growth rate of the seeded mode in the linear regime was in agreement with the expected growth rate of $\propto \exp(3.227t)$, as obtained numerically using the PSECAS code (Berlok & Pfrommer 2019). This is half the analytic estimate for an incompressible fluid with sharp interfaces. The SPH growth rate was linearly converging to the expected rate, and was accurate even at our lowest resolution of $n_x = 256$ particles.

(iv) The Kelvin–Helmholtz instability can be activated in SPH even for modest resolutions. For our lowest resolution calculation of $n_x = 256$ particles, the total entropy of the colour field increased monotonically along a similar curve to the reference solution. Even though small-scale features in the vortices were not as well represented at low resolutions, the overall degree of mixing experienced, with respect to the reference solution, was similar.

(v) The rate of convergence of \mathcal{L}_2 error was linear in the linear regime of the Kelvin–Helmholtz instability. That is, the \mathcal{L}_2 error decreased by a factor of 2 for each factor of 2 increase in resolution in each spatial direction. In the non-linear regime, the \mathcal{L}_2 errors were decreasing, but the convergence rate was sublinear.

(vi) Numerical dissipation was the primary factor in convergence rate. Performing calculations with reduced artificial viscosity, using $\alpha = 0$, yielded results that qualitatively and quantitatively better agreed with the reference solution. The filaments and spiral structure of the vortices more closely resembled the reference solution, and the \mathcal{L}_2 errors were lower at all times sampled. That the calculations could be performed with $\alpha = 0$ was permissible only because, at high resolution, the physical Navier–Stokes viscosity was as dissipative as the numerical dissipation and was sufficient to maintain particle regularity.

(vii) It was found that the septic spline was the minimum order of B-spline required for these Kelvin–Helmholtz calculations in order to accurately model both the linear and non-linear regimes. High-order kernels are needed to minimize the generation of spurious velocity noise from errors in the pressure gradient. This background velocity noise must be less than the initial velocity perturbation in order to correctly model the linear regime. Velocity noise can additionally excite other modes, disrupting the non-linear phase even when the linear phase is correctly modelled, as demonstrated using the cubic spline.

