

A validated non-linear Kelvin–Helmholtz benchmark for numerical hydrodynamics

D. Lecoanet,^{1,2*} M. McCourt,³ E. Quataert,^{1,2} K. J. Burns,⁴ G. M. Vasil,⁵
J. S. Oishi,⁶ B. P. Brown,⁷ J. M. Stone⁸ and R. M. O’Leary⁹

¹Physics Department, University of California, Berkeley, CA 94720, USA

²Astronomy Department and Theoretical Astrophysics Center, University of California, Berkeley, CA 94720, USA

³Institute for Theory and Computation, Center for Astrophysics, Harvard University, 60 Garden St, Cambridge, MA 02138, USA

⁴Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁵School of Mathematics and Statistics, University of Sydney, NSW 2006, Australia

⁶Department of Physics, Farmingdale State College, Farmingdale, NY 11735, USA

⁷Laboratory for Atmospheric and Space Physics and Department of Astrophysical and Planetary Sciences, University of Colorado, Boulder, CO 80309, USA

⁸Department of Astrophysical Sciences, Princeton University, Princeton, NJ 08544, USA

⁹JILA, University of Colorado and NIST, 440 UCB, Boulder, CO 80309-0440, USA

Accepted 2015 October 29. Received 2015 October 24; in original form 2015 September 3

ABSTRACT

The non-linear evolution of the Kelvin–Helmholtz instability is a popular test for code verification. To date, most Kelvin–Helmholtz problems discussed in the literature are ill-posed: they do not converge to any single solution with increasing resolution. This precludes comparisons among different codes and severely limits the utility of the Kelvin–Helmholtz instability as a test problem. The lack of a reference solution has led various authors to assert the accuracy of their simulations based on ad hoc proxies, e.g. the existence of small-scale structures. This paper proposes well-posed two-dimensional Kelvin–Helmholtz problems with smooth initial conditions and explicit diffusion. We show that in many cases numerical errors/noise can seed spurious small-scale structure in Kelvin–Helmholtz problems. We demonstrate convergence to a reference solution using both *ATHENA*, a Godunov code, and *DEDALUS*, a pseudo-spectral code. Problems with constant initial density throughout the domain are relatively straightforward for both codes. However, problems with an initial density jump (which are the norm in astrophysical systems) exhibit rich behaviour and are more computationally challenging. In the latter case, *ATHENA* simulations are prone to an instability of the inner rolled-up vortex; this instability is seeded by grid-scale errors introduced by the algorithm, and disappears as resolution increases. Both *ATHENA* and *DEDALUS* exhibit late-time chaos. Inviscid simulations are riddled with extremely vigorous secondary instabilities which induce *more* mixing than simulations with explicit diffusion. Our results highlight the importance of running well-posed test problems with demonstrated convergence to a reference solution. To facilitate future comparisons, we include as supplementary material the resolved, converged solutions to the Kelvin–Helmholtz problems in this paper in machine-readable form.

Key words: hydrodynamics – instabilities – methods: numerical.

1 INTRODUCTION

The Kelvin–Helmholtz (KH) instability results from a wide array of velocity–shear profiles in a continuous fluid, or across the interface between two distinct fluids. The instability is ubiquitous in nature, playing important roles in meteorology, oceanography,

and engineering. The KH instability plays a particularly prominent role in astrophysical systems ranging in scale from stellar interiors (e.g. Brüggén & Hillebrandt 2001) and protoplanetary discs (e.g. Johansen, Henning & Klahr 2006) to the evolution of the intergalactic medium (e.g. Nulsen 1982, 1986). Physically, the KH instability wraps up coherent sheets of vorticity into smaller, less organized structures. The small-scale motion then stretches and cascades to yet smaller scales. The instability therefore plays fundamental roles in fluid mixing and in the transition to turbulence.

*E-mail: dlecoanet@berkeley.edu

Because of its prevalence in nature and its physical significance, KH test problems are commonly used to evaluate the accuracy of different astrophysical hydrodynamics codes (e.g. Springel 2010; Hopkins 2015; Schaal et al. 2015): if a code can properly simulate the KH instability, it is presumed to capture mixing and turbulence in astrophysical simulations. Ideally, such an important test problem should stand against an analytic solution to ensure the veracity (not just reproducibility) of simulation results. Some analytic work addresses the KH instability with a sheet vortex model (Moore 1979), but only for incompressible fluid equations. For the compressible Navier–Stokes equations relevant to astrophysics, no analytic description of the non-linear KH instability currently exists.

Absent a non-linear analytic prediction, a resolved reference simulation provides the only reasonable approximation of the true solution. Comparing to a well-controlled and high-resolution benchmark gives a proxy for the true error of a given test. Robertson et al. (2010) and McNally, Lyra & Passy (2012) present careful studies of the early evolution of the KH instability. These authors also point out the numeric ill-posedness of contact-discontinuity simulations, in spite of existing analytical solutions in the linear and/or incompressible regimes. These works emphasize that converged non-linear simulations require well-resolved initial conditions (ICs). One limitation of these studies, however, is that Robertson et al. (2010) and McNally et al. (2012) only provide converged reference simulations for the linear (and possibly weakly non-linear) phase. In addition, converged non-linear solutions require solving dissipative equations. Many available astrophysical codes do not implement this essential feature. As a result, these works could only follow the instability for a few e-folding time-scales. Furthermore, although almost all astrophysical systems are three-dimensional, test problems are typical run in only two dimensions. Salvesen et al. (2014) performed KH convergence tests in both two and three dimensions. They found that simulations differed dramatically between resolutions in two dimensions, but found little resolution dependence in three-dimensional simulations. Porter, Pouquet & Woodward (1994) found similar results in three-dimensional turbulence simulations.

Not all works take the benchmark approach, however. In place of a non-linear reference solution, some authors use apparent small-scale structure as a proxy for the accuracy of their simulations (e.g. Springel 2010; Hopkins 2015). Presumably, more small-scale structure implies less numerical dissipation, and therefore greater accuracy. We find in the current paper that this intuition can, in some cases, lead to false conclusions. Mocz et al. (2015) show that small-scale structure is moving-mesh codes can be due to grid noise rather than physical effects, and describe methods to mitigate these errors. Some tests also abandon the smooth ICs of Robertson et al. (2010) and McNally et al. (2012), even though this choice precludes convergence of even the linear phase of the instability because the linear growth rates increase with wavenumber for an initially discontinuous velocity profile.

In this paper, we extend the work of McNally et al. (2012) by providing reference solutions for the *strongly non-linear* evolution of the KH instability. We use a smooth IC and explicit diffusion. We conduct simulations using both ATHENA (a Godunov code), and DEDALUS (a pseudo-spectral code that can solve the Navier–Stokes equations of compressible hydrodynamics) and find that both converge to the same reference solutions. We see agreement among different codes and different resolutions, with the validity of the reference solution limited only by (unavoidable) chaotic evolution at late times. We propose that future code tests include this KH

instability problem and compare to our validated, converged, reference solutions.

We organize the remainder of the paper as follows. Section 2 describes the equations, ICs, and codes used for our simulations. The results comprise two sections. In Section 3.1 we discuss the simpler simulations with constant initial density. Section 3.2 discusses the more complicated simulations with an initial density jump. Section 4 summarizes our results.

2 METHODS

2.1 Equations and initial conditions

We solve the hydrodynamic equations, including explicit terms for the diffusion of momentum and temperature:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0, \quad (1a)$$

$$\frac{\partial}{\partial t} (\rho \mathbf{u}) + \nabla \cdot (P \mathbf{I} + \rho \mathbf{u} \otimes \mathbf{u}) = -\nabla \cdot \mathbf{\Pi}, \quad (1b)$$

$$\frac{\partial E}{\partial t} + \nabla \cdot [(E + P) \mathbf{u}] = \nabla \cdot (\chi \rho \nabla T) - \nabla \cdot (\mathbf{u} \cdot \mathbf{\Pi}), \quad (1c)$$

along with the non-dimensionalized ideal gas equation of state, $P = \rho T$, with constant ratio of specific heats $\gamma = 5/3$. \mathbf{I} is the identity tensor, χ is the thermal diffusivity (with units $\text{cm}^2 \text{s}^{-1}$; $K = nk_b \chi$ is the thermal conductivity), and

$$\mathbf{\Pi} = -\nu \rho \left(\nabla \mathbf{u} + (\nabla \mathbf{u})^T - \frac{2}{3} \mathbf{I} \nabla \cdot \mathbf{u} \right) \quad (2)$$

is the viscous stress tensor with viscosity ν (with units $\text{cm}^2 \text{s}^{-1}$). We assume both ν and χ are constant.

We add a passive scalar to our simulations which we refer to as ‘dye’. The local fraction of dye particles c expresses dye concentration, and initially ranges from 0 to 1. The local conservation of dye is then

$$\frac{\partial}{\partial t} (\rho c) + \nabla \cdot (\rho c \mathbf{u}) = \rho \frac{dc}{dt} = -\nabla \cdot \mathbf{Q}_{\text{dye}}, \quad (3)$$

$$\mathbf{Q}_{\text{dye}} = -\rho \nu_{\text{dye}} \nabla c, \quad (4)$$

where d/dt represents the Lagrangian derivative, and ν_{dye} represents a diffusion coefficient for dye molecules (with units $\text{cm}^2 \text{s}^{-1}$). These equations conserve the total dye mass $\int \rho c dV$.

We define a dye entropy per unit mass $s \equiv -c \ln c$, along with its volume integral

$$S \equiv \int \rho s dV. \quad (5)$$

These evolve such that

$$\rho \frac{ds}{dt} - \nabla \cdot [(1 + \ln c) \mathbf{Q}_{\text{dye}}] = \rho \nu_{\text{dye}} \frac{|\nabla c|^2}{c} \quad (6)$$

$$\frac{dS}{dt} = \int \rho \nu_{\text{dye}} \frac{|\nabla c|^2}{c} dV \geq 0. \quad (7)$$

The second term on the left-hand side of equation (6) represents the entropy flux due to reversible diffusion of the dye. The right-hand side represents entropy generation due to non-reversible

dissipation.¹ The volume-integrated entropy S satisfies the following important properties.

(i) A fully unmixed fluid with $c = 0$ or 1 everywhere has zero entropy ($S = 0$).

(ii) A fully mixed fluid with $c^* = \int \rho c dV / \int \rho dV$ maximizes the entropy, $S_{\max} = -c^* \ln c^* \int \rho dV$.

