

A Mixing Study Using Coherent Structure Dynamics to Drive a Surrogate Fluid Dynamics Model

JAY BORIS

*Chief Scientist
Material Science and Component Technology Directorate*

July 29, 2021

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

1. REPORT DATE (DD-MM-YYYY) 29-07-2021		2. REPORT TYPE NRL Memorandum Report		3. DATES COVERED (From - To) 14-09-2019 – 05-21-2021	
4. TITLE AND SUBTITLE A Mixing Study Using Coherent Structure Dynamics to Drive a Surrogate Fluid Dynamics Model				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Jay Boris				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER 1G56	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory 4555 Overlook Avenue, SW Washington, DC 20375-5320				8. PERFORMING ORGANIZATION REPORT NUMBER NRL/6003/MR--2021/1	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Naval Research Laboratory 4555 Overlook Avenue, SW Washington, DC 20375-5320				10. SPONSOR / MONITOR'S ACRONYM(S)	
				11. SPONSOR / MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT DISTRIBUTION STATEMENT A: Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES This is the second NRL Memorandum Report in a series.					
14. ABSTRACT This paper uses a reduced-order “Coherent Structure Dynamics” model, introduced in an earlier paper, to compute the time-dependent interspecies mixing due to a non-equilibrium spectrum of fluid dynamic turbulence expressing interactions between the “coherent structures” observed in turbulence (Brown and Roshko, 1974). The Coherent Structure Dynamics (CSD) model addresses situations where the time-scale for changes in the fluid dynamics driving the turbulence is short. The turbulence will be far from equilibrium when the turbulent small scales do not have time to reach their equilibrium strength. True molecular mixing and thus chemical reactions of initially separate species depends on the short scales in the velocity spectrum to mix the species. Surrogate Fluid Dynamic (SFD) realizations of such a turbulent flow, driven by the time-dependent CSD turbulence spectrum, are computed in a triply-periodic cube of a tracer density to study the mixing of two chemical species. This paper measures the time delay in molecular-scale mixing caused by the non-equilibrium delay to establish the small turbulent scales that bring the species close enough to mix molecularly.					
15. SUBJECT TERMS Computational Fluid Dynamics Coherent Structure Dynamics Non-equilibrium turbulence Turbulent Spectrum Surrogate Fluid Dynamics Multiple Fluid Dynamic Realizations					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Jay Boris
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			U

This page intentionally left blank.

Glossary:

ρ	mass density of the fluid (gm/cc).
ν	Kinematic viscosity of the fluid (gm/(sec cm)).
L_{sys}	System scale length, 10 m in examples following.
R_k	Length scale (cm) of the rotors in scale size bin k .

The following variables change in time due to stiff evolution equations . . .

E_k	Energy density (ergs/cc) of rotors of size R_k ($k = 0, kmax$).
N_k	Number density of rotors (#/cc) of scale size R_k .
ϵ_k	Energy in a single rotor of size R_k . $\epsilon_k = 3\pi\rho R_k^3 V_k^2 = E_k/N_k$

The following derived quantities are also used . . .

P_k^f	Packing fraction for rotors of size R_k (dimensionless).
	$P_k^f \equiv 3\pi R_k^3 N_k$. Thus $E_k \equiv 3\pi\rho R_k^3 V_k^2 N_k = \rho P_k^f V_k^2$
V_k	Characteristic average velocity (energy weighted) of rotors of size R_k .

$$V_k \equiv \sqrt{E_k/(\rho P_k^f)}$$

S_k	Typical separation distance (cm) of rotors of size R_k .
-------	--

$$\frac{4\pi}{3} S_k^3 N_k = 1 \rightarrow S_k \equiv R_k [P_k^f]^{1/3}$$

V_k^{pre}	Rotor precession velocity, decreases with separation $V_k^{pre} \equiv V_k R_k / S_k$
-------------	---


τ_k	Interaction time (sec) of rotors of size R_k with each other. $\tau_k = 2\pi S_k / V_k^{pre}$
----------	---

1. Introduction

The Coherent Structure Dynamics (CSD) model is based on time-dependent nonlinear fluid-dynamic interactions between a numerical representation of the “coherent structures” that actually comprise turbulence (Brown and Roshko, 1974). The goal of CSD is to represent the scale-resolved non-equilibrium evolution of turbulence within a limited volume of space. CSD is distinct from vortex dynamics in that velocities are not determined from integrating the Biot-Savart law. Rather, the coherent structures interact nonlinearly in a scale-similar fashion by local merging, which backscatters energy into larger scales, and local shredding, which cascades energy forward to smaller scales. The CSD solutions are better than statistical in the sense that the time-dependence of the coherent structures, called rotors here, is predicted for a range of scale sizes spanning orders of magnitude below the scale of the macroscopic driving flow. The CSD model was presented and the numerical representation and algorithms discussed in a previous paper, ‘Paper 1’ (Boris, 2018). This model, summarized in Section 2, advances stiff ordinary differential equations for the turbulent kinetic energy density and the number density of the coherent structures (rotors) using discrete bins of rotor scale length that are spaced logarithmically.

The CSD model is a new approach, based on the nonlinear interaction of the coherent, fluid-dynamic structures that comprise turbulence. The kinetic energy density as a function of scale

size is resolved on a logarithmic grid to describe the entire turbulent spectrum. In addition, and in distinction from shell models (Biferale, 2003), a number density of coherent structures at each scale is introduced to account for the fact that the relative spacing of the rotors, particularly away from equilibrium, is not expected to be the same at all scales. Two cases are used to study and compare the small-scale mixing: 1) an equilibrium turbulence spectrum resulting from constant stirring at the system scale L_{sys} and 2) a highly-transient non-equilibrium case where the stirring at the system scale is the only non-zero component of the spectrum initially. The details of the CSD formulation and the numerical algorithms are discussed in the earlier paper. Here a shift-based, fully Lagrangian convection model provides independent 3D realizations of the CSD time-dependent spectrum to provide diffusion-free surrogate flow fields that drive turbulent mixing.



Grinstein, LANL - FCT

Woodward, et al. Minnesota - PPM

Figure 1.1: Illustrating detailed numerical simulations of the small-scale structure of turbulent flows. Left: a simulation of flow structure in a square jet using Flux-Corrected Transport (FCT) performed by Fernando F. Grinstein, now at LANL. Right: A simulation of compressible turbulence in a periodic box using the Piecewise Parabolic Method (PPM) courtesy of Paul Woodward and colleagues at the University of Minnesota.

Figure 1.1 above presents two flow visualizations from detailed Large-Eddy Simulations (LES) approximating high Reynolds-Number flow conditions. The panel on the left is a rendering of the vorticity in a spatially evolving square jet (Grinstein, 2001) and the panel on the right is vorticity magnitude observed during turbulence decay in a triply periodic box (Woodward, et al., 2001). This latter case is the idealized system often treated in high-resolution Direct Numerical Simulation (DNS) (e.g. Moin and Mahesh,1998; Moser, et al., 1999; A. Kajzer, et al.,2014). The resolution possible in today's simulations, though considerably greater than the pace-setting computations in Fig. 1.1, is still far from adequate to treat airflow on an urban scale and yet resolve strong, intermittent eddies with scales of five to fifty centimeters that affect birds, insects, and small UAVs. Turbulence, below the macroscopic scales, is seen to be dominated by filamentary vorticity structures. Brown and Roshko (1974) observed this behavior in

experiments, calling the filaments “coherent structures.” They distinguished between these coherent structures and the view of turbulence as a mishmash of small scale random flows as had previously been thought. The new time-dependent model, being called Coherent Structure Dynamics (CSD) here, evolves these coherent structures throughout the inertial range and down into the viscous dissipation range.

A Direct Numerical Simulation (DNS) of such mixing in a Reynolds-number 10^6 flow might require $10^6 \times 10^6 \times 10^6$ computational cells if uniformly resolved. Subgrid turbulence models for use with Computational Fluid Dynamics (CFD) have generally treated the inertial-range below the CFD-resolved scale as if obeying a renormalizable or scale-similar equilibrium described by the Kolmogorov spectrum with a spectral energy density that scales as $k^{-5/3}$ in the inertial range. This assumption is only accurate when the turbulent spectrum is in or near an equilibrium energy cascade. Mixing of chemical species in CFD requires very high resolution to limit the nonphysical numerical diffusion of separated species into each other before true molecular mixing has a chance to act. Appropriately resolved CFD, when it is possible at all to limit numerical diffusion, takes days or weeks on large parallel computers.

In Fig. 1.1 the coherent structures appear as relatively isolated filaments whereas the larger scales do not appear to show such correspondingly coherent, or at least visible, structures. Vortex dynamics models, pursued for decades for vortex filament problems like this, are deterministic and Lagrangian in representation. They are based on solving the motion of vortex filaments composed of a number of line segments or spherical vortex cores coupled by velocities that can be computed from the Biot-Savart law, e.g. Knio and Ghoneim (1990) or Pullin and Saffman (1998). Only one realization of the complex flow field is computed at a time and the computational cost for this one realization becomes prohibitive when multiscale complexity in the flow approaching true turbulence is considered.

CSD describes a more statistical, reduced-order treatment of the interaction and dynamics of these coherent structures to expose important but more general behaviors. CSD, by itself, does not seek to compute any particular deterministic turbulent realization, but rather expresses the coupling of turbulent kinetic energy between the different spatial scales comprising the turbulence. The determination of specific realizations must be determined after CSD has approximated the non-equilibrium, time-varying ensemble energy spectrum from which spatially and temporally correlated realizations can be drawn.

2. The Coherent Structure Dynamics Model

The Coherent Structure Dynamics (CSD) model treats the cascade of turbulent kinetic energy up and down the spectrum by integrating ordinary differential equations describing dynamic nonlinear interactions between coherent structures in adjacent scale-size bins. Coupling of widely disparate scales in Navier-Stokes flow exists but is argued to be small (e.g. Brasseur and Corrsin, 1987; Brasseur, 1991; Brasseur and Wei, 1994) and is not treated here. No assumption is made of equilibrium behavior in the spectrum. Rather, the model aims to approximate non-equilibrium effects caused by rapid changes in the turbulence forcing and is expected to predict the rate of relaxation to equilibrium under various conditions.

The set of coupled nonlinear, stiff, ordinary differential CSD equations satisfies expected dimensional scaling conditions and other overall physical constraints. These equations treat the turbulent kinetic energy density and the coherent structure (rotor) number density for each of a

discrete set of scale size bins differing by a factor of two from one to the next. The coupling of energy between the adjacent wavelength scales is treated in CSD by computing energy-conserving fluid-dynamic interactions between the filamentary coherent structures (rotors). The required scale-similarity in the behavior of the rotor interactions at different scales is maintained. The rotor number density is a new, non-dimensional variable in this representation. It allows for deviations from ideal scale similarity. In other words, the flow patterns at different scales in the inertial range may be distinguishable geometrically without a knowledge of length scale.

The CSD model is designed to drive a separate CFD convection model to provide macroscale-through-inertial scale fluid-dynamic forcing functions for the small-scale (subgrid) turbulence in a region of space containing a modest number of CFD cells. The goal is to eventually provide, for example, realistic time-dependent realizations of the non-equilibrium turbulent flow buffeting a 10 cm UAV generated by CFD calculations or experiments measured at the 5 or 10-meter scale. CSD is also designed to address one of the key questions in combustion research: how does the mixing and energy-releasing reactions of fuels and oxidizers drive and shape the turbulence at small scales that controls the energy release? CSD also provides a framework to evaluate the modifications to standard approximations for heat and momentum transfer that are caused by high Mach-number flow.




Figure 2.1: Schematic of the logarithmic k -space grid used in Coherent Structure Dynamics. The wavenumber increases and the corresponding turbulence scale size decreases to the right. The turbulent energy density (red) decreases from bin to bin with k , the bin index, while the rotor number density (brown) increases. Rotor merging and rotor shredding are local interactions.

The grid used to represent the rotor sizes is logarithmic in wave number, or scale space, as shown schematically in Figure 2.1 above. The rotor scale lengths in the CSD model are discretized into bins separated by a factor of two in size. This allows many orders of magnitude in scale to be represented with a few tens of bins and independent variables. The horizontal red lines in each size bin below show a decreasing energy density E_k on a log-log scale and the brown lines show an increasing rotor number density N_k . This compact representation provides a fast evaluation of the evolving non-equilibrium turbulence spectrum that can be evaluated at a number of separate locations in space. This speed could be used to estimate the high-resolution turbulence, for example, within a number of spatial cells in a region of a CFD computation. The

goal is an ongoing estimation of the turbulence at spatial scales well below those that could possibly be resolved by a macroscopic 3D LES or DNS model.