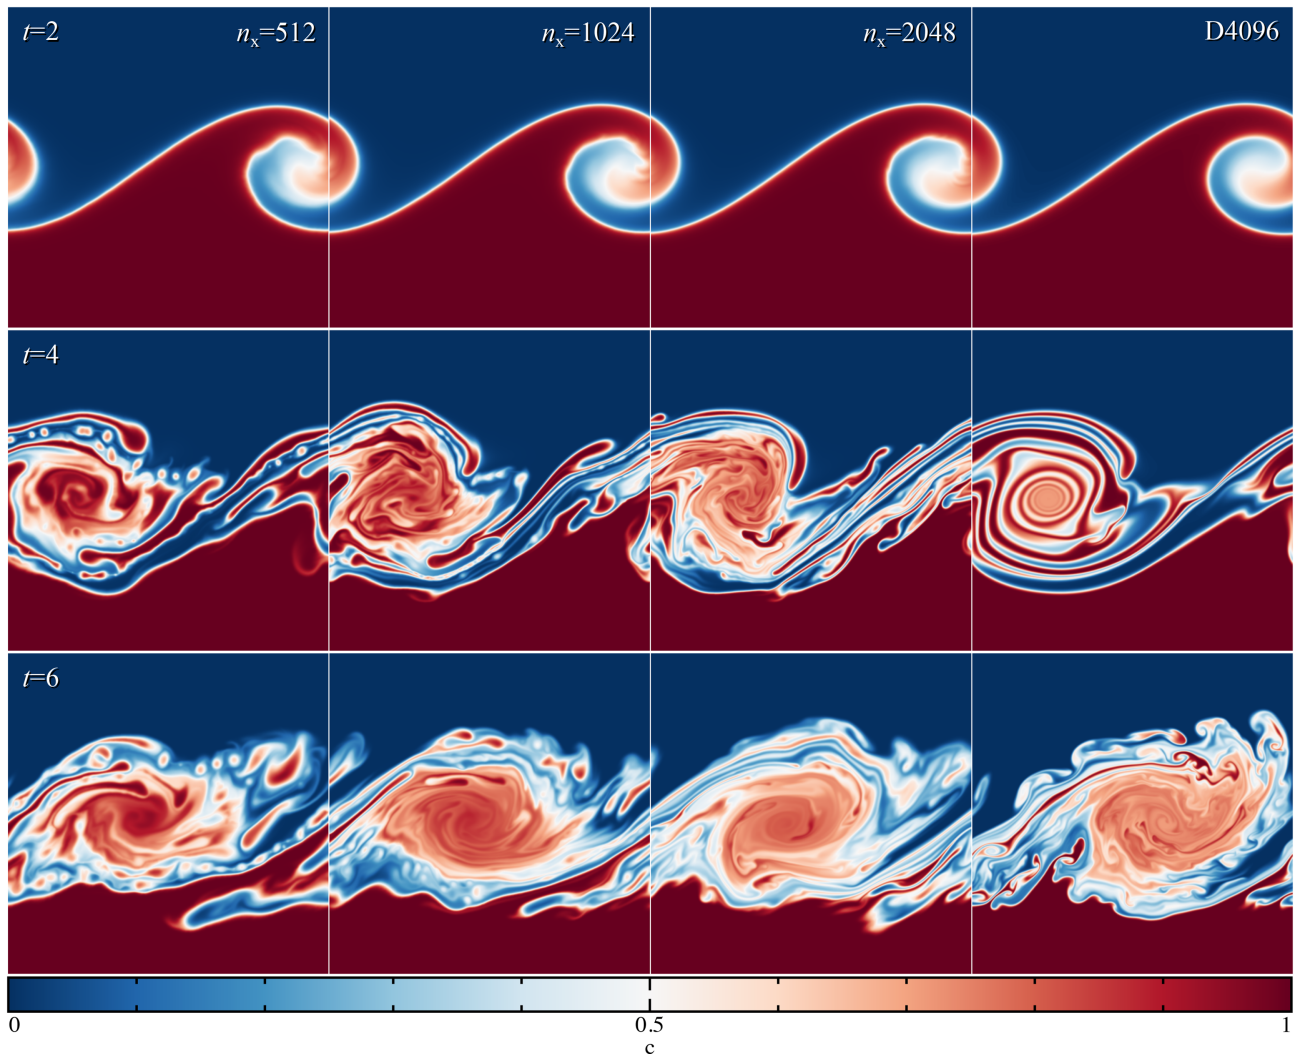


Figure 13. The colour field with a density contrast in the initial conditions ($\Delta\rho/\rho_0 = 1$) for SPH calculations of $n_x = 512$, 1024, and 2048 particles (first to third columns), with the DEDALUS reference solution for comparison (D4096) at $t = 2, 4$, and 6 (top to bottom rows). The instability is triggered in SPH even at modest resolutions. The vortex structure is trending towards the reference solution as the resolution increases, but would require significantly higher resolution to match the reference solution for $t \geq 4$.

(viii) Calculations employing an initial density contrast qualitatively agree with the reference solution. A single curl forms during the linear regime, which continues to wind producing substantial substructure in the non-linear regime. The interior structure of the vortex is overmixed compared to the reference solution, but similar behaviour was found for ATHENA calculations at similar resolutions (Lecoanet et al. 2016).

The true test of any numerical method is whether it converges to an agreed upon solution, and, in this work, it has been demonstrated that SPH is converging towards the Kelvin–Helmholtz solution of Lecoanet et al. (2016). It would be of interest to test the convergence properties of voronoi moving mesh or other particle-based schemes. Our core conclusion is that SPH can correctly model the Kelvin–Helmholtz instability.

ACKNOWLEDGEMENTS

Thank you to Daniel Lecoanet for providing the data from the DEDALUS simulations, Daniel Price for useful discussions on phys-

ical dissipation terms in SPH, and the anonymous referee whose comments helped improve the quality of this work. TST was supported by a CITA Postdoctoral Research Fellowship.