(iii) S increases monotonically with time if $v_{\text{dye}} > 0$, and stays constant otherwise.

We restrict our attention to periodic simulations. This avoids potential difficulties with imposing Dirichlet and/or Neumann boundary conditions. Our ICs are

$$\rho = 1 + \frac{\Delta\rho}{\rho_0} \frac{1}{2} \left[\tanh\left(\frac{z-z_1}{a}\right) - \tanh\left(\frac{z-z_2}{a}\right) \right], \quad (8a)$$

$$u_x = u_{\text{flow}} \left[\tanh\left(\frac{z-z_1}{a}\right) - \tanh\left(\frac{z-z_2}{a}\right) - 1 \right], \quad (8b)$$

$$u_z = A \sin(2\pi x) \left[\exp\left(-\frac{(z-z_1)^2}{\sigma^2}\right) + \exp\left(-\frac{(z-z_2)^2}{\sigma^2}\right) \right], \quad (8c)$$

$$P = P_0, \quad (8d)$$

$$c = \frac{1}{2} \left[\tanh\left(\frac{z-z_2}{a}\right) - \tanh\left(\frac{z-z_1}{a}\right) + 2 \right], \quad (8e)$$

where $a = 0.05$ and $\sigma = 0.2$ are chosen so that the IC is resolved in all of our simulations. We take $u_{\text{flow}} = 1$ and $P_0 = 10$ so that the flow is subsonic with a Mach number $M \sim 0.25$ in regions with $\rho = 1$ and $M \sim 0.35$ in regions with $\rho = 2$. The size of the initial vertical velocity perturbation is $A = 0.01$. The ATHENA simulations are initialized with these functions evaluated at cell-centres even though ATHENA data represent cell-averaged quantities (see Appendix A for more discussion of this effect).

We adopt a rectangular domain with x in $[0, L)$, and z in $[0, 2L)$, with $L = 1$, and $z_1 = 0.5$, $z_2 = 1.5$, with periodic boundary conditions in both directions. The simulations have a horizontal resolution of N grid points (in ATHENA) or modes (in DEDALUS) in the x direction, and $2N$ grid points/modes in the z direction. Our IC has a reflect-and-shift symmetry: taking $z \rightarrow 2 - z$ and $x \rightarrow x + 1/2$ changes the sign of u_z but leaves the other quantities invariant. Thus, the simulations solve for the same flow twice. This is a requirement when using periodic boundary conditions, but also provides a test of whether or not the numerical simulations can preserve the symmetry. Almost all simulations presented here maintain the symmetry. We therefore only show the lower half of the domain. We calculate volume-averaged quantities like the dye entropy or the L_2 norm with respect to the entire domain.

In equation (8a), the free parameter $\Delta\rho/\rho_0$ represents the density jump across the interface. We study simulations with $\Delta\rho/\rho_0 = 0$ in Section 3.1 and with $\Delta\rho/\rho_0 = 1$ in Section 3.2. We refer to this change in density as a ‘jump’ throughout, although the transition is smooth, set by the tanh in equation (8a). The Reynolds number Re quantifies diffusion:

$$\nu = \chi = \nu_{\text{dye}} = \frac{L\Delta u}{Re}, \quad (9)$$

where $\Delta u = 2u_{\text{flow}}$ is the change in velocity. Note that we set the thermal diffusivity χ constant; consequently, the thermal conductivity $K \propto \rho$. Throughout the paper we measure time in units of L/u_{flow} , so $t = 1$ corresponds to approximately one turnover time. Equations (1)–(9) specify our system, with the free parameters $\Delta\rho/\rho_0$ and Re . In the following section we detail our methods for solving this system of equations.

2.2 Numerical methods

We study the KH instability using two open-source codes employing very different numerical methods: ATHENA and DEDALUS.

ATHENA² is a finite-volume Godunov code (Gardiner & Stone 2008; Stone et al. 2008). The scheme represents all field quantities with volume-averaged values in each grid element. A Riemann problem solves for fluxes between elements. We use third-order reconstruction with limiting in the characteristic variables to approximate field values at the element walls (Colella & Woodward 1984), the Harten–Lax–van Leer–Contact (HLLC) Riemann solver (Toro 2013), the Corner Transport Upwind (CTU) integrator (Colella 1990), and super-time-stepping for the diffusive terms (Alexiades, Amiez & Gremaud 1996; Choi, Kim & Wiita 2009). We used the ‘-O3’ compiler flag using Intel 14.0.1.106 and Mvapi2 2.0b on the Stampede supercomputer. We repeated some runs using second-order reconstruction and/or the Roe Riemann solver and/or stricter compiler flags (e.g. ‘-O2 -fp-model strict’) – these choices did not qualitatively affect the solutions. We use a static, uniform mesh, and a CFL safety factor of 0.8.

ATHENA is second-order accurate in both space and time. The leading-order grid-scale errors are diffusive. For most simulations reported here, we include explicit diffusion. A sufficiently large explicit diffusion can dominate grid-scale errors and allow the simulation to remain close to the true solution. However, higher order grid-scale errors can introduce non-diffusive effects, such as dispersion. If higher order errors project on to unstable modes, they can cause large differences in the solution, despite being higher order. The grid-scale errors in ATHENA respect the reflect-and-shift symmetry of our problem up to floating point accuracy, so even non-converged simulations can maintain the initial symmetry of the flow. In practice, we find all simulations maintain the initial symmetry, except simulations with $\Delta\rho/\rho_0 = 1$ without explicit diffusion. Since ATHENA’s algorithm manifestly preserves this symmetry, we expect the error results from chaotic amplification of floating-point errors.

DEDALUS³ is a pseudo-spectral code (Burns et al., in preparation). All field variables are represented as Fourier series, and the simulation solves for the evolution of the spectral-expansion coefficients in time. The code evaluates non-linear terms on a grid with a factor $3/2$ more points than Fourier coefficients, i.e. the $2/3$ de-aliasing rule. Lecoanet et al. (2014, appendix D.1) describe our implementation of the Navier–Stokes equations. Our implementation of the dye evolution equation is

$$\partial_t c - \nu_{\text{dye}} (\partial_x^2 c + \partial_z c_z) = -u \partial_x c - w c_z + \nu_{\text{dye}} (\partial_x \Upsilon' \partial_x c + \partial_z \Upsilon' c_z), \quad (10a)$$

$$c_z - \partial_z c = 0, \quad (10b)$$

¹ Equation (6) can be made to look like the analogous equation for heat conduction with the definition of a new ‘temperature’ $T_{\text{dye}} \equiv -\frac{1}{1+\ln c}$.

² ATHENA is available at <https://trac.princeton.edu/Athena/>

³ DEDALUS is available at <http://dedalus-project.org>

where we use the same notation as Lecoanet et al. (2014). For time-stepping, we use a third-order, four-stage DIRK/ERK method (RK443 of Ascher, Ruuth & Spiteri 1997) with a total CFL safety factor of 0.6 (i.e. 0.15 per stage). This formulation allows implicit time-stepping of sound waves. Thus, our time-step size only adjusts with the flow velocity, not the sound speed. The excellent agreement between the highest resolution DEDALUS and ATHENA simulations shows that high-wavenumber sound waves have negligible influence on the solution.

The pseudo-spectral method produces almost no numerical diffusion. Stability concerns require explicit diffusion in non-linear calculations. In marginally resolved simulations, discretization errors manifest as Gibbs’ ringing, which is prominently visible in snapshots. The numerical method does not explicitly preserve the reflect-and-shift symmetry – numerical errors can put power into the asymmetric modes. However, we find that in resolved simulations these asymmetric modes never grow to large amplitudes. Thus, maintaining this symmetry gives a test for a simulation’s fidelity.

3 RESULTS

This section describes the non-linear evolution of the KH instability, provides reference solutions, and compares the performance of DEDALUS and ATHENA. Section 3.1 considers unstratified simulations with constant initial density; both codes handle this problem easily. Section 3.2 concerns simulations with a density jump across the shear interface. This problem shows rich behaviour and poses significant numerical challenges. In both cases, a central vortex which wraps up material from above and below the initial interface into thin filaments.

3.1 Unstratified simulations ($\Delta\rho/\rho_0 = 0$)

In this section, we discuss simulations with constant initial density ($\Delta\rho/\rho_0 = 0$). Fig. 1 visualizes the flow with the dye concentration field of the lower half of the domain for simulations with explicit diffusion at different resolutions and Reynolds number, Re . The flow

consists of coherent filaments of unmixed fluid with dye concentration close to zero or one. The filaments twist around the central vortex until they become thin enough to diffuse away. The central vortex stays coherent in all simulations, and exhibits a more gradual dye concentration gradient than in the filaments. This reflects the smooth velocity and dye IC.

The snapshots show the state at $t = 6$. Strong non-linearity begins at $t \sim 2$, so this corresponds to at least four turnover times after the initial saturation of the instability. The simulations are labelled by the code used (A for ATHENA; D for DEDALUS) and their horizontal resolution (see Table 1 for a list of simulations with $\Delta\rho/\rho_0 = 0$).

3.1.1 $Re = 10^5$

Many of the simulations with the same Re but different resolution look similar by eye. To more quantitatively assess convergence, we calculate the L_2 norm of the differences between dye concentration fields in different simulations:

$$L_2(c_X - c_Y) = \left[\int dV (c_X - c_Y)^2 \right]^{1/2}, \quad (11)$$

where c_X and c_Y represent the dye concentration fields in two simulations, X and Y. The ATHENA and DEDALUS grids are different, so we use spectrally accurate techniques to interpolate DEDALUS solutions to the ATHENA grid for direct comparison (Appendix A). We argue in Appendix B that all simulations converge to our highest resolution DEDALUS simulations; thus, we assume these simulations are a good approximation to the ‘true’ solution.