The turbulent flow field is treated as a superposition of rotors of many different scales and strengths that interact among themselves locally. These interactions will produce the nonlinear turbulent behavior that is defined, overall by the Navier-Stokes equation, and thus the fluid dynamics that we know. This is, I suppose, a kind of wavelet representation. Fig. 3.1 in Paper 1 shows a notional structure for these rotors, a simple idealization for the purpose of constructing a higher-level model. Here the rotor scales are discretized logarithmically, each bin a factor of two in scale different from the next longest and the next shortest wavelength.

Rotors at a given scale can all have different strengths and orientations but it is dynamically interesting to treat all rotors in a local volume as having a single kinetic energy at each scale. There clearly has to be some consideration given for nearby rotors having different energies but, for now, we consider this distribution to be a delta-function at the average energy density. The energy varies as kinetic energy enters and leaves a given bin due to nonlinear fluid dynamic interactions. Another assumption we are making is that interactions between rotors of different scales occurs between adjacent bins only. Brasseur and colleagues have shown that most of the energy transfer in turbulent cascade occurs locally in k space.

Two nonlinear rotor interactions, or processes, are considered in CSD:

Rotor Shredding/Forward Cascade: generates new rotors of size R_{k+1} from the interaction of rotors of the next larger size R_k .

Rotor Merging/Inverse Cascade: generates rotors of size R_{k-1} by merging two rotors of the next smaller size R_k . Inverse cascade is expected to be much weaker than forward cascade in nearly all situations.

The dimensionless scale factors γ and δ in Eq. 2.1 below govern the strength of forward cascade and inverse cascade respectively. Their ratio could be subsumed in adjusting the time scale. Written as below, γ and δ are constant across all bins as required by scale independence of the inviscid component of the fluid dynamics. Being constant also ensures rigorous energy conservation summed over all scales. What leaves one scale bin must enter the next smaller (or larger) bin, except for viscous damping and energy transfer beyond the largest and smallest rotor scales, which become “boundary” conditions for the model. The factor ν is dimensional and governs the strength of viscous dissipation, with units of m^2/sec . The full kinetic energy density equations being solved are written as:

$$\frac{dE_k}{dt} = \frac{dE_k^{ext}}{dt} + \gamma \frac{V_{k-1}}{R_{k-1}} [P_{k-1}^f]^{2/3} E_{k-1} + \delta \frac{V_{k+1}}{R_{k+1}} [P_{k+1}^f]^{2/3} E_{k+1} - (\gamma + \delta) \frac{V_k}{R_k} [P_k^f]^{2/3} E_k - \nu \frac{1}{R_k^2} E_k. \quad \text{Eq. 2.1}$$

In Eq. 2.1 the ‘packing fraction’ of these rotors is defined as

$$P_k^f \equiv 3\pi R_k^3 N_k \quad \text{Eq. 2.2}$$

and the characteristic velocity of a rotor at the edge of its core is

$$V_k \equiv \sqrt{E_k / (\rho P_k^f)}. \quad \text{Eq. 2.3}$$

The packing fraction is a non-dimensional number that describes how isolated the rotors at each scale are from each other. A small packing fraction means that the rotors are far apart relative to the characteristic scale size of each bin R_k . When the packing fraction is constant over all bins, scale similarity, in the sense usually meant is possible. Equation 2.3 shows that the characteristic velocity of rotors of scale k becomes large when the number density is small, for a given energy density, i.e. when the packing fraction is small.

Equation 2.1, including the definitions in Eq. 2.2 and 2.3, is intrinsically nonlinear. The first term on the right side is an external energy “stirring” source, which will typically be applied only at the longest wavelengths. The next two terms on the right side of Eq. 2.1 are production terms for larger rotors and smaller rotors respectively based on scale- k interactions. These are followed by a combined destruction term from rotors shredded and merged at scale k . Finally, on the right, is the viscous rotor dissipation term. These production and destruction terms can be evaluated given the principal quantities at any time but their sum, external to the numerical integration, will not necessarily coincide with the effective values actually applied by the stiff equation integrator, CHEMEQ2, being used (Mott, 1999; Mott and Oran, 2001; Young and Boris, 1973,1977; Young, 1981).

In addition to the nonlinear energy cascade Eq. 2.1 from above, we also need an equation to advance the rotor number densities in each scale bin. The current version of CSD uses

$$\begin{aligned} \frac{dN_k}{dt} = \frac{dN_k^{ext}}{dt} + \frac{\eta^s}{2} \gamma \frac{V_{k-1}}{R_{k-1}} [P_{k-1}^f]^{2/3} N_{k-1} + \frac{1}{2} \delta \frac{V_{k+1}}{R_{k+1}} [P_{k+1}^f]^{2/3} N_{k+1} \\ - (\gamma + \delta) \frac{V_k}{R_k} [P_k^f]^{2/3} N_k - \nu \frac{F_k^{vd}}{R_k^2} N_k. \end{aligned} \quad \text{Eq. 2.4}$$

The first term on the right is an external source of rotors which will be present if, for example, the external energy source is provided by stirring. This equation is also in the form that can be integrated using CHEMEQ2 although it is explicitly not in conservation form. The factor $\frac{1}{2}$ in the second term on the right indicates that one rotor is formed from two during merger. The term $\frac{\eta^s}{2}$ in the third term on the right indicates that two rotors shred each other to form η^s rotors in the next smaller-scale bin. These η^s new rotors will share the energy cascading down from the next larger scale with the already existing rotors at the smaller scale. The viscous term in Eq. 2.4 is included to allow the removal of rotors that have become dynamically insignificant in time due to viscous decay. The viscous decay factor F_k^{vd} for each scale has yet to be determined. These equations are discussed in greater detail in paper 1.

Clearly a whole range of different interactions could be defined, beyond the simple shredding and merging interactions used here, including non-local interactions connecting distant bins. However, these two coherent structure interactions suffice. Further, for the two interactions identified, merging and shredding, when distributions of rotor properties in each bin are admitted, much more detailed dynamics and varied rotor end states could be included in the model. A third process, viscous rotor decay is of paramount importance at small scales. Rotor decay reduces the kinetic energy E_k of rotors of size R_k diffusively due to viscosity without changing the number of rotors N_k or the packing fraction P_k^f . Thus, viscous dissipation terminates the Kolmogorov cascade.

3. Two Scenarios Considered in This Paper

Figure 3.1 below shows a highly non-equilibrium turbulence spectrum which is used as the initial condition for the CSD model in Scenario 1. The rotor distribution is resolved by 30 scale-size bins ranging from 10 meters down to about 10^{-4} centimeters but only the 10-meter scale is excited initially. The y-axis in the figure is a normalized logarithmic scale with each horizontal grid line marking a factor of 10, as shown at the left. All primary variables except the packing fraction are plotted on this scale. The rotor number density N_k in green increases with wavenumber from left to right whereas the other variables are all decreasing – even in this non-equilibrium condition. As described in Paper 1, the model was run with a fixed timestep from the initial condition of Fig. 3.1 using a constant stirring velocity of 1.0 m/s in bin 1. The characteristic time for the flow to cross this system, measured at the edge of the rotor cores, is 20 seconds. In this time interval, the large-scale rotors shred each other and the energy in the system has more than doubled, driven by the continual stirring. Appreciable excitations at the 15-cm and 20-cm scales have appeared, as shown in Paper 1 for a slightly faster problem with 2.0 m/s stirring speed, with turnover times an order of magnitude faster than the large-scale rotors.

This initial condition evolves nonlinearly for 60 seconds in Scenario 1 (Section 5 of this paper) but the mixing begins slowly and is delayed until an equilibrium spectrum with energetic short wavelengths is approached. Scenario 2, discussed in the next section also mixes the species turbulently for 60 seconds but the initial mixing is much stronger because the start of the simulation was delayed 300 seconds while the spectrum fully developed and equilibrated. The spectral energy density ε_k , the rotor energy density $E_k = k\varepsilon_k$, and the rotor velocity V_k in Scenario 1 have almost fully equilibrated in 20 to 30 seconds, reaching approximately the values shown in Fig. 3.2. However, it takes several minutes for the packing fraction P_k^f to settle down. As shown in paper 1, energy in the spectrum cascades to shorter and shorter wavelength until, 5 minutes later, the spectrum has equilibrated as shown in figure 3.2 below. A 60-second mixing simulations driven by that equilibrated spectrum comprises Scenario 2 (Section 6 of this paper).

In Figs. 3.1 and 3.2 below the light blue straight line is the rotor energy density in the ideal Kolmogorov spectrum $E_k \sim k^{-2/3}$ and the dark blue curve in Fig. 3.2, switching to black as it enters the viscous dissipation zone, is $E_k(t)$. The dark green lines in Fig 3.1 and 3.2 respectively are the rotor Number densities as initialized and computed after 300 seconds. The light green straight lines up and to the right in each figure show the number density of rotors at each scale that would exist with a packing fraction P_k^f varying as $k^{-\alpha}$. The packing fraction is shown in blue-green. The solid blue-green horizontal line in Fig. 3.1 with black squares at the bin locations is the initial condition with each bin having the packing fraction 0.1. In Fig. 3.2 the packing fraction has relaxed to its equilibrium configuration, roughly approximated by the blue-green squares in the lower-left marking $P_k^f \sim k^{-1/6}$. Note that P_k^f oscillates in the rotor size range 10^{-3} to 10^{-4} cm at the end of the dissipation range where E_k , the dark blue to black curve, and V_k , the purple curve, become vanishingly small. The packing fraction is hard to define when the coherent structures are disappearing due to viscosity.

The orange and red curves in Figs. 3.1 and 3.2 show the ideal and actual Kolmogorov spectral energy density as the dissipation zone is approached and entered at small scales.




Figure 3.1: Initial conditions for an equilibrating turbulent spectrum computed by the CSD model. All the energy is initially in the rotors of 10-meter scale (bin 1). The initial packing fraction is 0.9; the entire volume is essentially packed with the driven largest-scale rotors. The energies and velocity are 9 orders of magnitude smaller in all bins except the first. The constant stirring speed is 1.0 m/s at the 10-meter scale.

Figure 3.2 below shows the fully equilibrated turbulence spectrum computed using CSD with 30 bins whose sizes range from 10 meters down to about 10^{-4} centimeters.