REFERENCES

- Abel T., 2011, *MNRAS*, 413, 271
 Agertz O. et al., 2007, *MNRAS*, 380, 963
 Artymowicz P., Lubow S. H., 1994, *ApJ*, 421, 651
 Bauer A., Springel V., 2012, *MNRAS*, 423, 2558
 Berlok T., Pfrommer C., 2019, *MNRAS*, 485, 908
 Brandenburg A., Dobler W., 2002, *Comput. Phys. Commun.*, 147, 471
 Brookshaw L., 1985, *Proc. Astron. Soc. Aust.*, 6, 207
 Bryan G. L. et al., 2014, *ApJS*, 211, 19
 Burns K. J., Vasil G. M., Oishi J. S., Lecoanet D., Brown B. P., 2016, *Astrophysics Source Code Library*, record ascl:1603.015
 Burns K. J., Vasil G. M., Oishi J. S., Lecoanet D., Brown B. P., 2019, preprint ([arXiv:1905.10388](https://arxiv.org/abs/1905.10388))
 Cabezón R. M., García-Senz D., Figueira J., 2017, *A&A*, 606, A78
 Chandrasekhar S., 1961, *Hydrodynamic and Hydromagnetic Stability*, Clarendon Press, Oxford

Cha S.-H., Inutsuka S.-I., Nayakshin S., 2010, *MNRAS*, 403, 1165
 Cleary P. W., Monaghan J. J., 1999, *J. Comput. Phys.*, 148, 227
 Cullen L., Dehnen W., 2010, *MNRAS*, 408, 669
 Cummins S. J., Rudman M., 1999, *J. Comput. Phys.*, 152, 584
 Dehnen W., Aly H., 2012, *MNRAS*, 425, 1068
 Duffell P. C., MacFadyen A. I., 2015, *MNRAS*, 449, 2718
 Flebbe O., Muenzel S., Herold H., Riffert H., Ruder H., 1994, *ApJ*, 431, 754
 Frontiere N., Raskin C. D., Owen J. M., 2017, *J. Comput. Phys.*, 332, 160
 García-Senz D., Cabezón R. M., Escartín J. A., 2012, *A&A*, 538, A9
 Goffrey T. et al., 2017, *A&A*, 600, A7
 Heß S., Springel V., 2010, *MNRAS*, 406, 2289
 Hopkins P. F., 2013, *MNRAS*, 428, 2840
 Hopkins P. F., 2015, *MNRAS*, 450, 53
 Hubber D. A., Falle S. A. E. G., Goodwin S. P., 2013, *MNRAS*, 432, 711
 Hubber D. A., Rosotti G. P., Booth R. A., 2018, *MNRAS*, 473, 1603
 Junk V., Walch S., Heitsch F., Burkert A., Wetzstein M., Schartmann M., Price D., 2010, *MNRAS*, 407, 1933
 Kawata D., Okamoto T., Gibson B. K., Barnes D. J., Cen R., 2013, *MNRAS*, 428, 1968
 Kravtsov A. V., Klypin A., Hoffman Y., 2002, *ApJ*, 571, 563
 Kravtsov A. V., Klypin A. A., Khokhlov A. M., 1997, *ApJS*, 111, 73
 Lecoanet D. et al., 2016, *MNRAS*, 455, 4274
 Lodato G., Price D. J., 2010, *MNRAS*, 405, 1212
 Maron J. L., McNally C. P., Mac Low M.-M., 2012, *ApJS*, 200, 6
 McNally C. P., Maron J. L., Mac Low M.-M., 2012a, *ApJS*, 200, 7
 McNally C. P., Lyra W., Passy J.-C., 2012b, *ApJS*, 201, 18
 Monaghan J. J., 1989, *J. Comput. Phys.*, 82, 1
 Monaghan J. J., 1992, *ARA&A*, 30, 543
 Monaghan J. J., 1997, *J. Comput. Phys.*, 136, 298
 Monaghan J. J., 2002, *MNRAS*, 335, 843
 Monaghan J. J., 2005, *Rep. Prog. Phys.*, 68, 1703
 Morris J. P., Monaghan J. J., 1997, *J. Comput. Phys.*, 136, 41
 Murray J. R., 1996, *MNRAS*, 279, 402
 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., 2005, in Plewa T., Linde T., Gregory Weirs V., eds, *Adaptive Mesh Refinement-theory and Applications*, Lecture Notes in Computational Science and Engineering, vol 41. Springer, Berlin, Heidelberg, p. 341
 Price D. J., 2008, *J. Comput. Phys.*, 227, 10040
 Price D. J., 2010, *MNRAS*, 401, 1475
 Price D. J., 2012, *J. Comput. Phys.*, 231, 759
 Price D. J. et al., 2017, preprint ([arXiv:1702.03930](https://arxiv.org/abs/1702.03930))
 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
 Robertson B. E., Kravtsov A. V., Gnedin N. Y., Abel T., Rudd D. H., 2010, *MNRAS*, 401, 2463
 Saitoh T. R., Makino J., 2013, *ApJ*, 768, 44
 Schoenberg I. J., 1946, *Q. Appl. Math.*, 4, 45
 Shadloo M. S., Yildiz M., 2011, *Int. J. Numer. Methods Eng.*, 87, 988
 Sijacki D., Springel V., 2006, *MNRAS*, 371, 1025
 Sijacki D., Vogelsberger M., Kereš D., Springel V., Hernquist L., 2012, *MNRAS*, 424, 2999
 Springel V., 2010a, *ARA&A*, 48, 391
 Springel V., 2010b, *MNRAS*, 401, 791
 Springel V., Hernquist L., 2002, *MNRAS*, 333, 649