Fig. 2 shows the L_2 norm of the difference between dye concentration fields of D2048 and other simulations with $Re = 10^5$. Because we believe D2048 closely represents the true solution (Appendix B), we call this the L_2 norm of the error. Solutions from both codes approach D2048 as resolution increases. At late times, A2048 and D1024 have roughly eight times smaller errors than A1024 and D512, respectively. That is, both codes exhibit third-order convergence. This indicates that interpolation produces the dominant error in ATHENA, which is the only third-order part of the algorithm. The DEDALUS simulations are spatially resolved, so

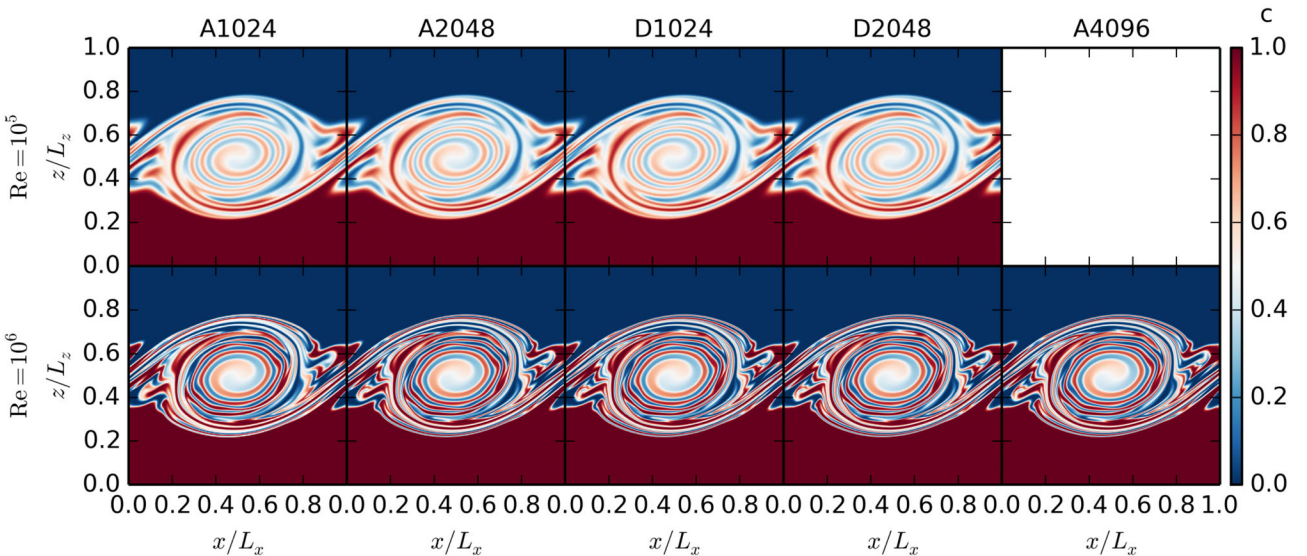


Figure 1. Snapshots of the dye concentration field in several simulations with $\Delta\rho/\rho_0 = 0$ at $t = 6$. The upper (lower) row shows simulations with $Re = 10^5$ (10^6). All the simulations with $Re = 10^5$ are well resolved. Small differences exist between the lower resolution ATHENA simulations at $Re = 10^6$ and the highest resolution ATHENA simulation and DEDALUS simulations [e.g. near $(x, z) = (0.9, 0.6)$, see Fig. 7].

Table 1. List of simulations with $\Delta\rho/\rho_0 = 0$. The character refers to the code used (D for DEDALUS, A for ATHENA, and N for ATHENA with no explicit diffusion), and the number is the horizontal resolution.

Re = 10^4	Re = 10^5	Re = 10^6	No explicit diffusion ^a
D512	D512	D512	N1024
A512	D512dt ^b	D1024	N2048
A1024	D1024	D2048	N4096
	D2048	A1024	
	A1024	A2048	
	A2048	A4096	

Notes. ^aRun in ATHENA.

^bRun with half the CFL safety factor.




Figure 2. L_2 norm of dye concentration errors for $\Delta\rho/\rho_0 = 0$ and $Re = 10^5$. We take D2048 as the ‘true’ solution (see Appendix B). Both DEDALUS and ATHENA exhibit third-order convergence. D512dt is run with half the time-step size as D512. Its error is similar to D1024, showing that the higher accuracy of D1024 is mostly due to a smaller time-step size rather than higher spatial resolution.

time-stepping produces the dominant error source in the DEDALUS simulations, which is also third order. We also plot errors from D512dt, which is run with a horizontal resolution of 512, but with half the CFL safety factor. D512dt is almost as accurate as D1024, showing that the higher accuracy of D1024 is mostly due to taking smaller time-steps. There are certain times (most notably near $t = 3.5$) where the flow develops smaller structures, and extra spatial resolution is required. The errors in quantities other than dye concentration (e.g. density) follow similar behaviour to that shown in Fig. 2.

We calculate the volume-integrated dye entropy for each simulation (equation 5). Fig. 3 plots the entropy as a function of time. Because all simulations are well resolved, there are no visible differences in the entropy between the different simulations.

3.1.2 $Re = 10^6$

The unmixed filaments are much thinner for $Re = 10^6$ than for $Re = 10^5$, challenging the codes. Unlike the $Re = 10^5$ case, some minor visible differences appear between the solutions for $Re = 10^6$. The lower resolution simulations do not fully resolve the flow (one such feature is highlighted in Fig. 7).

To assess convergence, we again plot the L_2 norm of the error in dye concentration with respect to D2048 (Fig. 4). A1024 has the largest errors of any simulation. At late times, the errors interact non-linearly, whereas the errors in the higher resolution ATHENA




Figure 3. Volume-integrated dye entropy (equation 5) as a function of time for the four simulations with $Re = 10^5$ shown in Fig. 1. All simulations are well resolved, so the dye entropies are almost equal.




Figure 4. L_2 norm of dye concentration errors for $\Delta\rho/\rho_0 = 0$ and $Re = 10^6$. A1024 is not well resolved so its errors follow a different pattern than the other ATHENA simulations. The errors in A4096 are smaller than the errors in A2048 by ≈ 6 . The errors in D1024 are smaller than the errors in D512 by about 100. This demonstrates the fast (exponential) convergence of spectral methods.

simulations stay linear and the temporal variation of the error is the same independent of the magnitude of the error. The ratio of errors of the two higher resolution ATHENA simulations is about 6 – in between second- and third-order convergence. This suggests that the size of interpolation errors roughly match the size of other errors in the code (e.g. from the Riemann problem or time-stepping).

The difference in errors between D512 and D1024 is about 100 – much larger than the difference in errors between the ATHENA simulations. D512 (not shown in Fig. 1) underresolves the flow and includes some low-amplitude Gibbs’ ringing. Increasing the resolution from 512 to 1024 eliminates spatial errors because of the exponential convergence of spectral methods. This allows for very large error reduction with only modest resolution changes. The exponential nature of spectral methods makes convergence practically binary: simulations with Gibbs’ ringing are not converged; simulations without Gibbs’ ringing very likely are converged.

We plot volume-integrated dye entropy for $Re = 10^6$ in Fig. 5. Like for $Re = 10^5$, all well-resolved simulations produce similar entropy. However, the underresolved A1024 produces slightly more entropy. This agrees with the heuristic that extra numerical diffusion leads to excess entropy generation.




Figure 5. Volume-integrated dye entropy (equation 5) as a function of time for the five simulations with $\text{Re} = 10^6$ shown in Fig. 1. The entropy of all simulations is very similar except for A1024; this is another indication that A1024 is not well resolved.




Figure 6. Volume-integrated dye entropy (see Section 2.2) as a function of time with $\Delta\rho/\rho_0 = 0$, for three resolved simulations with different Re , as well as three ATHENA simulations with no explicit diffusion (dashed lines; labelled with N, for no explicit diffusion, and their horizontal resolution). The entropy of N1024 and the simulation with $\text{Re} = 10^6$ are very similar. Their flow fields show minor differences (see Fig. 7). Note that the entropy decreases with increasing resolution in the simulations without explicit diffusion. This is not the case in simulations with an initial density jump (see Fig. 16).

3.1.3 An effective Reynolds number?

We now describe ATHENA simulations without any explicit diffusion. An important question is, does the numerical diffusion in ATHENA act like an explicit diffusion? Put another way, does ATHENA have an effective Reynolds number at a given resolution for this problem? As we describe below and in Section 3.2, the answer to this question is very problem dependent.

To test this, we plot the converged volume-integrated dye entropy for several Reynolds numbers, along with the volume-integrated dye entropy for ATHENA simulations without explicit diffusion (Fig. 6); simulations run without explicit diffusion are labelled with N. The entropy evolution of N1024 is similar to the entropy evolution for $\text{Re} = 10^6$. This might lead one to think that the effective Reynolds number of this ATHENA simulation is about 10^6 .

However, a closer investigation shows that N1024 and the $\text{Re} = 10^6$ simulation have different dye concentration fields which, by chance, result in similar volume-integrated entropies (Fig. 7). Instead, the dye concentration field of N1024 looks like the dye concentration field of the (underresolved) A1024 simulation with




Figure 7. Snapshots of the dye concentration field between $0.89 < x < 0.95$ and $0.55 < z < 0.61$, at $t = 6$ for $\Delta\rho/\rho_0 = 0$. All simulations use ATHENA, either with $\text{Re} = 10^6$ (left-hand column) or no explicit diffusion (right-hand column). The three rows have different resolutions. This zoom-in of Fig. 1 highlights the differences between simulations at different resolutions – however, for the most part, the simulations look very similar. A2048 and A4096 represent resolved simulations with $\text{Re} = 10^6$. Although the entropies for N1024 (upper right-hand plot) and A4096 (lower left-hand plot) track each other (Fig. 6), the dye concentration fields exhibit minor differences.

$\text{Re} = 10^6$. Fig. 5 shows A1024 has a higher entropy than the true $\text{Re} = 10^6$ solution. By removing the explicit diffusion, the flow evolution remains similar to A1024 (and different from the resolved $\text{Re} = 10^6$ solution), but the interfaces between filaments are sharper, which decreases the entropy. The effects of having the incorrect flow field (increasing entropy), but sharper interfaces between filaments (decreasing entropy) happen to cancel out, so the entropy of N1024 is similar to that of $\text{Re} = 10^6$.

Although we have highlighted the differences between N1024 and the converged solutions with $\text{Re} = 10^6$, it is worth reiterating that the two solutions are in fact remarkably similar. This shows that N1024 roughly has an effective Reynolds number of 10^6 . In detail, however, the remaining modest differences between N1024 and the $\text{Re} = 10^6$ solution demonstrate that the numerical dissipation in ATHENA is not exactly equivalent to physical dissipation via viscosity and thermal conduction.