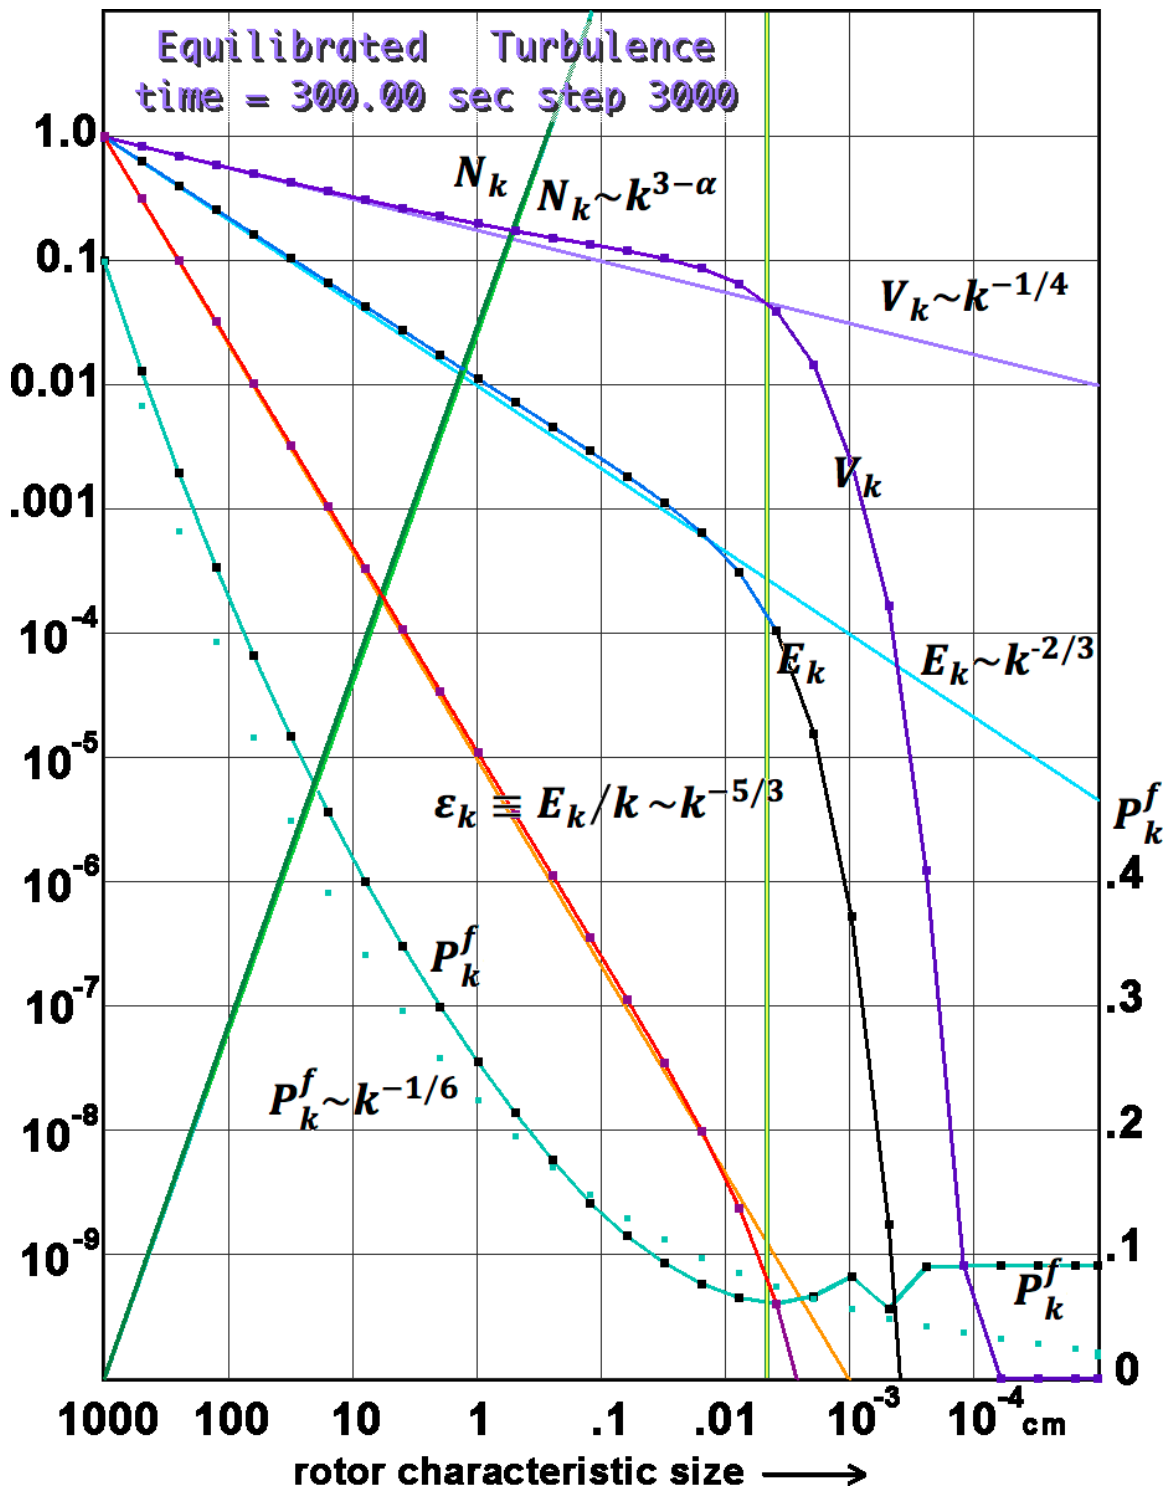


Figure 3.2: An equilibrated turbulent spectrum computed by the CSD model beginning from the initial condition of Fig. 3.1 and run for 300 seconds. The spectral energy density ϵ_k , the rotor energy density $E_k = k\epsilon_k$, and the rotor velocity V_k are nearly equilibrated in 20 to 30 seconds but it takes several minutes for the packing fraction P_k^f to settled down.

The two vertical yellow/green lines in Fig. 3.2 mark the beginning of the viscous dissipation zone around 0.05 cm. The “bottleneck,” also sometimes called the pre-dissipative bump, is clearly seen just to the left of the yellow/green lines. It is most easily seen in the lavender V_k curve, which deviates above the power-law $V_k \sim k^{-1/4}$ before dropping sharply in the dissipation zone. This steady, fully-developed spectrum is used to define the turbulence for Scenario 2. In each scenario, two fully separated species are mixed by a realization of the turbulent flow corresponding to the two scenarios. Paper 1 presents accuracy studies and convergence studies for the CSD component of these two scenarios. Section 4 describes the diffusion-free Surrogate Fluid Dynamics (SFD) model. Section 5 and 6 consider the effect of developing non-equilibrium turbulence to delay the small-scale “molecular” mixing that can lead to chemical reactions.

4. The Dissipation-Free Mixing Model

The Coherent Structure Dynamics (CSD) model is being used here to investigate molecular mixing of initially separated species in the context of a high Reynold’s-number, turbulent flow. This section describes the 3D axis-rectified convection model, call Surrogate Fluid Dynamics (SFD) that was developed to produce physically reasonable flow realizations that are driven by the time-dependent turbulent velocities computed in each of the CSD bins. The approximate convection model in SFD uses axis-aligned shifts of the tracer (chemical) species densities and thus is entirely Lagrangian without numerical diffusion. To cause molecular mixing at all, SFD adds a small amount of grid-scale diffusion every other timestep. Section 9 returns to the subject of the added diffusion in Surrogate Fluid Dynamics after demonstrating convergence with timestep and with grid resolution in Sections 7 and 8 respectively.

Turbulence continually convolutes and stretches the geometric interface that separates two species. Actual molecular mixing, however, can only occur within a few mean-free-paths of this evolving interface. Therefore, only a small fraction of the volume should become well mixed for a relatively long time – at least until the turbulent interface has moved close to most locations in the volume. Unless the computational cells are smaller than the density gradients that would contribute most heavily to molecular mixing, resolution-based numerical errors in CFD will give the appearance of mixing prematurely two species that are entirely unmixed initially. To use a grid-based flow algorithm in solving the continuity equations for different species requires very high resolution to keep the species physically separated for as long as they should be. Interface tracking approaches offer little help in resolving turbulence with diffusive scales that are orders of magnitude smaller than the macroscopic flow scales that drive the turbulence.

In the current Sections 4, 5 and 6, the question treated is “how much is mixing delayed when the turbulence is not fully developed?” The viscosity is given the value 0.01 cm²/second, approximating its value for water. This makes the Reynold’s number nominally of order several millions in a 10-meter triply periodic domain. Detailed Numerical Simulation (DNS) would not be possible for these cases; millions of cells would be needed in each CFD dimension. To maintain the species separation for as long as possible in SFD, the purely convective aspect of the turbulent flow is represented by shifting the density profile a finite number of cells parallel to each of the coordinate axes of the grid. This is reproducing convection in an idealized Lagrangian way without any numerical dissipation.

To remove numerical diffusion entirely from SFD, motions parallel to the X, Y, and Z axes are only allowed to traverse an integral number of cells every timestep on the Cartesian grid. Any

required fractional cell motions are not enacted but rather accumulated for use in a timestep when an integral shift threshold is reached. Thus interpolation back onto the grid is unnecessary. Thus all motions on the SFD grid use efficient shift operations that are a factor of 30 to 50 times faster than common 3D CFD algorithms. These shifts are direction split and cycled through the three dimensions. The length of the shift on each line is determined by the rotor velocities in each bin at the time the shift is made. In 2D this was once called a “cookie-cutter model.” In 3D the relative, shearing shifts between adjacent lines generate density gradients. The term “fine-graining” is some-times used to characterize how a density can become highly variegated by the continuing action of longer wavelength flow components.

In SFD’s simplest form, a single species tracer density is defined on a regular, triply periodic grid. Four grids are used to study resolution, 128^3 , 256^3 , 512^3 and 1024^3 . The SFD tracer density values can be convected over 10 systems lengths on the 1024^3 grid in an hour or so on a single MacBook Pro. Even at the highest resolution, however, this density field represents only the longest wavelengths of the 25-30 factor-of-two CSD bins used to represent the turbulence evolution. The highest-resolution 1024^3 grid does not actually penetrate the viscous-dissipation range of the driving flow for Reynold’s numbers greater than 5 or 10 thousand. Nevertheless, this resolution covers a wide enough portion of the inertial range below the CFD grid scale to study the mixing in both equilibrium and non-equilibrium turbulence, as will be shown in the following sections. Since the convective shifts can be performed at many different scales down to the cell size simultaneously, orders of magnitude in convective scales are being represented without numerical diffusion - albeit relatively crudely. Mixing is forced to occur only at the grid scale by an explicit diffusion operation added at the grid spacing. Only this diffusion mixes the tracer density values representing chemical species.

A two-species mixing problem can be represented by a single real array, defined as either 1.0 or 0.0 initially in the large-scale checkerboard pattern shown in the upper left of Figs. 4.1 and 4.2. In these two figures, the first 5.0 seconds, shown in the top six panels, are computed using $1/16^{\text{th}}$ second timesteps with the added grid scale diffusion turned off to better illustrate the composite effect of cyclic-shift convection at all resolved scales. For the remaining 6 panels out to 10 seconds, 160 timesteps on the 1024^3 grid and 80 timesteps of $1/8^{\text{th}}$ second on the 512^3 grid, the added diffusion is turned on and “molecular” mixing begins.

Figure 4.1 below shows the tracer density at one specific layer of the tracer density, Z layer 24. The quantity plotted in Fig. 4.2 is the average of all density values in each X-Y column along Z perpendicular to the figure. Black indicates density values between 0.0 and 0.0001. Magenta indicates values between 0.9999 and 1.0 in subsequent panels of these visual snapshots of the solution. Once the layer value or average value being plotted moves away from these limits later in time, the plot uses a continuous color scale from blue near zero to red near one. Grey dithering is applied in a thin range around 0.5 to show where the fluid has become essentially fully mixed.

Layer 24, and most of the other layers, show unmixed tracer densities (magenta and black) in some regions for as long as 15 or 20 seconds after mixing begins. The vertically-averaged density in Fig. 4.2, however, is less than 1.0 and greater than 0.0 in the entire plotted cube average after 10 seconds, as shown by the absence of magenta or black in Figs. 4.2 and 4.3. Shearing flows in the X- and Y-directions have moved up to half the system length in each direction so the vertical averages include at least a few cells of different densities. Also note that the averaged tracer density is smoother than the layer density as expected.