Stone J. M., Gardiner T. A., Teuben P., Hawley J. F., Simon J. B., 2008, *ApJS*, 178, 137
 Tricco T. S., Price D. J., 2012, *J. Comput. Phys.*, 231, 7214
 Valcke S., de Rijcke S., Rödigier E., Dejonghe H., 2010, *MNRAS*, 408, 71
 Valdarnini R., 2012, *A&A*, 546, A45
 Wadsley J. W., Keller B. W., Quinn T. R., 2017, *MNRAS*, 471, 2357
 Wadsley J. W., Veeravalli G., Couchman H. M. P., 2008, *MNRAS*, 387, 427
 Wang L. F., Ye W. H., Li Y. J., 2010, *Phys. Plasmas*, 17, 042103
 Watkins S. J., Bhattal A. S., Francis N., Turner J. A., Whitworth A. P., 1996, *A&AS*, 119, 177
 Yue T., Pearce F., Kruisbrink A., Morvan H., 2015, *Int. J. Numer. Methods Fluids*, 78, 283

APPENDIX A: B-SPLINE KERNELS

The family of B-spline kernels (Schoenberg 1946) can be written in the functional form

$$W_{ab}(h_a) = \frac{\sigma}{h^p} w(q), \quad (\text{A1})$$

where σ is the normalization, p is the number of dimensions, and $q = r_{ab}/h_a$. Each spline is a piecewise polynomial. Table A1 lists the M4 (cubic) to M10 (nonic) polynomial spline kernels. In the table, the kernel function, $w(q)$, is for the range $R - 1 \leq q < R$, where R is the radial extent of the kernel. The full piecewise function can be constructed from this.

For clarity, we also explicitly write the cubic spline,

$$w(q) = \begin{cases} (2-q)^3 - 4(1-q)^3 & q < 1, \\ (2-q)^3 & 1 \leq q < 2, \\ 0 & q \geq 2, \end{cases} \quad (\text{A2})$$

quintic spline,

$$w(q) = \begin{cases} (3-q)^5 - 6(2-q)^5 + 15(1-q)^5 & q < 1, \\ (3-q)^5 - 6(2-q)^5 & 1 \leq q < 2, \\ (3-q)^5 & 2 \leq q < 3, \\ 0 & q \geq 3, \end{cases} \quad (\text{A3})$$

and septic spline,

$$w(q) = \begin{cases} (4-q)^7 - 8(3-q)^7 + 28(2-q)^7 & q < 1, \\ -56(1-q)^7 & 1 \leq q < 2, \\ (4-q)^7 - 8(3-q)^7 + 28(2-q)^7 & 2 \leq q < 3, \\ (4-q)^7 & 3 \leq q < 4, \\ 0 & q \geq 4. \end{cases} \quad (\text{A4})$$