One difficulty with the notion of an effective Reynolds number is that it is extremely problem dependent, even at fixed resolution. In the next section, we introduce a small (by astrophysical standards) density jump into the IC. This completely changes the problem by introducing secondary instabilities which enhance mixing, producing very clear differences between resolved simulations and ATHENA simulations without explicit diffusion (Fig. 15). For the constant-density problem described here, omitting diffusion produces less entropy. Including a density jump reverses this trend:

Table 2. List of simulations with $\Delta\rho/\rho_0 = 1$. The character refers to the code used (D for DEDALUS, A for ATHENA, and N for ATHENA with no explicit diffusion), and the number is the horizontal resolution.

Re = 10^5		No explicit diffusion
DEDALUS	ATHENA	ATHENA
D2048	A1024	N512
D2048r ^a	A2048	N1024
D2048p ^b	A4096	N2048
D3072	A8192	N4096
D4096	A16384	

Notes. ^aRestarted from A2048 at $t = 3.2$.

^bDifferent initial perturbation.

simulations with only numerical diffusion undergo *more* mixing than simulations with explicit diffusion. Although assigning an effective Reynolds number to ATHENA simulations without explicit diffusion may be reasonably accurate for the constant-initial-density problem, this does not carry over to the problem with an initial density jump.

3.2 Simulations with a density jump ($\Delta\rho/\rho_0 = 1$)

Both the qualitative features of the flow and the convergence properties of the simulations change dramatically once we introduce an initial density jump ($\Delta\rho/\rho_0 \neq 0$). Unlike the unstratified case, secondary instabilities of the filaments produce small-scale structures in the flow. These secondary instabilities, and the resulting small-scale features, depend on the resolution and the code used. As a result, simulations with a non-zero density jump require far more computational resources than the unstratified simulations presented in the previous section. We limit the simulations with explicit diffusion to Re = 10^5 (Table 2) – our finite-computing budget precludes solutions for Re = 10^6 . The largest simulations required roughly 10^6 core hours.

Fig. 8 shows the dye concentration for different simulations at different times. In both DEDALUS simulations, and the highest resolution ATHENA simulation, the outer filaments (i.e. those outside the central vortex) become unstable to a sausage-like mode (see the panel in Fig. 10 for an example). Lower resolution ATHENA simulations also undergo a separate instability of the inner filaments of the vortex. We refer to these two instabilities at the outer-filament instability (OFI) and the inner-vortex instability (IVI; see Fig. 10

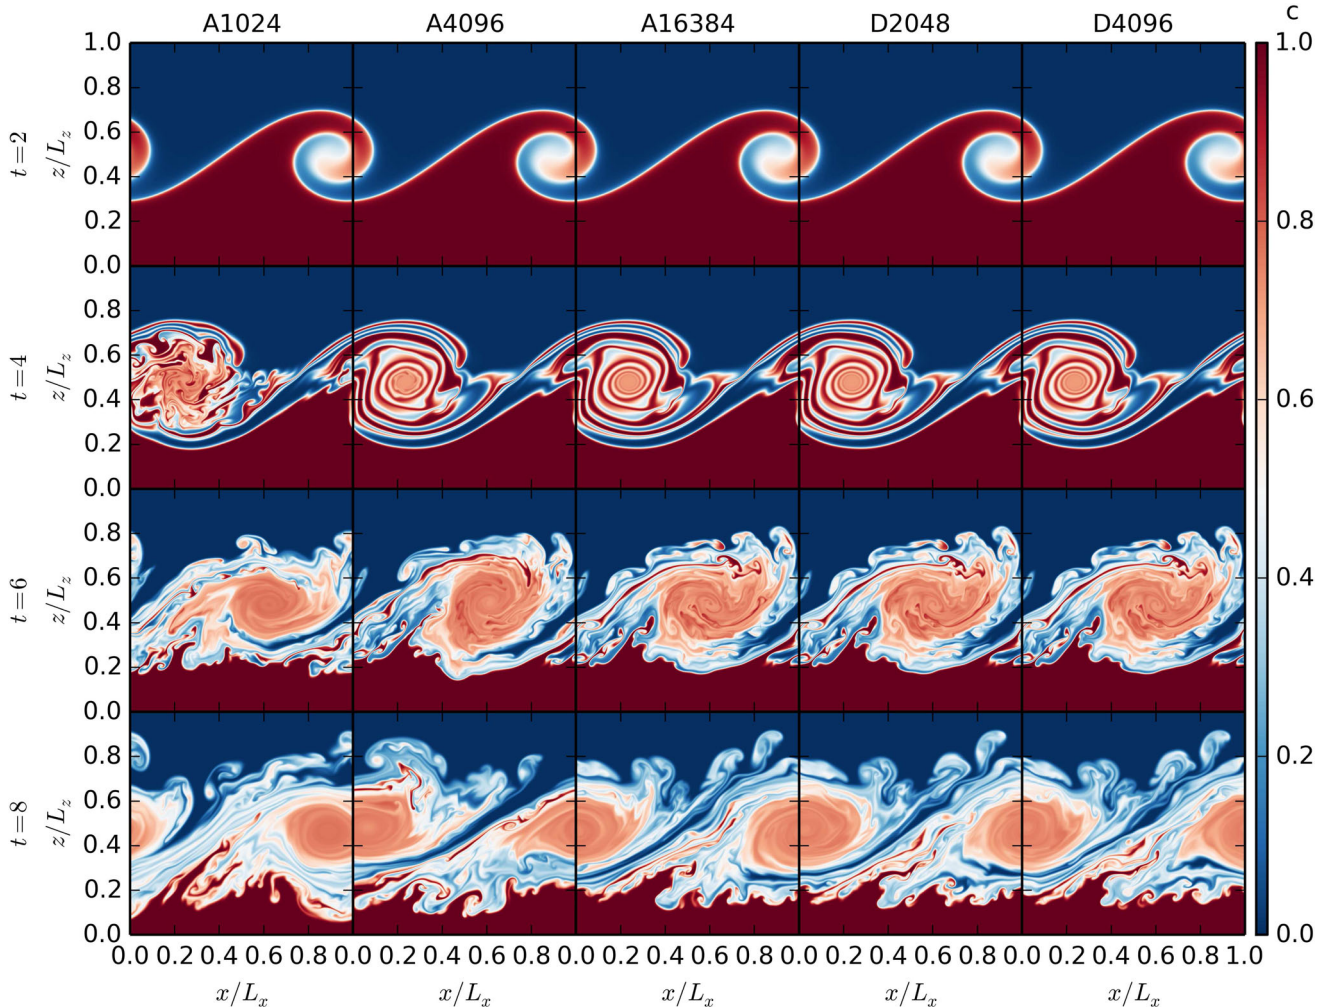


Figure 8. Snapshots of the dye concentration field in several simulations with $\Delta\rho/\rho_0 = 1$ and Re = 10^5 . Each row corresponds to a different time. The low-resolution ATHENA simulations suffer from a secondary instability (seen at $t = 4$) in the middle of the vortex, which is not present in the DEDALUS simulations nor in the A16384. This causes substantial differences at later times. A16384 and both DEDALUS simulations stay very similar at late times, although small differences develop from chaos (see Section 3.2.2).




Figure 9. L_2 norm of dye concentration errors for $\Delta\rho/\rho_0 = 1$ and $\text{Re} = 10^5$. D3072 and D4096 are the closest pair of simulations, suggesting that D4096 is a good approximation to the true solution. All ATHENA simulations except A16384 diverge away from D4096 exponentially with a rate of 8, suggesting the growth rate of the inner vortex instability (IVI) (see Fig. 10) is also 8. The errors in the lower resolution DEDALUS simulation and A16384 grow exponentially with a rate of about 2–3. We interpret this divergence as due to chaos (see Section 3.2.2). D3072 has errors smaller than D2048 by ≈ 4 , consistent with third-order convergence set by our choice of time-stepping algorithm.

for examples). These instabilities are similar to the baroclinic secondary instabilities discussed in Reinaud, Joly & Chassaing (2000) and Fontane & Joly (2008). The competition between these two instabilities plays a crucial role in the evolution of the system.

We plot the L_2 norm of the error in dye concentration with respect to D4096 in Fig. 9. As described in Appendix B, we believe D4096 approximates the true solution. The difference between D3072 and D4096 is smaller than the differences between any other pair of simulations. At later times, even the errors between D3072 and

D4096 become large. In Section 3.2.2 we attribute this late-time behaviour to chaos.

Fig. 9 shows that at early times, the low-resolution ATHENA simulations diverge exponentially from D4096 with an inferred growth rate of about 8. The IVI produces this divergence. Furthermore, the four ATHENA simulations with resolutions between 1024 and 8192 are all equally spaced horizontally in Fig. 9. The horizontal axis spacing is $\log 2/2$ time units. This suggests that the same instability exists independent of resolution, but the amplitude of the perturbation that seeds the instability drops by 16 when the resolution doubles. Though numerical errors seed the growth, the constant growth rate of the IVI suggests it is a physical instability (we demonstrate this in Section 3.2.1).

The IVI is a robust feature of low-resolution ATHENA simulations. Using the Roe integrator, second-order reconstruction, or shifting the IC by half a grid point does not affect the development of this instability (as confirmed using the L_2 error), but can cause visible differences in the flow evolution. This demonstrates that grid-scale errors drive the IVI. Using first-order reconstruction suppresses the IVI, but the enhanced numerical diffusion causes large errors. We have also tried adding low-amplitude (up to 10^{-4}) white noise to the initial density or pressure. These do not cause any visible changes to the IVI. The flow forgets some of the detailed information of its IC (see Section 3.2.3).

The highest resolution ATHENA simulation (A16384) does not develop the IVI. This demonstrates that the IC is in fact stable to the IVI; the problem is well posed. Rather, numerical errors seed the IVI at some later time, during the evolution of the flow. Although some numerical errors are still inevitably present, A16384 does not develop the IVI because the ‘base state’ of spiralling filaments of unmixed fluid also succumbs to the OFI. In this case, the OFI disrupts the inner vortex before the IVI grows to large amplitudes (see Fig. 10).