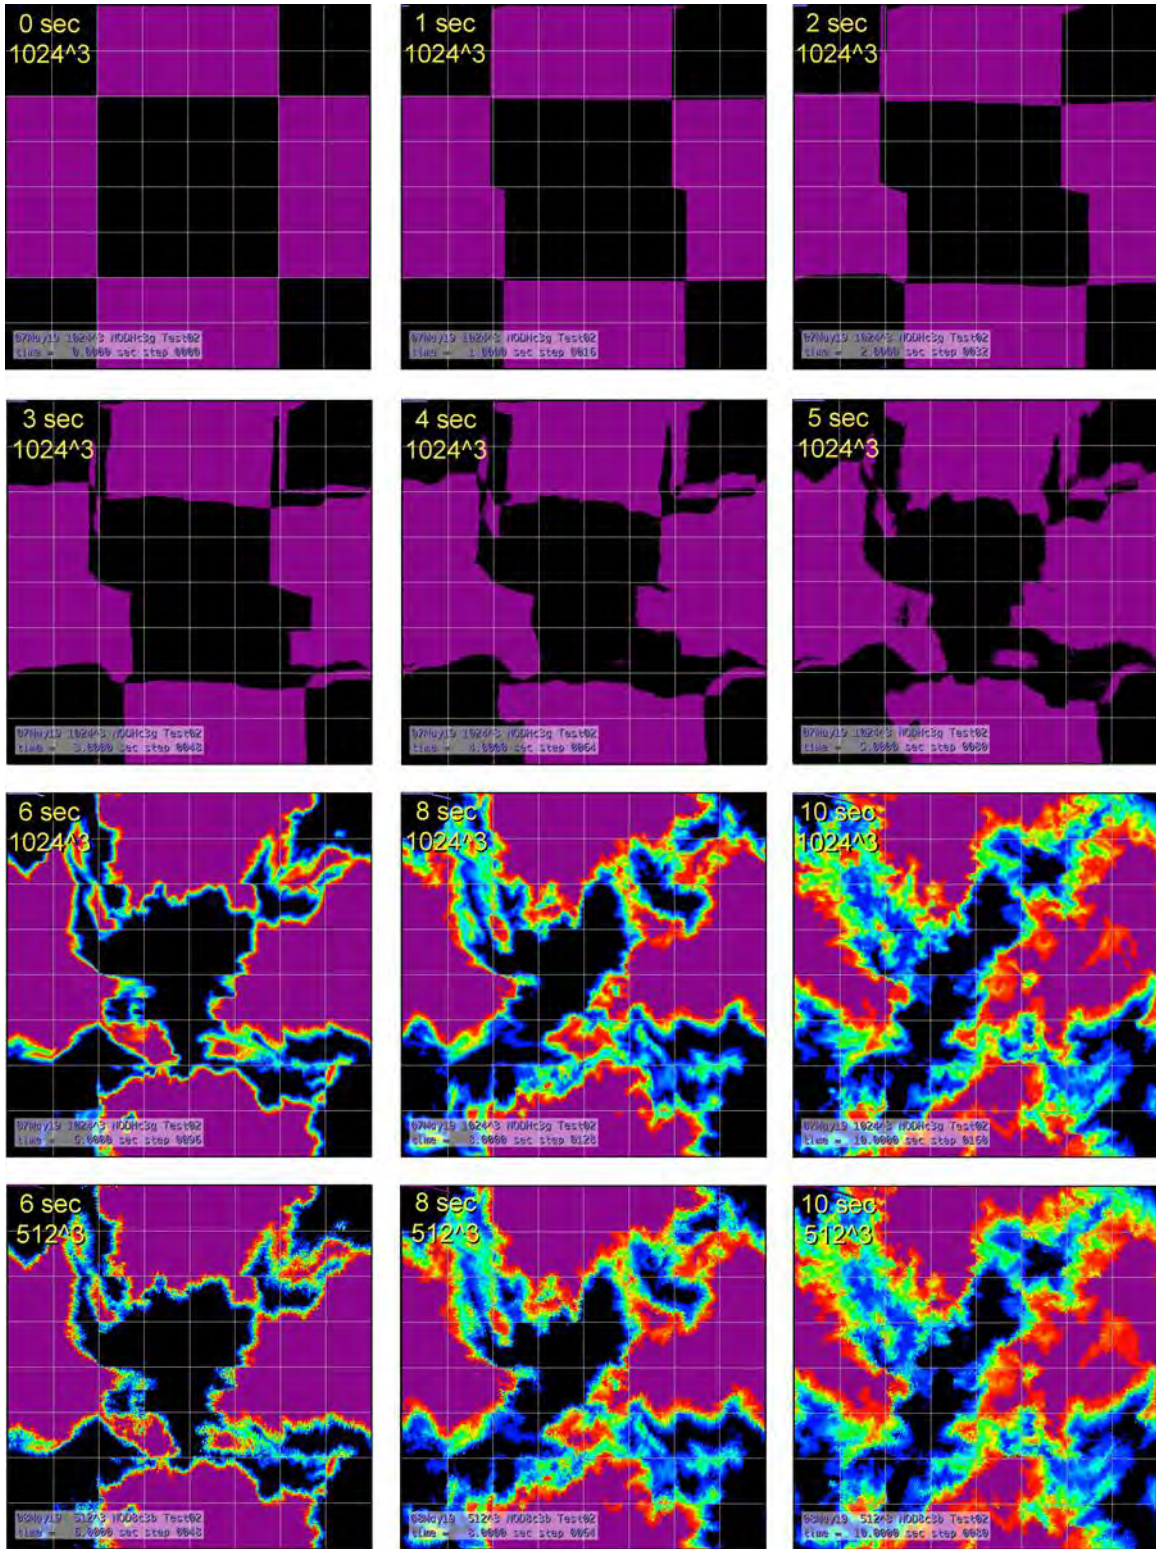


Figure 4.1: Passive tracer density on Z plane 24 with initial values 1.0 (magenta) and 0.0 (black). The times of each snapshot are shown in the upper left. When grid-scale diffusive mixing sets in after 5 seconds, the images show the maximum density (< 1.0) on this single plane as red and the minimum (> 0.0) as dark blue. The top three rows were computed with a 1024^3 grid. The bottom row, for comparison with the images on the row directly above, used Z plane 12 on a 512^3 grid.

The detailed solution of 3D continuity equations for this problem is possible in principle using Direct Numerical Simulation (DNS) but at prohibitively high resolution. This simple SFD mixing model is more computable and has much in common with turbulent convective mixing. Movies of the evolution shown in Fig. 4.1 appear realistic despite the inaccuracies of using system-wide cyclic shifts to represent convection. Different realizations of the small scales in turbulent flows cause otherwise identical initial conditions to begin to deviate after one or two large-eddy turnover times. The phasing of even large-scale fluctuations can deviate quickly by 100% although macroscopic averages will generally track each other for longer times, even at different spatial resolutions.

The panels of Fig. 4.1 suggest that mixing does not begin until the turbulence small scales have become populated. The diffusion, applied here after 5 seconds, is coupled to the smallest grid scale on the assumption that smaller scales down to the physical diffusive-mixing scale will become populated as soon as the smallest grid-resolved scales have become energized. Because the eddy turnover time for unresolved scales (Oran and Boris, 1993) is much faster than for the smallest resolved scales in the inertial range, we are safe in assuming that actual mixing has occurred. The top two rows in Fig. 4.1 show the plane 24 density at $t = 0$ through 5 seconds after the run begins. The CSD spectrum for this case was initialized with zero velocity in all but the largest scale, which is then stirred continually with an average rotor velocity of 1 meter/sec. After 5 seconds, the velocity at the resolved scale, bin 10 on this 1024^3 grid, is moving the fluid far enough for diffusion to begin taking place so the added diffusion is turned on.

The panels at 1 through 5 seconds in Fig 4.2 below seem to indicate some immediate mixing. What is actually being shown is the vertical averaging of different grid cells with 1.0 and 0.0 density along each vertical line through the grid perpendicular to the image. The bottom two rows of Fig. 4.1 at 6, 8, and 10 seconds show portions of layer $k = 24$, or layer 12 on the 512^3 grid, that still have the values of 1.0 and 0.0 with no evidence of local mixing. The interfaces where mixing is occurring in Fig. 4.1 are clearly seen and are only a couple of grid cells thick. Only very small areas of the images at 6, 8 and 10 seconds in Fig. 4.2 show the value 1.0 (magenta) and 0.0 (black) extending all the way through the cube.

Note again that the bottom row in these two figures, computed on the 512^3 grid rather than the 1024^3 grid, are essentially identical to the images at higher resolution shown immediately above. The large-scale flow at these two resolutions are identical except for the smallest resolved scale because the fluid dynamics of the SFD realization do not modify the turbulence. These interactions are computed in the stiff CSD integrations providing the time-dependent spectrum. This is a convenient feature of the CSD-SFD combination when looking for convergence in mixing computations because the detailed species densities do not affect the applied turbulent spectrum. In a CFD simulation the inclusion of smaller scales rather quickly changes the larger scale components of the flow, with or without chemical energy release, which makes the comparison of different-resolution computations devilishly difficult.

Rotor shredding populates the successively smaller bins at the rate specified by the CSD model equations. Therefore the rate of mixing is nearly zero early on when the interface area between the tracer density 1.0 and 0.0 is minimal. The interface area changes rapidly once the smaller scales of the spectrum are active because the curvature of the interfaces grows quickly and thus the small scales can cause an ever faster interface area increase and correspondingly faster mixing than the large scales.

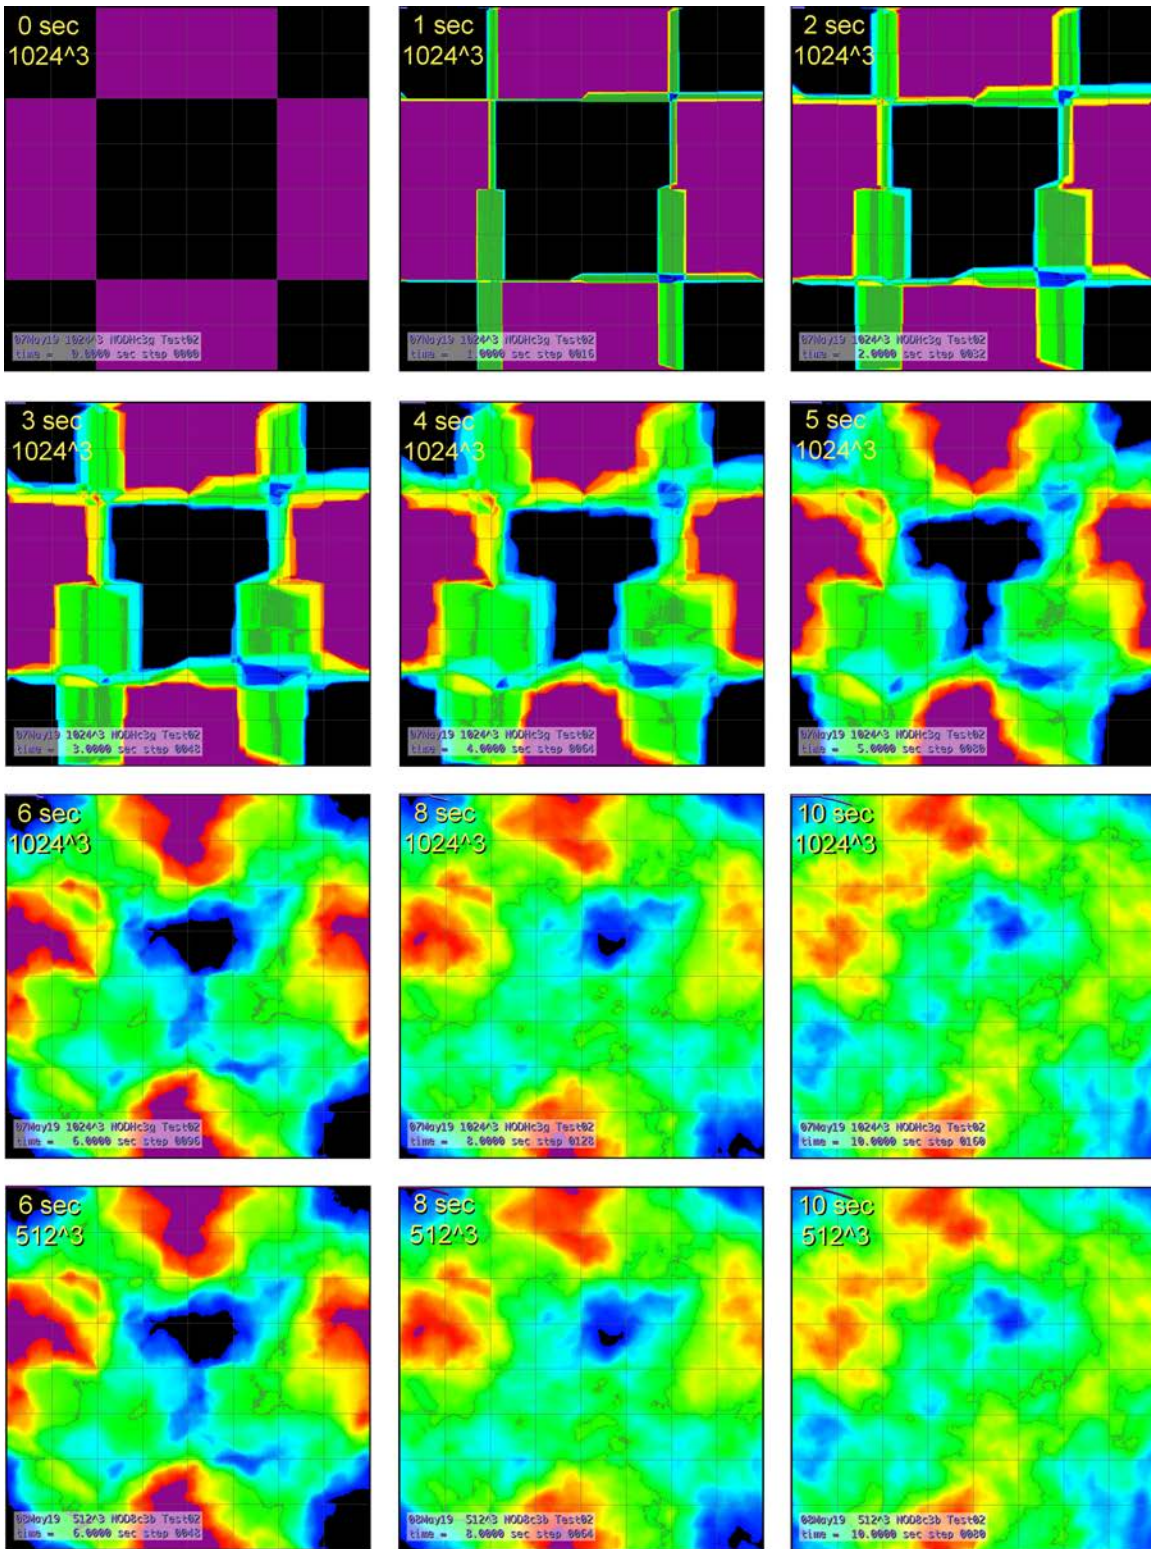


Figure 4.2: Vertically averaged passive tracer density. Initial values are 1.0 (magenta) and 0.0 (black). Snapshot times are shown in the upper left. Grid-scale diffusive sets in around 5 seconds. The images then show the maximum density (< 1.0) on the cube as red and the minimum density (> 0.0) as dark blue. The top three rows used a 1024^3 grid and the bottom row used a 512^3 grid.