The increased computational cost for higher order kernels compared to the cubic spline can be estimated by the increased number of particles within the radial extent of the kernel. In two dimensions, such as the calculations in this paper, it would imply that the increased computational cost of switching to the quintic spline over the cubic spline would be $2.25 \times$. For the septic spline over the cubic spline, this would be $4 \times$. There is a further floating point operation cost incurred to compute the more complex splines.

Table A1. List of bell-shaped spline kernels, with the functional form, normalization (σ), and radial extent (R). The functional form, $w(q)$, is for the piece $R - 1 \leq q < R$, from which the full piecewise function can be constructed.

Name	$w(q)$	σ (1D)	σ (2D)	σ (3D)	R (q)
Cubic	$(2 - q)^3 - 4(1 - q)^3$	3!	$5/(14\pi)$	$6/(4!\pi)$	2
Quartic	$(\frac{5}{2} - q)^4 - 5(\frac{3}{2} - q)^4 + 10(\frac{1}{2} - q)^4$	4!	$96/(1\ 199\pi)$	$6/(5!\pi)$	2.5
Quintic	$(3 - q)^5 - 6(2 - q)^5 + 15(1 - q)^5$	5!	$7/(478\pi)$	$6/(6!\pi)$	3
Sextic	$(\frac{7}{2} - q)^6 - 7(\frac{5}{2} - q)^6 + 21(\frac{3}{2} - q)^6 - 35(\frac{1}{2} - q)^6$	6!	$256/(113\ 149\pi)$	$6/(7!\pi)$	3.5
Septic	$(4 - q)^7 - 8(3 - q)^7 + 28(2 - q)^7 - 56(1 - q)^7$	7!	$9/(29\ 749\pi)$	$6/(8!\pi)$	4
Octic	$(\frac{9}{2} - q)^8 - 9(\frac{7}{2} - q)^8 + 36(\frac{5}{2} - q)^8 - 84(\frac{3}{2} - q)^8 + 126(\frac{1}{2} - q)^8$	8!	$512/(14\ 345\ 663\pi)$	$6/(9!\pi)$	4.5
Nonic	$(5 - q)^9 - 10(4 - q)^9 + 45(3 - q)^9 - 120(2 - q)^9 + 210(1 - q)^9$	9!	$11/(2\ 922\ 230\pi)$	$6/(10!\pi)$	5

APPENDIX B: DERIVATION OF THE NAVIER-STOKES HEATING TERM

The heat gained from the two first derivatives implementation of Navier–Stokes viscosity in SPH can be derived as follows. The total energy is the sum of kinetic and thermal energy from all particles,

$$E = \sum_a m_a \left(\frac{1}{2} v_a^2 + u_a \right). \quad (\text{B1})$$

Conservation of energy requires the total energy to remain constant, thus

$$\frac{dE}{dt} = \sum_a m_a \left(\mathbf{v}_a \cdot \frac{d\mathbf{v}_a}{dt} + \frac{du_a}{dt} \right) = 0. \quad (\text{B2})$$

Rearranging yields

$$\sum_a m_a \frac{du_a}{dt} = - \sum_a m_a \mathbf{v}_a \cdot \frac{d\mathbf{v}_a}{dt}. \quad (\text{B3})$$

Substituting in the Navier–Stokes term from the SPH momentum equation (equation 20) yields, for the RHS,

$$\sum_a \sum_b m_a m_b v_a^i \left[\frac{\Pi_a^{ij}}{\Omega_a \rho_a^2} \nabla_a^j W_{ab}(h_a) + \frac{\Pi_b^{ij}}{\Omega_b \rho_b^2} \nabla_b^j W_{ab}(h_b) \right]. \quad (\text{B4})$$