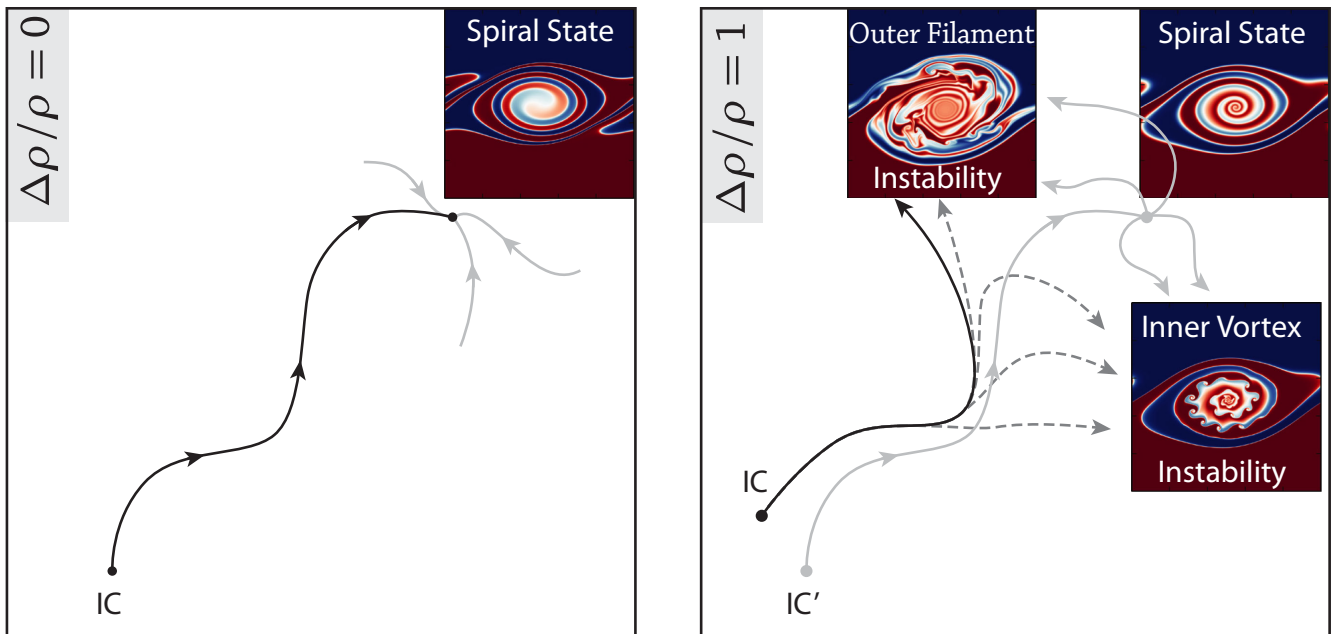


Figure 10. Schematic phase–space diagram for $\Delta\rho/\rho_0 = 0$ (left) and $\Delta\rho/\rho_0 = 1$ (right). For constant initial density, the system has a stable state with ever-narrowing spiral filaments. We hypothesize that there is an initial condition (IC; right-hand panel) leading to a similar spiral state for $\Delta\rho/\rho_0 = 1$. But this state is now unstable to the outer filament instability (OFI) and inner vortex instability (IVI). Our chosen IC’s trajectory (solid black line) approaches the spiral state, but becomes unstable to the OFI. Errors introduced by the numerical hydrodynamics may cause deviations in the trajectory leading to the IVI (dashed grey lines).




Figure 11. Volume-integrated dye entropy (equation 5) as a function of time for simulations with $\Delta\rho/\rho_0 = 1$ and $\text{Re} = 10^5$. The top panel plots the entropy, and the bottom panel plots the entropy deviation from D4096. The entropy of all the simulations diverge from D4096, but the less-accurate simulations diverge faster. For each ATHENA simulation, the entropy initially increases faster than D4096 when it starts to diverge. At later times, the entropy sometimes drops below the entropy of D4096.

The absence of the IVI is a robust feature of our DEDALUS simulations. We confirmed the stability of the base state by re-running D2048 with low-amplitude white noise added to the IC; we also re-initialized D2048 from the ATHENA IC. This introduces small but non-random grid representation differences (Section 3.2.1). In both cases, we recover the same evolution. However, we can trigger the IVI in DEDALUS with a large (~ 10 per cent by energy) perturbation to the IC (Section 3.2.3).

Fig. 10 summarizes the relation between the two secondary instabilities in this problem. For a constant initial density (left-hand panel), the system evolves towards a stable state characterized by spiralling filaments. Small differences in ICs, integration algorithms, presence of dissipation, etc., cause only minor changes in the evolution. We hypothesize that a similar spiral state also exists for $\Delta\rho/\rho_0 = 1$, and that it could be reached from some IC. However, our simulations demonstrate that the spiral state is now unstable. Thus, small errors lead to the large differences in evolution.

Small perturbations to the hypothetical IC of Fig. 10 would lead to trajectories that either develop the OFI or the IVI. However, our chosen IC is squarely in the attracting basin of the OFI. Thus, infinitesimal perturbations to IC will still lead to the OFI. Errors introduced by numerical hydrodynamics cause the codes to not follow the correct trajectory (solid black line). Certain types of errors can cause trajectories to diverge from the correct solution, sometimes towards the IVI (dashed grey lines). Alternatively, sufficiently large initial perturbations can also knock the system into the attracting basin of the IVI (Section 3.2.3).

We note that the phase-space for this problem is very high dimension, and that the OFI and IVI represent two (likely non-parallel) unstable directions of the spiral state’s stable manifold. Thus, both instabilities can act simultaneously, which sometimes occurs in simulations.

Fig. 11 shows the volume-integrated dye entropy of the simulations shown in Fig. 8. The entropy follows a similar evolution in every simulation. To visualize the small deviations, the bottom




Figure 12. Plots of dye concentration (c), mass density (ρ), the divergence of the velocity ($\nabla\cdot\mathbf{u}$), and the vorticity ($\omega = \mathbf{e}_z \cdot \nabla \times \mathbf{u}$) in D4096 with $\Delta\rho/\rho_0 = 1$, $\text{Re} = 10^5$ at $t = 6$. The divergence of the velocity and the vorticity are measured in units of u_{flow}/L_x . The dye concentration and mass density fields are almost inverses of each other. The divergence of the velocity is largest at the interfaces between filaments, whereas the vorticity shows the location of vortices.

panel shows the entropy with reference solution D4096 subtracted off. All the simulations diverge from D4096, but more accurate simulation diverge later, with D2048 and A16384 developing small differences later than any other simulation. The relation between entropy and resolution is more complicated for $\Delta\rho/\rho_0 = 1$ than for $\Delta\rho/\rho_0 = 0$ (Figs 3 and 5).

Apart from the dye concentration field, many of the other flow quantities follow similar patterns. Fig. 12 shows several quantities from D4096 at $t = 6$. The mass density is almost the inverse of the dye concentration. This indicates that compression is not an important part of the large-scale dynamics. Lacking mass diffusion, the density shows sharper gradients than the concentration field. Temperature diffusion and rapid sound waves regularize the density evolution. These effects limit large temperature gradients, and keep the flow in local pressure equilibrium.

The velocity divergence field is characterized by a large-scale quadrupole centred at the vortex, and large amplitude, small-scale features near the boundaries of filaments. The most prominent feature of the vorticity field is the central vortex, which is a remnant of the initial shear. Small-scale vortex sheets and filaments perhaps result from the incomplete roll-up of the IC due to secondary instabilities.

Throughout this paper, we compare different solutions by calculating the L_2 norm of the difference between dye concentration fields. We have made similar comparisons between simulations with $\text{Re} = 10^5$ and $\Delta\rho/\rho_0 = 1$ using the L_1 norm of the difference between dye concentration fields, and using the L_2 norm of the




Figure 13. Snapshots of dye concentration field for $\text{Re} = 10^5$ and $\Delta\rho/\rho_0 = 1$. D2048r is a DEDALUS simulation restarted with the A2048 output at $t = 3.2$. At this time, the inner vortex instability (IVI) is still in the linear phase, so there are no visible differences between the three simulations. At $t = 4$, the IVI is very non-linear, producing large differences between D2048 and A2048. This instability also takes place in D2048r, and the dye concentration fields of A2048 and D2048r are nearly identical. This demonstrates that the IVI is physical, but is seeded by errors in the lower resolution ATHENA simulations that are not present in the DEDALUS simulations or the highest resolution ATHENA simulations.

difference between the three other fields shown in Fig. 12. We find the results to be qualitatively similar in all cases. This is expected given the similarity between the fields.

3.2.1 Inner-vortex instability

To determine the origin (physical versus numerical) of IVI, we initialize a DEDALUS simulation with horizontal resolution 2048 with the output from A2048 at $t = 3.2$. We call this simulation D2048r. Fig. 9 shows that A2048 is still in the linear phase of the IVI at this time. In Fig. 13, we plot the dye concentration field at $t = 3.2$ and 4 for D2048, A2048, and D2048r. At $t = 3.2$, the simulations all look the same. However, the instability becomes non-linear by $t = 4$, producing large changes in the dye concentration field. D2048 shows no signs of the IVI. However, D2048r looks almost identical to A2048. The L_2 norm of the difference of dye concentration fields between D2048r and D4096 almost exactly follows the norm of the difference between A2048 and D4096.

This shows that the IVI is a physical instability of this system. It is not seen in the DEDALUS simulations or the highest resolution ATHENA simulation because the IC does not project sufficiently on to its unstable modes. Errors in low-resolution ATHENA simulations incorrectly excite perturbations unstable to the IVI. DEDALUS simulations, and the highest resolution ATHENA simulation, suppress noise well enough the instability never becomes non-linear. In our phase-space diagram (Fig. 10), the lower resolution ATHENA simulations do not properly follow the black line, and instead meander to the right, becoming unstable to the IVI. D2048r is initialized to the right of IC, so it develops the IVI just like A2048.

As a final test, we started a DEDALUS simulation from the output of an ATHENA simulation at $t = 0$. This tests whether dynamical evolution causes the IVI, rather than small differences between the implementation of the ICs. Although this introduced root-mean-squared differences in the horizontal velocity of $\approx 4 \times 10^{-4}$ at $t = 0$, the DEDALUS simulation did not develop the IVI.

3.2.2 Chaos

At around $t \approx 4$, D2048, D3072, and A16384 start to diverge exponentially from D4096 (Fig. 9). The differences increase with a growth rate of about 2–3, much lower than the growth rate of 8 of the IVI found in the lower resolution ATHENA simulations. We interpret the differences between the simulations as due to chaos. The faster divergence discussed in Section 3.2.1 is inconsistent with chaos since it is resolution dependent and only seen in low-resolution ATHENA simulations.