Figure 4.3 below extends the simulation in Fig. 4.1 out to 60 seconds of run time. The lavender curve superimposed on each panel starting at the upper left of Fig. 4.3 is the Mixed Volume Index,

MVI(t), plotted as a function of time from 0.0 seconds on the left to 60 seconds on the right. With a single tracer, as used in Sections 4 through 6, MVI(t) measures the decreasing unmixed fraction of the inter-penetrating species, a quantity discussed further in the next section. Each horizontal grey fiducial line marks a square root of 10 (~ 3.16) reduction in this index.

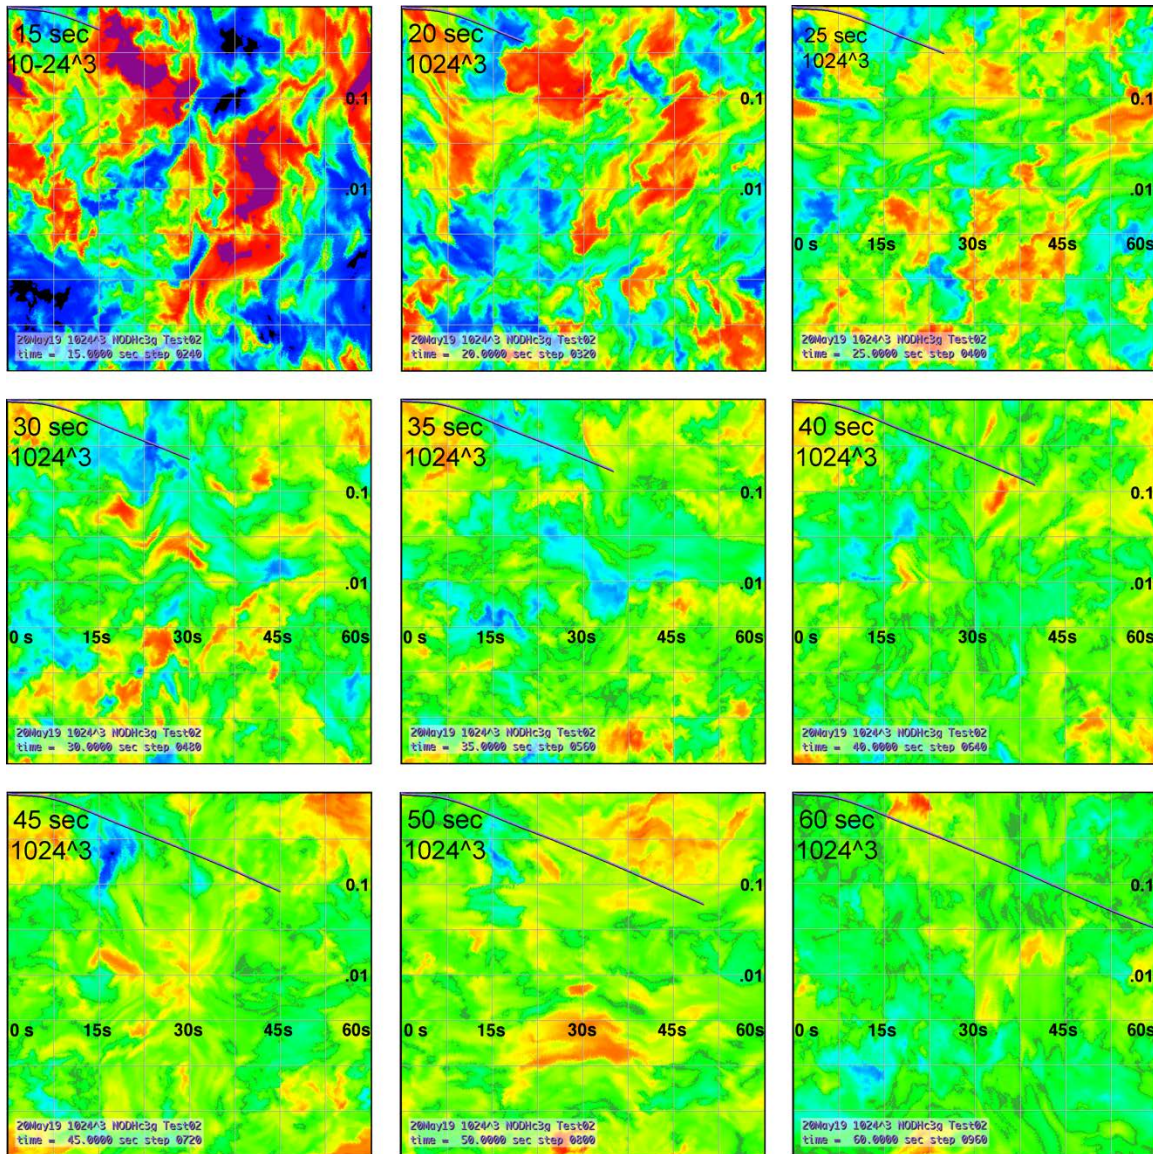


Figure 4.3: Mixing tracer density on a Z plane 24 with initial values 1.0 (magenta) and 0.0 (black). The snapshot times are shown in the upper left. The resolution is 1024^3 . With grid-scale diffusive mixing, the images show the maximum density (< 1.0) on this plane as red and the minimum (> 0.0) as dark blue. Each horizontal grey line marks $1/8^{\text{th}}$ of the $N_{xyz} = 1024$ cells and also indicates a square-root of 10 reduction in the lavender Mixed Volume Index curve entering from the upper left.

Because the turbulence spectrum is non-equilibrium, with a 1.0 m/s rotor speed in bin 1 at $t = 0$ seconds, the slope of MVI(t) at $t = 0$ is nearly zero. Essentially no mixing is taking place because only the longest wavelength has energy and the interface has its smallest area. Nonlinear cascade due to rotor shredding progressively energizes the smaller and smaller scales until a steady Kolmogorov cascade equilibrates. This equilibrium spectrum becomes the driving term for SFD in Section 6.

5. The First Scenario: Mixing Begins Slowly as the Turbulence Evolves

In this section the turbulent spectrum evolves from zero during the SFD mixing. Only the largest scale has energy at $t = 0$ before the spectrum fills in. The species $MVI(t)$ is defined here as

$$MVI(t) \equiv \frac{\sum |\rho(x, y, z, t) - 0.5|}{\sum |\rho(x, y, z, 0) - 0.5|} \quad \text{Eq. 5.1}$$

where 0.5 is the final fully mixed tracer density and the sums are taken over all N_{xyz}^3 cells in the cube. The absolute values appear in Eq. 5.1 so the deviations of the tracer density, both above and below the domain average value, are being summed.

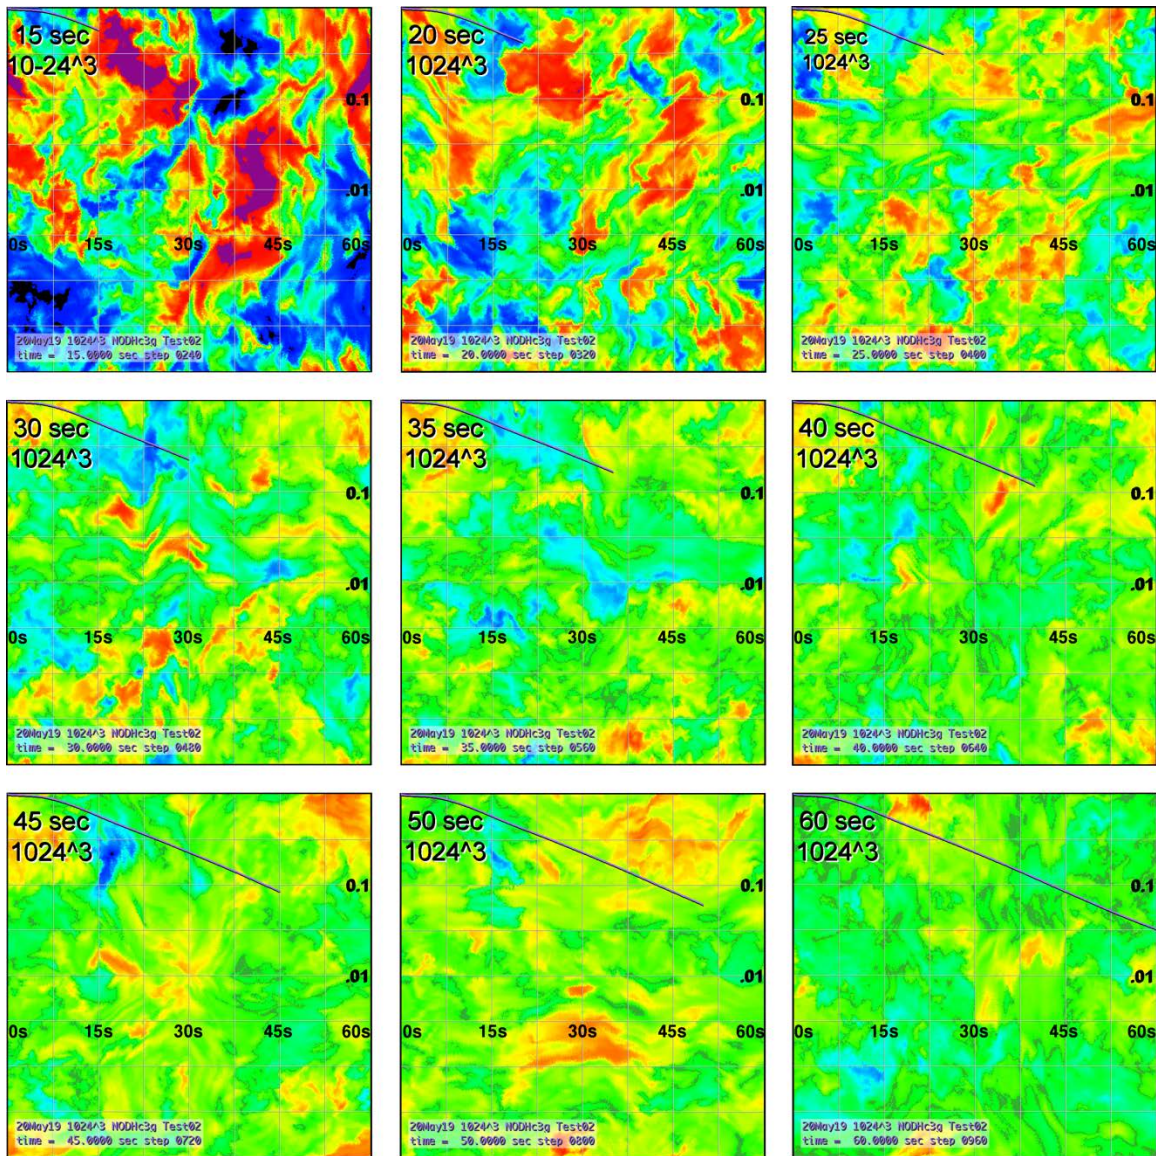


Figure 5.1: Tracer density on Z plane 24 in evolving non-equilibrium turbulence from 15 sec to 60 sec. Each panel plots the maximum density at that time as red (< 1.0) and the minimum as dark blue (> 0.0). The lavender curve from the upper left on each panel is the Mixed Volume Index, $MVI(t)$. Each horizontal greyline marks a square root of 10 (3.162...) reduction in MVI.

Figure 5.1 above has been computed using the 1024^3 grid and a timestep of $1/16^{\text{th}}$ second. Thus runs to 60 seconds require 960 timesteps. The timescale for the MVI(t) curve is shown from 0 to 60 seconds across the lower third of the figure and the scale for the monotonically decreasing MVI values is shown on the right edge of the plots, indicating the lines where MVI(t) has dropped from 1.0 to 0.1 and later to 0.01. The MVI curve seems to be approaching a straight line on this plot, indicating an exponential decrease of MVI toward zero. Since the spectrum has equilibrated half way through the run, the mixing which continues has a structure that does not change and continues to reappportion the residual density variations to smaller and smaller levels. The process can be expected to be mixed by the turbulence in the same manner from second to second without any change in the general structure of the residual unmixed component, just an exponential decrease in amplitude.