Substituting in the definition the Navier–Stokes stress tensor (equation 4) yields

$$\sum_a \sum_b m_a m_b v \left[v_a^i \left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} - \frac{2}{3} \frac{\partial v_a^k}{\partial x_a^k} \delta^{ij} \right) \frac{\nabla_a^j W_{ab}(h_a)}{\Omega_a \rho_a} + v_a^i \left(\frac{\partial v_b^i}{\partial x_b^j} + \frac{\partial v_b^j}{\partial x_b^i} - \frac{2}{3} \frac{\partial v_b^k}{\partial x_b^k} \delta^{ij} \right) \frac{\nabla_b^j W_{ab}(h_b)}{\Omega_b \rho_b} \right]. \quad (\text{B5})$$

Swapping the summation indices on the second term produces

$$\sum_a \sum_b m_a m_b v \left[v_a^i \left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} - \frac{2}{3} \frac{\partial v_a^k}{\partial x_a^k} \delta^{ij} \right) \frac{\nabla_a^j W_{ab}(h_a)}{\Omega_a \rho_a} \right] + \sum_a \sum_b m_a m_b v \left[v_b^i \left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} - \frac{2}{3} \frac{\partial v_a^k}{\partial x_a^k} \delta^{ij} \right) \frac{\nabla_b^j W_{ab}(h_a)}{\Omega_a \rho_a} \right]. \quad (\text{B6})$$

The kernel derivative in the second term can be substituted using the identity $\nabla_b W_{ab} \equiv -\nabla_a W_{ab}$, allowing the two terms to be combined to produce

$$\sum_a \sum_b m_a m_b v \left[\left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} - \frac{2}{3} \frac{\partial v_a^k}{\partial x_a^k} \delta^{ij} \right) v_{ab}^i \frac{\nabla_b^j W_{ab}(h_a)}{\Omega_a \rho_a} \right]. \quad (\text{B7})$$

By utilizing the definition of the difference derivative estimate (equation 19), one can arrive, after simplification, to

$$\sum_a m_a v \left[\left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} \right) \left(\frac{\partial v_a^i}{\partial x_a^j} \right) - \frac{2}{3} \left(\frac{\partial v_a^k}{\partial x_a^k} \right)^2 \right]. \quad (\text{B8})$$

The second term represents the divergence of the velocity, $\nabla \cdot \mathbf{v}$. The first term may be simplified by splitting the sum into two equal halves and swapping the i, j indices on the second half,

$$\sum_a m_a v \left[\frac{1}{2} \left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} \right) \left(\frac{\partial v_a^i}{\partial x_a^j} \right) - \frac{1}{3} \left(\frac{\partial v_a^k}{\partial x_a^k} \right)^2 \right] + \frac{1}{2} \left(\frac{\partial v_a^j}{\partial x_a^i} + \frac{\partial v_a^i}{\partial x_a^j} \right) \left(\frac{\partial v_a^j}{\partial x_a^i} \right) - \frac{1}{3} \left(\frac{\partial v_a^k}{\partial x_a^k} \right)^2, \quad (\text{B9})$$

which after combining yields

$$\sum_a m_a v \left[\left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} \right) \left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} \right) - \frac{2}{3} \left(\frac{\partial v_a^k}{\partial x_a^k} \right)^2 \right]. \quad (\text{B10})$$

The heat gain from the Navier–Stokes viscosity is obtained by substituting the preceding into equation (B3), yielding

$$\frac{du_a}{dt} = \frac{v}{2} \left[\left(\frac{\partial v_a^i}{\partial x_a^j} + \frac{\partial v_a^j}{\partial x_a^i} \right)^2 - \frac{2}{3} \left(\frac{\partial v_a^k}{\partial x_a^k} \right)^2 \right]. \quad (\text{B11})$$

This is guaranteed to yield positive definite increases in thermal energy, provided that the velocity derivatives are computed using difference derivative estimates.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.