A system is chaotic if small differences between ICs grow exponentially in time. To confirm the system is chaotic, we calculate a ‘local-in-time’ Lyapunov exponent (i.e. growth rate). We pick a time and simulation, and look for linearly unstable perturbations. This requires solving an eigenvalue problem. The largest unstable eigenvalue is the Lyapunov exponent. Appendix C details this procedure.

This calculation does not include base-state time evolution (i.e. we consider a ‘local-in-time’ calculation). The most unstable eigenvector at a time t_0 might differ significantly from the most unstable eigenvector at a nearby time $t_0 + \Delta t$. Then it would be impossible for perturbations to grow at the Lyapunov exponent over times $\sim \Delta t$. We interpret our ‘local-in-time’ Lyapunov exponents as an upper bound on the growth rate of perturbations due to chaos (up to logarithmic corrections), and as a heuristic measure of the strength of chaos in this problem.

We calculated the Lyapunov exponent for D2048 with $\text{Re} = 10^5$ and $\Delta\rho/\rho_0 = 1$ at two times, $t = 2.5$ and 4.5. We find Lyapunov exponents of $\lambda_{t=2.5} \approx 2.1$, and $\lambda_{t=4.5} \approx 3.7$. Thus, the exponential growth of differences between either D2048, D3072, or A16384, D4096 is consistent with chaos. However, the growth rate of the differences between the lower resolution ATHENA simulations and D4096 is much larger than the Lyapunov exponent. These differences are inconsistent with chaos, instead being due to the IVI (Section 3.2.1).

The simulations with $\Delta\rho/\rho_0 = 0$ do not appear to diverge from one another in the same way. The highest resolution DEDALUS simulations converge at late times. We also calculate the Lyapunov exponent for D1024 with $\text{Re} = 10^6$ and $\Delta\rho/\rho_0 = 0$ at $t = 6$. We find $\lambda_{t=6} \approx 0.4$. Although this seems inconsistent with our finding that the DEDALUS simulations approach each other with time, recall that this ‘local-in-time’ calculation gives an upper bound on the growth rate due to chaos (up to logarithmic corrections). Because the turnover time is 1, a Lyapunov exponent less than 1 suggests that small perturbations cannot grow before the background state changes substantially. To show definitively that the $\Delta\rho/\rho_0 = 0$ solution is not chaotic, one should maximize the amplification of an initial perturbation over several turnover times, e.g. between $t = 6$ and 9 (for instance, using the adjoint method; e.g. Kerswell, Pringle & Willis 2014).

3.2.3 Initial condition

Although our chosen IC does not lead to the IVI for converged simulations, one might wonder if other ICs do lead to this instability. We performed several DEDALUS simulations that add low-amplitude white noise to the IC (e.g. see Section 3.2.1). None of these simulations develops the IVI.

We now consider a simulation in which we include perturbations to the IC with order unity amplitude and large wavelengths. Equations (8) still hold for all quantities except the vertical velocity,

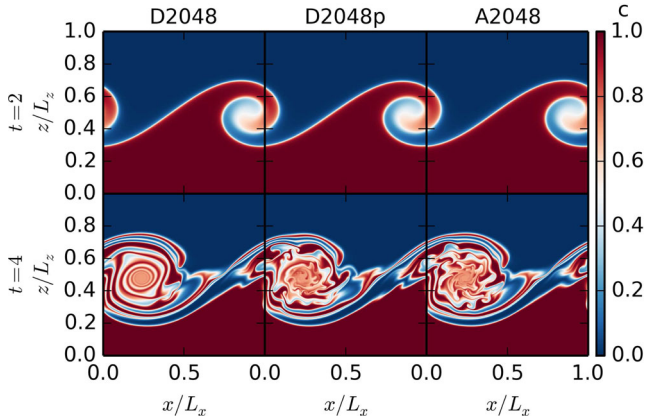


Figure 14. Snapshots of dye concentration field for $\text{Re} = 10^5$ and $\Delta\rho/\rho_0 = 1$. D2048p is a DEDALUS simulation with an initial vertical velocity that includes power over a range of Fourier modes (equation 12), in contrast to the single mode ICs focused on throughout the rest of this paper. At $t = 2$ all solutions look the same, indicating that longest wavelength mode has the largest growth rate. At $t = 4$, D2048p has developed the inner vortex instability (IVI), as well as other deviations from the DEDALUS and ATHENA simulations away from the vortex.

which we now take to be

$$u_z = A (\sin(2\pi x) + f(x)) \times \left[\exp\left(-\frac{(z-z_1)^2}{\sigma^2}\right) + \exp\left(-\frac{(z-z_2)^2}{\sigma^2}\right) \right], \quad (12)$$

where $f(x)$ includes Fourier modes 2–10. Each mode receives a random phase and random amplitude uniformly distributed between -0.05 and 0.05 . Thus, $f(x)$ represents about a 10 per cent perturbation to the single sine mode IC.

Fig. 14 shows snapshots of the dye concentration field for this simulation, denoted D2048p, along with D2048 and A2048 for comparison. At $t = 2$, all three simulations look identical. This indicates that the lowest wavenumber Fourier mode grows faster than the other modes included in our IC.

By $t = 4$ the perturbations from the other Fourier modes produce significant changes to the dye concentration field. D2048p now displays the IVI. In addition, large differences appear away from the vortex, where the DEDALUS and ATHENA simulations look almost identical. Because the new IC does not respect the shift-and-reflect symmetry of the problem, the two half domains have different features (we only show the bottom half).

3.2.4 Simulations without explicit diffusion

Lastly, Fig. 15 compares the resolved simulations at $\text{Re} = 10^5$ with an ATHENA simulation with horizontal resolution 4096 without explicit diffusion (N4096). The simulation without explicit diffusion exhibits many secondary instabilities early in the evolution (between $t = 2$ and 4). Unlike the lower resolution simulations at $\text{Re} = 10^5$, the secondary instability is not limited to the IVI. Instead, instabilities grow throughout the domain at locations of strong shear.

These instabilities shred apart the vortex, leading to vigorous mixing. Fig. 16 compares the volume-integrated dye entropy of ATHENA simulations with no explicit diffusion at different resolutions with D4096. Simulations without explicit diffusion produce almost no entropy until $t \approx 3.5$. At this time, the secondary instabilities start to cause diffusion at the grid-scale. This generates entropy more rapidly than the explicit diffusion of D4096 (or any of the




Figure 15. Snapshots of dye concentration field for $\Delta\rho/\rho_0 = 1$. N4096 is an ATHENA simulation with no explicit diffusion. For comparison, we also plot D4096 ($\text{Re} = 10^5$). Secondary instabilities occur very early at many locations in N4096. By $t = 6$, the simulation has broken its initial symmetry (we only plot the bottom half). The secondary instabilities produce significant mixing, leading to greater entropy generation than in simulations with explicit diffusion (Fig. 16).

other simulations with explicit diffusion). For $t > 5$, the entropy of the simulations without explicit diffusion is larger than the entropy of D4096. Paradoxically, the entropy increases as the resolution increases. Our expectation is that the entropy generation should decrease as Re increases. However, we do not have any resolved simulations with higher Re for comparison, so we cannot present evidence that this additional mixing is spurious. But this problem shows that introducing an explicit diffusion in ATHENA can *decrease* the diffusion in the simulation.

4 CONCLUSION

This paper describes several converged, non-linear solutions to the KH problem. By using a smooth IC and explicit diffusion, we demonstrate that solutions remain virtually identical (for constant initial density) or very similar (for an initial density jump of one) with resolution above a certain threshold. This permits a well-defined reference solution for this problem, against which errors can be accurately estimated. We verify this using two codes, DEDALUS and ATHENA, with very different numerical methods (pseudo-spectral




Figure 16. Volume-integrated dye entropy (equation 5) as a function of time for simulations with $\Delta\rho/\rho_0 = 1$. D4096 is run at $\text{Re} = 10^5$, and all simulations labelled with N are run with ATHENA with no explicit diffusion. At early times, the highest resolution runs without explicit diffusion have the lowest entropy. However, at around $t = 5$, the lower resolution runs without explicit diffusion have lower entropy. D4096 has the lowest entropy at late times. This indicates that simulations without explicit diffusion have *greater* numerical mixing compared to simulations with explicit diffusion. This becomes more prominent as the resolution increases. By contrast, in the simulations without an initial density jump, explicit diffusion leads to more mixing, and for simulations without explicit diffusion, increasing resolution decreases mixing (Fig. 6).

and Godunov, respectively). Previous KH test problems either did not use smooth ICs, or did not include explicit diffusion. Absent these two choices, the KH problem cannot be quantitatively compared between codes because the solutions depend sensitively on grid-scale errors and do not converge with increasing resolution.

We first study simulations with a constant initial density (Section 3.1). We find converged solutions to this relatively easy problem with Reynolds numbers (Re) as high as 10^6 . The solution is characterized by the continual roll-up of the initial vortex sheet, producing alternating filaments of unmixed material (Fig. 1). We find third-order convergence in both DEDALUS & ATHENA for simulations with $\text{Re} = 10^5$ (Fig. 2), and better than second-order convergence in both codes for simulations with $\text{Re} = 10^6$ (Fig. 4).

To quantify mixing in the simulations, we calculate the volume-integrated dye entropy as a function of time for several Reynolds numbers, as well as for ATHENA simulations without explicit diffusion (Fig. 6). As the Reynolds number increases, the entropy generation decreases monotonically. Similarly, as the resolution of ATHENA simulations without explicit diffusion increases, the entropy generation also decreases monotonically. The entropy of one ATHENA simulation without explicit diffusion is very close to the entropy of the $\text{Re} = 10^6$ simulation, although the solutions show minor differences (Fig. 7). These small differences indicate that the numerical diffusion in ATHENA does not act precisely as a physical diffusion from viscosity and/or thermal conductivity. For certain applications, however, assigning an effective Reynolds number to ideal fluid simulations may suffice. This does not appear to be the case for KH simulations with density jumps, as we now discuss.