Near $t = 0$ in the upper left, MVI(t) is almost flat which shows that mixing of the tracer with the background value zero, in the black regions, is initially very slow. As the interface area between the high density and low density regions increases due to the growing turbulence and continued distortion of the evolving interface,, the mixing rate and thus the slope of MVI increases. At 15 seconds in the upper left panel, there are still noticeable regions of tracer density 1.0 (magenta) and 0.0 (black). By 20 seconds these regions, at least in plane 24 as shown, have vanished. By this time, as well, the exponential decay of MVI is well established and changes very little if at all from 15 to 60 seconds.

6. The Second Scenario: Mixing Is Faster with Fully Equilibrated Turbulence

The second scenario, presented in this section, begins with an initial tracer checkerboard identical to the first scenario presented in Section 5. However, the run begins after the spectrum of turbulence has fully evolved to the equilibrium condition shown in Fig. 3.2. For the 60 seconds (960 timesteps) of the run shown in the 9 panels of Fig. 6.1, the spectrum of turbulent speeds and the rotor number density in each scale-size bin is held fixed but the direction and spatial disposition of these speeds varies continuously in time throughout the cubical domain. The difference of the mixing, captured in the MVI(t) curves in the two runs, arises because the mixing is slower initially in Scenario 1 as the turbulence builds up.

The times of the nine snapshots in Fig. 6.1 are the same as those in Fig. 5.1 but the tracer density plots are noticeably similar only at 15 seconds in the upper left of each figure. Once 20 to 25 seconds have elapsed, shown in the upper center and upper right panels, the disposition of the tracer density on Z plane 24 in the two panels is quite different. In this second scenario the MVI(t) slope is significant near $t = 0$ and increases up to about 30 seconds elapsed time. Thus mixing proceeds more slowly in Fig. 5.1 than in Fig. 6.1. This visible delay occurs because the turbulent small scales are churning from the beginning of the Scenario 2 run but the interfaces are initially flat so it takes a while for the small scales to have their maximum mixing effect since they are not yet causing actual mixing throughout the volume, only near the evolving interface.

Both Figs. 5.1 and 6.1 are tracer densities on a single X-Y plane at constant Z chosen to be at integral levels in all four resolutions. The panels in neither figure are particularly smooth as a result of being a single plane. This is most clearly shown by the grey markings in the middle of the green shading where the cell tracer density lies within 1% of the average value between the maximum and the minimum density on the entire Z plane with $k = 24$. The bottom row of panels in each figure show that this most fully mixed region is irregular and splotchy in many areas.

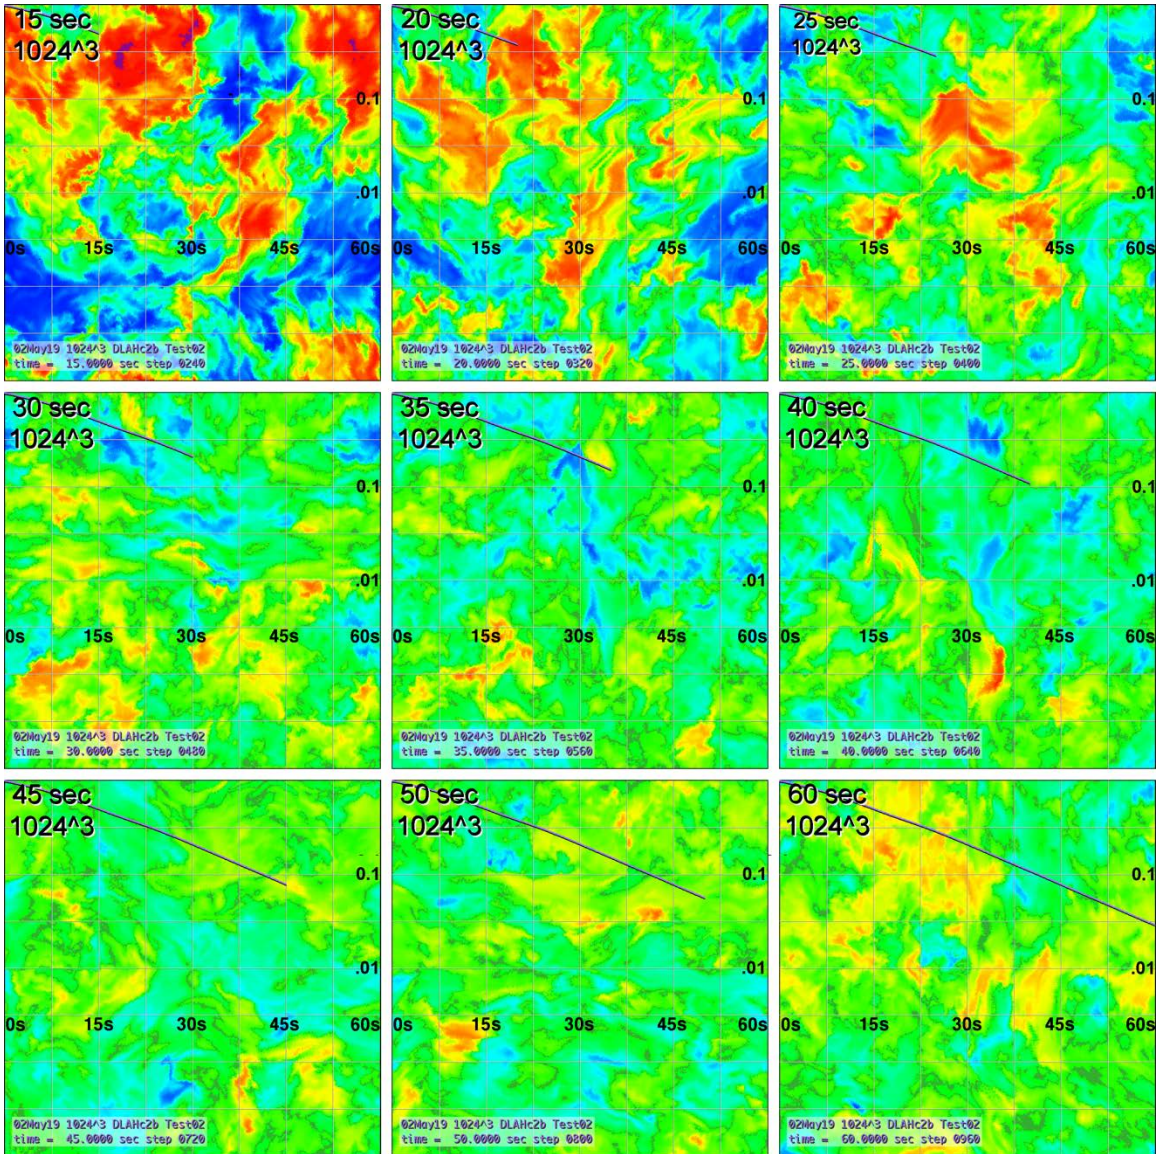


Figure 6.1: Tracer density on a single Z plane ($k = 24$) in fully equilibrated turbulence as a function of time. Initial values (upper left) 1.0 (magenta) and 0.0 (black). Each panel plots the maximum density at that time as red (< 1.0) and the minimum as dark blue (> 0.0). The lavender curve from the upper left on each panel is the Mixed Volume Index, $MVI(t)$. Each horizontal grey fiducial line marks a square root of 10 (3.162...) reduction in MVI.

Figure 6.2 below shows nine plots in which the tracer density has been averaged in the Z direction, normal to the plane of the figure, over the entire periodic cube. The $MVI(t)$ curves in Figs. 6.2 are from the same Scenario 2 data as in Fig. 6.1. The color contours are smoother than in Figs. 5.1 and 6.1, however, due to the average in each X-Y cell over all 1024 of the constant Z planes. What stands out in Fig. 6.2 are the well-defined contours, highlighted by the grey x marks in each panel that identify the $\rho_{ij}(t) = 0.5$ level indicating that the tracer is fully mixed.

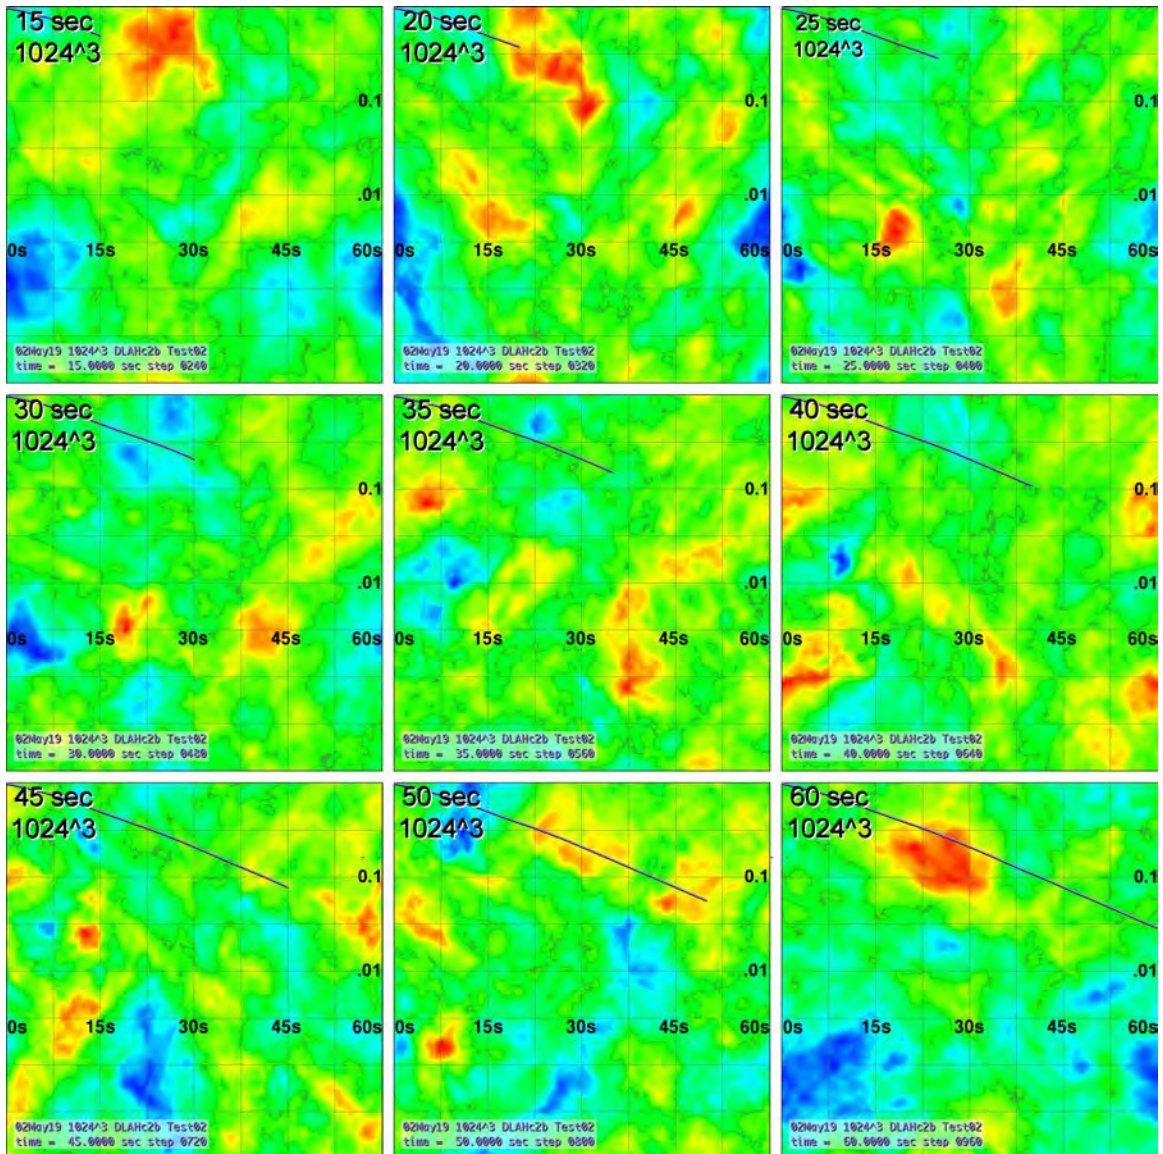


Figure 6.2: Vertically averaged passive tracer density with initial densities 1.0 (magenta) and 0.0 (black), as shown in the upper left panel. Each subsequent panel is normalized using the maximum average density (red) and the minimum average density (dark blue). The lavender curve from the upper left on each panel is the Mixed Volume Index, $MVI(t)$. Each horizontal grey fiducial line marks a square root of 10 (3.162...) reduction in MVI.

Figure 6.2 was computed using a tracer density represented as a single array of 4-byte floating point reals. At the end of the run the fractional non-conservation was 1.4×10^{-6} . The round-off non-conservation affects the estimated cancellation due to mixing. If the average density varies from 0.5 due to non-conservation, mixing would not be able to cancel the two (effective) species out completely. The value of $MVI(t)$ would go to a finite value based on the amount of non-conservation. In Section 7, 8, and 9 to follow, two separate species, a fuel and oxidizer, are treated to allow a surrogate chemical reaction as well as turbulent and 'molecular' mixing. In those sections the species are represented as 16 bit integers. The truncation error is larger but conservation is exact so there is no accumulating round-off to pollute the convection and added diffusion algorithms.




Figure 6.3: Mixing delay caused by population of small turbulence space scales. The lavender MVI curve is from the non-equilibrium evolution of the turbulent spectrum. The white MVI curve began with fully equilibrated turbulence at $t = 0$.

Figure 6.3 above overlays and compares the MVI(t) curves for the two simulated scenarios of Section 5 and Section 6 from $t = 0$ to 60 seconds. The horizontal distance between these curves gives the delay time in seconds for a given level of mixing to have occurred, because the spin up of the non-equilibrium turbulence (lavender curve - Scenario 2) mixes more slowly than when fully equilibrated turbulence acts from the start (white curve - Scenario 1). Again, the lavender curve begins in the upper left with near zero slope because the ‘molecular’ mixing from the added diffusion is restricted to flat interfaces with minimal area and the turbulent spectrum only has energy in the longest wavelengths so the interface area is not increasing quickly. It takes 8 to 10 seconds for the small scale motions to approach equilibrium (lavender curve), at which point the interface begins stretching rapidly.

The white curve, where the simulation begins only after the turbulent spectrum is in equilibrium, has a substantial slope in the upper left of Fig. 6.3 at $t = 0$. Mixing is faster. The

horizontal red-line segments and corresponding labels measure the time delay in reaching different levels of mixing at different elapsed times in the two scenario runs. This difference is recorded in Table 6.1. The elapsed time is the time at the center of each red line segment in Fig. 6.3. This is average between the elapsed times at which the equilibrated turbulence scenario and the non-equilibrium turbulence scenario both reach the indicated level of mixing.

In this pair of runs, the largest mixing delay occurs at about 8.7 seconds when tracer mixing has occurred in only 25% of the volume. The peak velocity in the flow is something over one meter per second, the large-scale stirring velocity, in the equilibrated turbulence. Thus the flow can cross more than half of the system size during this delay and only about 25% of any chemical energy could have been released during this time. Such estimates are only crude approximations. If the energy release were strong, the flow would be changed and mixing would be accelerated.

Elapsed Time	8.7 s	13.8 s	23.9 s	32.1 s	41.7 s	50.9 s	59.2 s
Mixing Delay	3.6 s	2.5 s	1.7 s	1.6 s	1.6 s	1.2 s	1.9 s
Tracer % Mixed	25 %	44 %	68 %	78 %	90 %	94%	97 %

The Table 6.1 shows that 90% of the mixing will have occurred after 41.7 seconds (column 6 of the table). Travelling at 1 m/s, the tracer could cross the system up to four times before small scale mixing will have allowed most of the energy to be released. The table also shows that the mixing delay decreases, at least in this case, after the first few seconds. Because the two scenarios have different flow fields after $t = 0$ seconds, there is no reason to expect that the mixing delay will approach a constant value or even increase or decrease monotonically.

More complete treatment of mixing that leads to chemical reactions requires more than one specie with, treatment of reaction products and consideration of finite-rate reactions. In the next three sections other features of the SFD model for turbulent mixing are demonstrated including convergence with timestep, convergence with grid resolution, and the nature of added numerical diffusion, as a surrogate for unresolved turbulent mixing, acting on the grid-resolved mixing.

7. Convergence of the Mixed Volume Index with Timestep

Sections 7, 8 and 9 consider convergence of Surrogate Fluid Dynamics (SFD) solutions with timestep, with grid resolution, and with the strength of the diffusion that must be added to enable microscale mixing. The underlying scenario is now changed to a two-species, high-speed flow with separate fuel and oxidizer. These species are initially separated into two flat layers rather than the checkerboard layout used in Sections 4, 5 and 6. A reaction algorithm is added so the fuel and oxidizer are forced to react as soon as they occupy the same grid cell. This is a simple approximation to instantaneous reaction. The product does not need to be solved separately because the flow is incompressible. The product mass density is the difference between the initial total density and the sum of the current fuel plus oxidizer densities. The system length L_{sys} is also reduced to 30 cm from the 10 m used in previous sections and the large-scale stirring velocity is 5×10^4 cm/sec rather than 100 cm/sec used previously. This flow is

still subsonic if the gas is hot enough and the turbulent velocities at smaller scale than the stirring velocity are substantially subsonic. Time scales are now in milliseconds (ms) rather than seconds as previously. All of the following simulations were run to 10.0 ms total elapsed time.

In this expanded problem definition, the species densities are now defined as 2-byte scaled integers rather than the single 4-byte floats used for the tracer density in Sections 4 through 6 above. Convection by cyclic shifts can be performed for both species in the same time and computer memory by treating the two two-byte integers as a single word. In Sections 7 through 9 the added diffusion and chemical reaction, which consumes the fuel and oxidizer in equal proportions, is calculated every other timestep to reduce both the residual diffusion and the running time by about 40%. The reaction rate, not varied in this paper, is considered in the next paper in this series. That paper also considers further the effects of the added diffusion strength on convergence in mixing-reaction problems and discusses the issue of measuring the roles of physical/numerical diffusion in turbulent flows characterized by an ensemble of quite different possible solutions which make convergence hard to even define.

All solutions in Sections 7 through 9 use a definition of the Mixed Volume Index to account for separate species, modified from that used earlier (Eq. 5.1). $MVI(t)$ is defined to be

$$MVI(t) \equiv \frac{\sum \rho(t)_{fuel}}{\sum \rho(0)_{fuel}} \quad \text{Eq. (7.1)}$$

where the sums are taken over all N_{xyz}^3 cells in the cube. A separate MVI could be defined for the oxidizer as well but the fuel MVI is used here because a 1% excess of fuel was initialized in this problem to prevent the fuel $MVI(t)$ from going to $-\infty$ on a logarithmic scale. As in the previous sections, the left and/or right edges of the plots are marked with the $MVI(t)$ values 0.1 and 0.01 indicating 90% mixing and 99% mixing respectively.

Figure 7.1 below displays four solutions run to 10 ms using four different timesteps. In each panel the background color contour plot is the Z plane 16 fuel density on the 512^3 grid for the different timestep solutions. These colored snapshots differ only because of the timestep. Panel (a) in Fig. 7.1 was run for 100 timesteps with a step size of 0.1 ms. This is five times longer than the timestep 0.02 ms found to give the fastest fully converged result in panel (c). Since diffusion and reaction are implemented every other timestep in all four runs, panel (a) sports an unacceptable 50-step staircase. The colored cross-section of plane 16 in panel (a) is also very granular and shows relatively little large-scale structure relative to the shorter-timestep runs. The $MVI(t)$ curve also approaches the fully reacted state, characterized by a horizontal line at $MVI(t) = 0.01$, sooner than in panels (b) through (d). This suggests that mixing progresses faster with larger timesteps, even with a fixed spatial resolution.

Panel (b) used $dt = 0.05$ ms, requiring 200 steps to reach 10.0 ms. The staircase is far less noticeable but the solution still deviates noticeably from the shape of $MVI(t)$ in the other panels. The solution with $dt = 0.05$ ms is still not close enough to panels (c) and (d) to be considered converged. The solutions shown in panels (c) and (d), with timesteps of 0.02 ms and 0.01 ms respectively, are close enough to each other to be considered converged. Both the color snapshots of the layer fuel density cross-sections and the shape and location of the $MVI(t)$ curves are essentially the same. The locations of red maxima and blue minima in the low level, residual fuel densities at 10.0 ms are also close. The distribution of the small grey x marks in panel (d) are

more evenly distributed than in panel (c). For the following convergence studies with grid resolution and added diffusion strength, the timestep 0.02 ms is selected to shorten the runs.

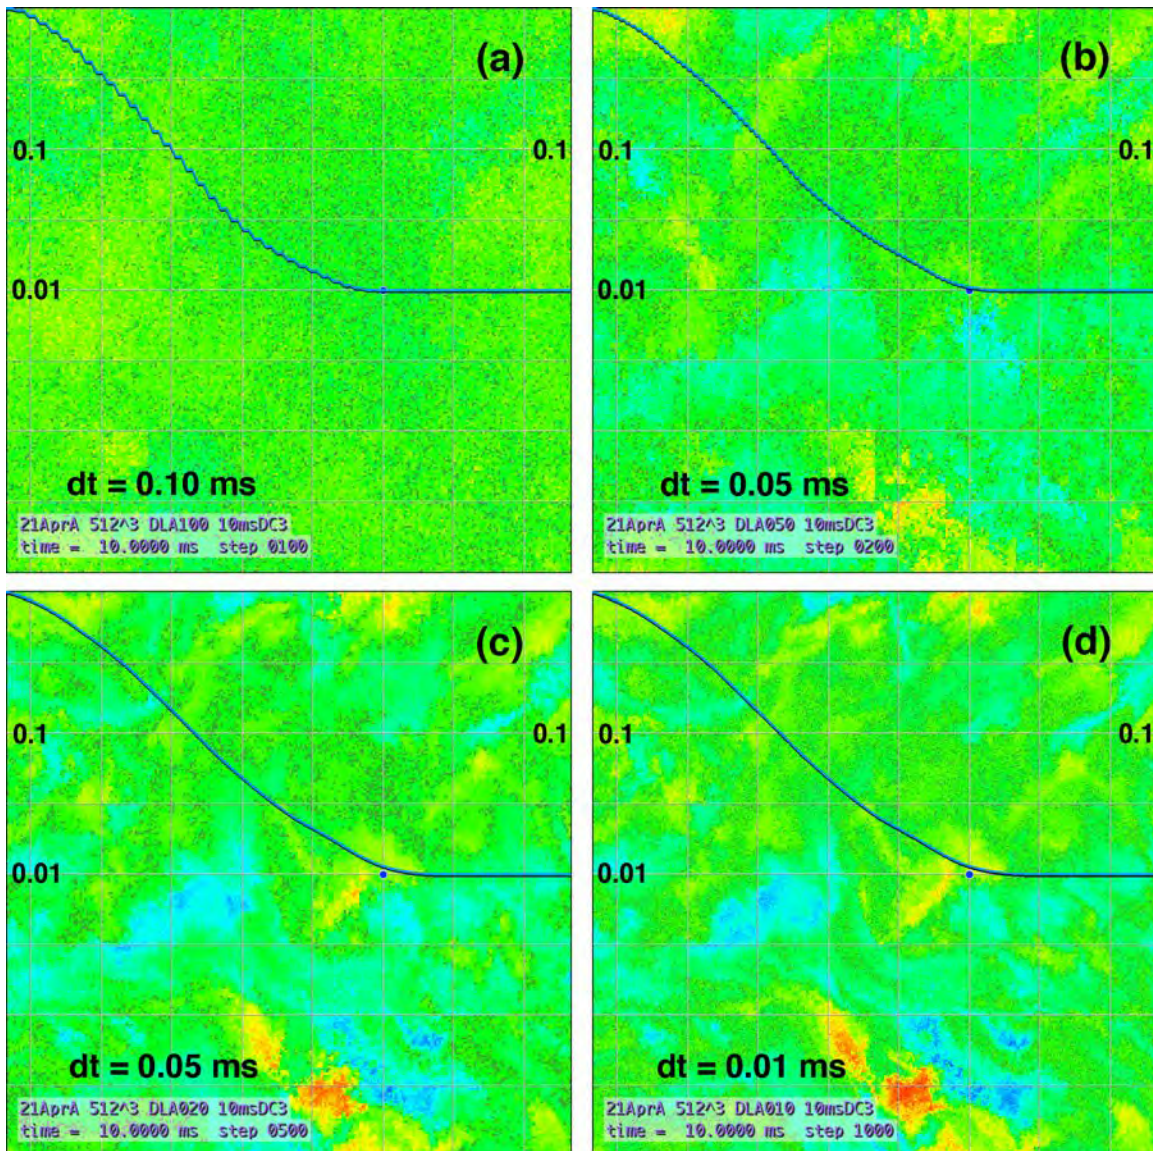


Figure 7.1: Fuel density on layer 16 of a 512^3 grid after 10 ms elapsed time. These panels show simulations with four different timesteps. The maximum value is red and the minimum is blue in each panel. The curve on each panel is the logarithm of the mixed volume index (MVI) as a function of time. Four of the horizontal grey lines span a factor of 100 reduction in fuel density.