Including an initial density gradient aligned with the velocity gradient makes the problem much richer (Section 3.2). The rolled-up vortex-sheet filaments become unstable in at least two ways: the inner vortex instability (IVI) and/or the outer filament instability (OFI) (Figs 8 and 10). The DEDALUS simulations and highest resolution ATHENA simulation only exhibit the OFI, whereas the lower resolution ATHENA simulations also exhibit the IVI. Adding small amplitude noise to the IC does not produce the IVI in DEDALUS,

demonstrating that our chosen IC is not susceptible to this instability; instead, numerical errors seed the IVI throughout the evolution of the ATHENA simulations. It is not surprising that DEDALUS is more accurate than ATHENA for this smooth flow – the Godunov method is designed for simulating flows with shocks. However, it is not well appreciated that the pseudo-spectral method is able to solve the full Navier–Stokes equations with Mach number order unity.

We use the L_2 norm to quantify the difference between dye concentration fields of different simulations, and find the IVI grows at a rate of ≈ 8 , independent of resolution (Fig. 9). Furthermore, a DEDALUS simulation initialized with an ATHENA state in the linear phase of the IVI develops the instability in the same way as ATHENA (Fig. 13), demonstrating the physical, rather than numerical, nature of the instability.

Adding a large (~ 10 per cent by energy) perturbation with multiple Fourier modes to the initial velocity in DEDALUS can seed the IVI (Section 3.2.3). Although this suggests that the IVI is possibly generic for KH instabilities in astrophysics, we believe the single-mode IC discussed throughout the rest of this paper is still particularly valuable for a test problem. Because small numerical errors can produce large differences in the solution, one can assess by eye the fidelity with which a code is solving the fluid equations. This KH test problem is difficult, which we believe makes it interesting. In contrast, an unresolved KH problem is not a good test of fluid codes, because noise due to numerical errors can masquerade as higher fidelity solutions.

The DEDALUS simulations and highest resolution ATHENA simulation also diverge from each other exponentially at late times, but with a much smaller growth rate ≈ 2 –3. In Section 3.2.2 we calculate the maximum Lyapunov exponent of the flow, and argue that chaos drives the divergence. The Lyapunov exponent represents the maximum possible rate of divergence of solutions due to chaos (up to logarithmic corrections). At late times when the DEDALUS simulations and highest resolution ATHENA simulation begin to diverge, the Lyapunov exponent is ≈ 3.7 , so the divergence we see is consistent with chaos. Because the system is chaotic, our solutions are not as accurate as the solutions with constant initial density. We still find power-law convergence in the DEDALUS simulations at fixed time (Fig. 9). However, the amount of time that a solution maintains a fixed level of accuracy increases only logarithmically with resolution.

For the IC with a density jump, we also compare a high-resolution ATHENA simulation without explicit diffusion to our converged (within the limits of chaos) simulations with $\text{Re} = 10^5$. Secondary instabilities pervade the simulation without explicit diffusion (Fig. 15). The secondary instabilities cause enhanced mixing, and at late times, the simulations without explicit diffusion have higher entropy than the $\text{Re} = 10^5$ simulation (Fig. 11). Introducing explicit diffusion into ATHENA can *reduce* the diffusion in the simulation. For this reason, we hypothesize (but cannot prove) that this small-scale structure is likely unphysical, and would not develop for any reasonable IC or Reynolds number. This highlights that a solution with more small-scale structure is not necessarily better.

Another possible comparison is to experiments. Unfortunately, the problems described here cannot be studied experimentally because they are two-dimensional, periodic, and do not include gravity. Furthermore, the flow is close to sonic and at high Reynolds number. Experimentally, it is far easier to study the incompressible KH instability (Worster 2009). Even with this simplification, results depend sensitively on ICs (Slessor, Bond & Dimotakis 1998). This is reminiscent of difficulties when comparing simulations and experiments of the Rayleigh–Taylor instability (Dimonte et al. 2004),

which have only recently been resolved by carefully controlling the IC in experiments (Wilkinson & Jacobs 2007; Hutchinson & Rosner, in preparation).

Although we only describe simulations with an initial density ratio of 1, we have experimented with larger initial density ratios (e.g. 4). Preliminary investigation suggests that vigorous secondary instabilities become increasingly prominent as the density ratio increases, greatly enhancing mixing. Though it's a common practice to leave out explicit dissipation to model the high Reynolds numbers relevant in astrophysics, our results suggest that including explicit diffusion may provide a very effective way to reduce diffusion in astrophysical simulations with very large density ratios. We stress that these large density ratios are common in astrophysical problems such as star formation or galaxy formation. Our results demonstrate just how subtle and computationally challenging it is to correctly capture mixing in these environments (even restricting ourselves to hydrodynamics, which is likely a poor approximation).

There are many remaining questions left unanswered in this paper. It is unclear how the ATHENA algorithm seeds the IVI. We did not search for the critical perturbation amplitude that will cause a DEDALUS simulation to exhibit the IVI. Because of limited computer time, we did not find converged DEDALUS or ATHENA simulations with $\Delta\rho/\rho_0 = 1$ and $\text{Re} = 10^6$. Perhaps, contrary to expectation, increasing the Reynolds number of the system does increase the entropy production, as found in the ATHENA simulations without explicit diffusion. Future work should also test the Galilean invariance of these simulations, test ICs with an interface at an angle to the grid, and extend this analysis to larger density ratios.

We hope this study provides a well-posed test problem for future codes used in astrophysics. It would be valuable to carry out this test problem with unstructured/meshless methods (e.g. Springel 2010; Duffell & MacFadyen 2011; Hopkins 2015) to understand their convergence properties on this challenging problem. Towards this goal, we include the *reference solutions* to these KH problems in the supplementary material accompanying this paper. Introducing smooth ICs and explicit diffusion allows us to calculate a converged reference solution and compare between codes. The competing secondary instabilities for ICs with a density jump of one provides a stringent test of the fidelity with which a code solves the Navier–Stokes equations, making it a great test problem.

ACKNOWLEDGEMENTS

The authors would like to thank Ramesh Narayan, Phil Hopkins, Paul Duffell, and Cédric Beame for helpful discussions. DL is supported by the Hertz Foundation. MM was supported by the National Science Foundation grant AST-1312651. EQ is supported in part by a Simons Investigator Award from the Simons Foundation. GMV acknowledges support from the Australian Research Council, project number DE140101960. JSO is supported by a Provost's Research Fellowship from Farmingdale State College. The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this paper. This work used the Extreme Science and Engineering Discovery Environment (XSEDE allocations TG-AST140039, TG-AST140047, and TG-AST140083), which is supported by National Science Foundation grant number ACI-1053575. Resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center and the NASA Center for Climate Simulation (NCCS) at Goddard Space Flight Center.

This project was supported by NASA under TCAN grant number NNX14AB53G.

REFERENCES

- Alexiades V., Amiez G., Gremaud P.-A., 1996, *Commun. Numer. Methods Eng.*, 12, 31
- Ascher U. M., Ruuth S. J., Spiteri R. J., 1997, *Appl. Numer. Math.*, 25, 151
- Brüggen M., Hillebrandt W., 2001, *MNRAS*, 320, 73
- Choi E., Kim J., Wiita P. J., 2009, *ApJS*, 181, 413
- Colella P., 1990, *J. Comput. Phys.*, 87, 171
- Colella P., Woodward P. R., 1984, *J. Comput. Phys.*, 54, 174
- Dimonte G. et al., 2004, *Phys. Fluids*, 16, 1668
- Duffell P. C., MacFadyen A. I., 2011, *ApJS*, 197, 15
- Fefferman C. L., 2000, in Carlson J., Jaffe A., Wiles A., eds, *The Millennium Prize Problems*. Clay Mathematics Institute, Cambridge, MA, p. 57
- Fontane J., Joly L., 2008, *J. Fluid Mech.*, 612, 237
- Gardiner T. A., Stone J. M., 2008, *J. Comput. Phys.*, 227, 4123
- Hopkins P. F., 2015, *MNRAS*, 450, 53
- Johansen A., Henning T., Klahr H., 2006, *ApJ*, 643, 1219
- Kerswell R. R., Pringle C. C. T., Willis A. P., 2014, *Rep. Progress Phys.*, 77, 085901
- Lecoanet D., Brown B. P., Zweibel E. G., Burns K. J., Oishi J. S., Vasil G. M., 2014, *ApJ*, 797, 94
- McNally C. P., Lyra W., Passy J.-C., 2012, *ApJS*, 201, 18
- Mocz P., Vogelsberger M., Pakmor R., Genel S., Springel V., Hernquist L., 2015, *MNRAS*, 452, 3853
- Moore D. W., 1979, *R. Soc. Lond. Proc. Ser. A*, 365, 105
- Nulsen P. E. J., 1982, *MNRAS*, 198, 1007
- Nulsen P. E. J., 1986, *MNRAS*, 221, 377
- Porter D. H., Pouquet A., Woodward P. R., 1994, *Phys. Fluids*, 6, 2133
- Reinaud J., Joly L., Chassaing P., 2000, *Phys. Fluids*, 12, 2489
- Robertson B. E., Kravtsov A. V., Gnedin N. Y., Abel T., Rudd D. H., 2010, *MNRAS*, 401, 2463
- Salvesen G., Beckwith K., Simon J. B., O'Neill S. M., Begelman M. C., 2014, *MNRAS*, 438, 1355
- Schaal K., Bauer A., Chandrashekar P., Pakmor R., Klingenberg C., Springel V., 2015, *MNRAS*, 453, 4278
- Slessor M. D., Bond C. L., Dimotakis P. E., 1998, *J. Fluid Mech.*, 376, 115
- Springel V., 2010, *MNRAS*, 401, 791
- Stone J. M., Gardiner T. A., Teuben P., Hawley J. F., Simon J. B., 2008, *ApJS*, 178, 137
- Toro E., 2013, *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction*. Springer-Verlag, Berlin
- Wilkinson J. P., Jacobs J. W., 2007, *Phys. Fluids*, 19, 124102
- Worster G., 2009, *Understanding Fluid Flow*. AIMS Library of Mathematical Sciences, Cambridge Univ. Press, Cambridge

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article:

(<http://mnras.oxfordjournals.org/lookup/suppl/doi:10.1093/mnras/stv2564/-/DC1>).

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

APPENDIX A: INTERPOLATION TO A COMMON GRID

The grid points used in DEDALUS and ATHENA differ slightly. For a periodic simulation between 0 and L with spacing Δx , the DEDALUS

grid points are $\{0, \Delta x, 2\Delta x, \dots, L - \Delta x\}$, whereas the ATHENA grid points are $\{\Delta x/2, 3\Delta x/2, \dots, L - \Delta x/2\}$. We use two spectrally accurate methods for interpolating DEDALUS and ATHENA data to a common grid. In several cases we test both methods and find excellent agreement.

Our first method is spectral interpolation. The DEDALUS data can be viewed as either $N = L/\Delta x$ values on grid points or N Fourier coefficients. We can pad the Fourier coefficients with zeros and transform to a grid of any uniform spacing. Going from N to $2N$ points, we can compare every other entry to the ATHENA data. In a second method, we multiply the Fourier coefficients by $\exp(ik_x \Delta x/2)$. A Fourier transform then shifts the grid points by $\Delta x/2$, to align the DEDALUS grid with the ATHENA grid. We follow the same procedure in the z direction.

Throughout this paper, we treat the ATHENA data (including ICs) as cell-centred data. However, the data are actually volume averaged. The lowest order differences between cell-centred and volume-averaged quantities scales as $\sim \Delta x^2$. Thus, any errors associated with these differences should decrease with order 2. In all cases studied here (i.e. Figs 2, 4, and 9), we find better than second-order convergence. This suggests that differences due to interpreting data as cell centred rather than volume averaged are not the dominant source of error.

APPENDIX B: CONVERGENCE TO A ‘TRUE’ SOLUTION

This paper describes a series of calculations of the non-linear evolution of the KH instability, as a function of resolution and Re . Without an analytic solution, we must assess the quality of the solutions carefully. We make two assumptions to help interpret our results.

- (i) DEDALUS and ATHENA converge to the same solution at fixed Re as the resolution increases. We refer to this unattainable ‘Platonic ideal’ solution as the true solution.
- (ii) The distance (given a choice of norm) between two solutions at different resolutions (for the same code) is larger than the distance between the higher resolution simulation and true solution.

Our simulations support these assumptions, but it is very difficult, if not impossible, to prove these statements. The existence and uniqueness of solutions to the Navier–Stokes equations remain an active field of research (Fefferman 2000).

To support these assumptions, Fig. B1 plots the relative differences between simulations with $Re = 10^6$ and $\Delta\rho/\rho_0 = 0$ at $t = 6$ (described further in Section 3.1.2). The top panel assumes our highest resolution DEDALUS solution is the true solution. The bottom panel assumes our highest resolution ATHENA solution is the true solution. To assess the deviations, we plot the L_2 norm of the difference of dye concentration fields. This allows us to define an error (alternatively a distance) between two solutions X and Y as

$$e(X, Y) = L_2(c_X - c_Y), \quad (\text{B1})$$

where c_X and c_Y are the dye concentration fields of solutions X and Y , respectively. Fig. B1 remains mostly unchanged if we compare lower Reynolds number simulations with $Re = 10^5$ and $\Delta\rho/\rho_0 = 0$, although the picture is more complicated for $\Delta\rho/\rho_0 = 1$ due to chaos (see Section 3.2).

Both the ATHENA simulations and the lower resolution DEDALUS simulations are converging to D2048. The top panel of Fig. B1 therefore suggests a true solution lives very close to D2048 [assumption (i)]. The ATHENA simulations converge slower than the

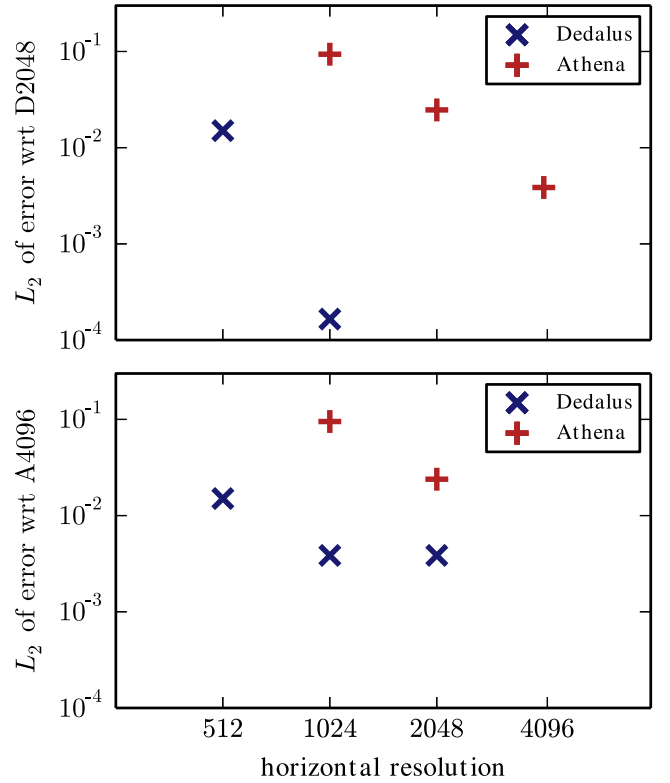


Figure B1. Differences between different solutions for $Re = 10^6$ and $\Delta\rho/\rho_0 = 0$ at $t = 6$. In the top panel, the DEDALUS simulation with horizontal resolution 2048 (D2048) is assumed to be the true solution, and in the bottom panel the ATHENA simulation with horizontal resolution 4096 (A4096) is assumed to be the true solution. The error with respect to the assumed true solution is the L_2 norm of the difference of the dye concentration fields (equation B1). The top panel shows that both ATHENA and DEDALUS are converging to the high-resolution DEDALUS solution, supporting the assumption that it is close to the true solution. In the bottom panel, the ATHENA solutions are converging to the high-resolution ATHENA simulation, but the DEDALUS solutions are not. This suggests that the ATHENA solution is further from the true solution than the DEDALUS solutions.

DEDALUS simulations because in DEDALUS spatial errors decrease exponentially.

The bottom panel of Fig. B1 shows that A4096 is a worse approximation to the true solution. This is because the DEDALUS simulations are not converging to A4096.

One could argue that perhaps the ATHENA simulations are converging to a solution near A4096 and the DEDALUS simulations are converging to a different solution near D2048. However, this would require the error of the ATHENA simulations with respect to D2048 to stay constant as the resolution increases, contrary to the top panel. Thus, we believe that both codes are converging to a true solution close to D2048 [assumption (i)].

Presumably if ATHENA were run at very high resolutions, it would become closer to the true solution than D2048. In this case, we hypothesize that both DEDALUS and ATHENA simulations would converge to this very high-resolution ATHENA simulation. For the range of resolutions examined in this paper, our highest resolution DEDALUS simulation is always closest to the true solution.

The main idea behind assumption (ii) is the convergence properties of the algorithms used in ATHENA and DEDALUS. Specifically, both codes are better than first-order accurate. Imagine we somehow know the true solution to our problem, T . If we run a high-resolution

simulation, S1, we calculate the error,

$$e(\text{S1}, \text{T}) \equiv E_1. \quad (\text{B2})$$

Now suppose we run another simulation S2 at double resolution. If S1 and S2 are converging to T, then

$$e(\text{S2}, \text{T}) \equiv E_2 < \frac{E_1}{2}, \quad (\text{B3})$$

where better than first-order accuracy implies the inequality. ATHENA is between second- and third-order accurate, so we expect $E_1/4 \leq E_2 \leq E_1/8$ for ATHENA. DEDALUS is exponentially accurate in space, and third-order accurate in time. Thus, for DEDALUS, we should expect $E_2 \leq E_1/8$. Nevertheless, equation (B3) implies, via the triangle inequality,

$$e(\text{S1}, \text{S2}) > \frac{E_1}{2} > e(\text{S2}, \text{T}), \quad (\text{B4})$$

which shows that assumption (ii) holds. One can check visually that equations (B3) and (B4) hold for the simulations described in Fig. B1, assuming that T is very close to D2048.

APPENDIX C: LYAPUNOV EXPONENT CALCULATION

One can write the equations of motion (equations 1) as

$$\partial_t U = F(U), \quad (\text{C1})$$

where $U = (\rho, \mathbf{u}, E)$ is the state vector. Then infinitesimal perturbations to U evolve according to the equation

$$\partial_t \delta U = \left. \frac{\delta F}{\delta U} \right|_U \delta U, \quad (\text{C2})$$

where $\delta F/\delta U$ is the Fréchet derivative, evaluated at $U(t)$.

To calculate the ‘local-in-time’ Lyapunov exponent, we fix the state vector to its value at a specific time $t = t_0$. The maximum

Lyapunov exponent is the greatest eigenvalue of $(\delta F/\delta U)_{U(t_0)}$. It is impractical to solve this eigenvalue problem directly – a two-dimensional problem with resolution greater than 1000 in each direction generates very large matrices. Instead, we solve an initial value problem by picking $\delta U(\tau = 0)$, and evolving

$$\partial_\tau \delta U = \left. \frac{\delta F}{\delta U} \right|_{U(t_0)} \delta U, \quad (\text{C3})$$

where τ should not be thought of as time, as we have fixed the background state U to $t = t_0$. The maximal Lyapunov exponent is

$$\lambda = \lim_{\tau \rightarrow \infty} \log \left(\frac{\|\delta U(\tau)\|}{\|\delta U(0)\|} \right) \quad (\text{C4})$$

for some norm $\|\cdot\|$. We choose $\sqrt{\|\mathbf{u}\|^2}$. This is equivalent to the power method.

We solve equation (C3) in DEDALUS using two methods. Both methods give very similar Lyapunov exponents. In the first, we directly evolve the linearized equations C3. We treat terms independent of U_0 implicitly, and treat all other terms explicitly. The second method uses an iteration. On each iteration, we evolve the full equations of motion (equation C1) for $U_0 + \delta U_i$ for a time $\Delta t \ll t_0$ to get a state we call $\tilde{U}_i(t_0 + \Delta t)$. The initial perturbation for the next iteration becomes $\delta U_{i+1} \propto \tilde{U}_i(t_0 + \Delta t) - U_0$, but with norm 10^{-8} . $\Delta U_0 = U(t_0 + \Delta t) - U(t_0)$ is the change in the unperturbed solution U_0 over the time Δt . We normalize after each iteration to ensure the perturbations stay linear.

In both cases, we initialize the calculation with a guess of random noise. After substantial evolution, the system begins to execute limit cycles (or seems to be close to a limit cycle for $\Delta\rho/\rho_0 = 1$ at $t = 2.5$). We report the growth averaged over a limit cycle.

This paper has been typeset from a $\text{T}_{\text{E}}\text{X}/\text{L}^{\text{A}}\text{T}_{\text{E}}\text{X}$ file prepared by the author.