In figure 7.2 below, the panel (d) from Fig. 7.1 is used as a basis for overlaying the MVI(t) curves from panels (a), (b), and (c) to compare convergence. Each curve is shown in a different shade of blue and the corresponding timesteps are indicated on the figure. The two longer timestep solutions not only display staircasing from the decision to apply diffusion and reaction every other timestep, they are significantly displaced from the shorter timestep solutions.

In CFD of turbulent flows, and for that matter experiments, the common lore is to use global measures, for example lift and drag on a body in motion, to gauge convergence because global measures are less susceptible to the aleotric fluctuations that affect local metrics when an

ensemble of different realizations characterize the results. There is no converged solution because the flows are intrinsically unsteady. The Mixed Volume Index is such a global measure.




Figure 7.2: Convergence of the Mixed Volume Indices $MVI(t)$ for the fuel density on Z plane 16 from 0.0 ms to 10.0 ms as a function of timestep. Each 512^3 panel is normalized by the maximum plane density (red) and the minimum plane density (blue). The MVI begins at the upper left of each panel and ends at 10.0 ms at the middle right. Four of the horizontal grey fiducial lines span a factor of 100 reduction of the $MVI(t)$ as the oxidizer is fully consumed.

The result of this convergence study is the decision to use $dt = 0.02$ ms as the preferred timestep for this problem. Figure 7.2 shows that the system-averaged metric $MVI(t)$ with a shorter timestep, $dt = 0.01$ ms, is virtually indistinguishable but the calculations take twice as long. The next section considers convergence with grid resolution. This test is made possible in SFD because the time-dependent flow field produced by CSD for SFD is defined for a range of scales several orders of magnitude more resolved than any grid we can afford to use and independent of the actual resolution used. This flow field does not change for the larger scales when additional small scales of the turbulence are being resolved.

8. Convergence of the Mixed Volume Index with Grid Resolution

In the previous section, the four simulations showing timestep convergence were performed using the high-speed flow case and the 512^3 resolution grid. Here the grid resolution is varied from 128^3 up to 1024^3 in powers of 2 using the single timestep 0.02 ms. These four MVI(t) curves are superimposed on the Z-plane 16 fuel density plot for the 512^3 grid. The small grey x marks are placed on the average fuel density contour, which changes in time.




Figure 8.1: Fuel density on Z-layer 16 of the reacting 512^3 fuel-oxidizer simulation at 10 ms elapsed time. This density is color normalized with the maximum layer 16 density as red and the minimum average density as dark blue. The Mixed Volume Index as a function of time $MVI(t)$, for four different grid resolutions, is overlaid starting in the upper left. Four of the horizontal grey fiducial lines span a factor of 100 reduction in $MVI(t)$.

The $MVI(t)$ curves for the 512^3 grid (blue) and the 1024^3 grid (lavender) are basically indistinguishable. The 256^3 result in green is also quite close, suggesting that this smaller grid may be adequate for initial scoping studies. Only the 128^3 MVI metric in yellow deviates appreciably from the two highest resolution results in the figure above.

9. Diffusion in Surrogate Fluid Dynamic Convection

The convection components of SFD have no numerical diffusion – by design. The cyclic-shift operations that implement convection on the triply period cube are fully Lagrangian and fully conservative – also by design. For decades the interaction between real “molecular” diffusion, real turbulent “diffusion” and unavoidable “numerical” diffusion has been confusing and confused. In SFD the only microscale “molecular” mixing that can occur must be added to each species, separate from and independent of the convection. The SFD model implements a small “added” diffusion as a controllable surrogate for molecular diffusion to bring fuel and oxidizer into direct contact. To reduce run times and the accumulated effect of diffusion in SFD, this small added diffusion is imposed on the convected species every other timestep. Since the finest grids we can afford are still too coarse to resolve the scale at which molecular diffusion actually takes place in high Reynold’s number flows, this added diffusion is numerical in origin but necessary to the problems we are interest in. This section shows that the macroscopic mixing-reaction results are insensitive to the amount of added diffusion as long as it is small enough.

The formula used to diffusively spread each species density $\rho_i^n(x, y, z, t)$ on the grid scale is

$$\rho_i^n(x, y, z, t) = \rho_i^o(t) + \frac{(\rho_{i+1}^o - \rho_i^o)}{N_D} - \frac{(\rho_i^o - \rho_{i-1}^o)}{N_D} \quad \text{Eq. 9.1}$$

where N_D is an integer between 2 and ∞ and superscripts n and o indicate ‘new’ and ‘old’ values of the density respectively. When implemented with an integer representation of the density, the telescoping sum on the right ensures conservation to full round-off accuracy. The operational issue is to ensure that the value chosen for N_D does not cause macroscopic solutions on the resolved grid to depend significantly on the grid resolution due to numerical error. Studying these errors is important and can be done within SFD because the numerical diffusion associated with conventional CFD is absent from the SFD model. Added diffusion can be turned on and off.

Figure 9.1 below shows plots of the fuel density on Z layer 32 of the 1024³ grid for several runs varying the amount of added diffusion. MVI(t) is also shown, averaged throughout the cube, from runs to 10 ms in which the added diffusion is only turned on at the 250th timestep. The simulations performed in this section use the standard high-speed scenario introduced in Sections 7 and 8 run with $N_D = 5$ and $dt = 0.02$ ms. Panel (a) shows the fuel density in this 32nd layer at 4.96 ms (step 248) just before the added diffusion is turned on. The image shows an apparently random distribution of magenta pixels, each indicating a density of 2.0, on a light blue background, where each pixel indicates 0.0 fuel density. Using black rather than blue to indicate zero density in this case would obscure the magenta pixels, which account for about 25% of the image in this problem. Panels (b), (c), and (d) are all plotted at 10 ms after the added diffusion has thoroughly mixed the two species and the fast reaction algorithm has consumed all of the oxidizer.

Panel (b) below was plotted from the layer 32 fuel density with an added diffusion 1.67 stronger than standard, that is $N_D = 3$ rather than $N_D = 5$. This results in the leftmost dark blue MVI(t) curve. The background image in panel (b) shows the distribution of the 1% residual fuel, which is almost uniformly distributed in the cube once mixing and reaction consumes all the oxidizer. Just to the right of the blue curve is the lavender MVI(t) curve for the standard case with $N_D = 5$. The curve is labelled diffusion = 1.0. The white MVI curve on panels (b), (c) and (d)

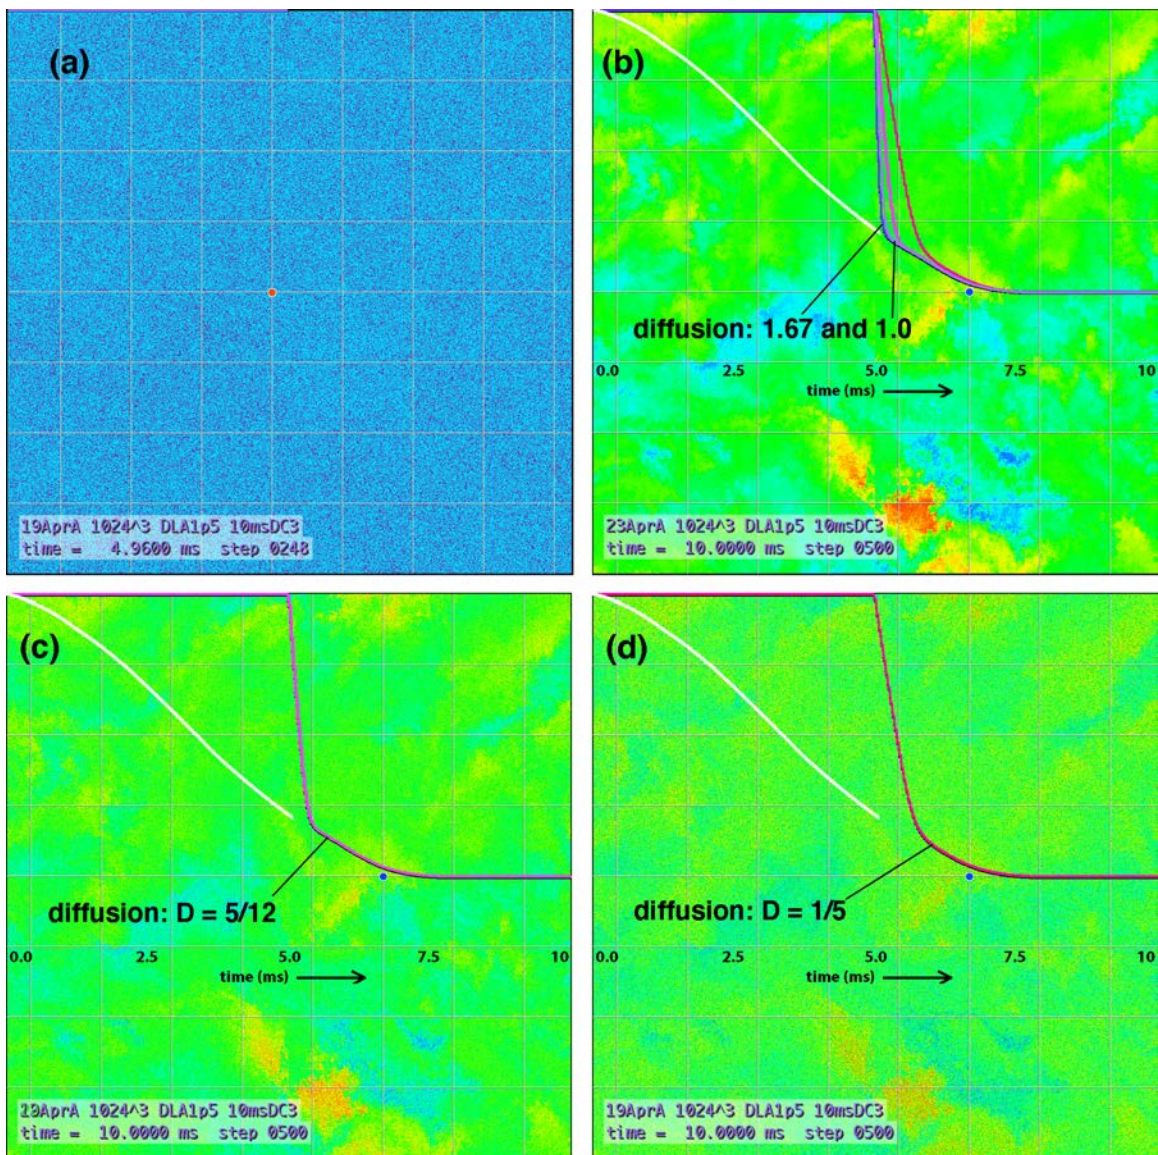


Figure 9.1: Fuel density on Z layer 32 of the 1024^3 grid using $dt = 0.02$ ms. “Added” diffusion is turned on at step 250, halfway through these reacting fuel-oxidizer simulations. (a) No mixing whatsoever, $D = 0$ at step 248. (b) Added diffusion $D = 1.67$, was turned on at step 250 (5.0 ms) and compares with $MVI(t)$ for $D = 1$, $D = 5/12$, and $D = 1/5$. (c) Reduced diffusion: $D = 5/12$ standard. (d) D reduced further: $D = 1/5^{\text{th}}$ standard case. Panels (c) and (d) have their full $MVI(t)$ curves overlaid for comparison on panel (b). Four horizontal grey fiducial lines define a factor of 100 reduction in MVI .

is from standard scenario on the 1024^3 grid in which the added diffusion is active throughout the run so reaction begins immediately.

This added diffusion is small in the sense that it does not actually transport the species more than a cell or two before the fast reaction converts the interpenetrated fuel and oxidizer to product. The difference between the two left most $MVI(t)$ curves, $N_D = 3$ and $N_D = 5$, is almost invisible because the added diffusion is strong enough that the transition from the unmixed but coarse-grained interpenetration in (a) to the “molecularly” mixed and reacted state occurs in a few steps. The diffusion is also small enough that the grid-resolved solution is not changed.

Panel (c) shows a case where the added diffusion is reduced to 5/12 the standard case, i.e. $N_D = 12$. The MVI(t) curve is superimposed on panel (b) as the lighter lavender curve second from the right. The main effect of the reduced diffusion is to delay the transition from the unmixed state before step 250 to the mixed reacted state that would have prevailed if the added diffusion were turned on throughout the run. Although the rate of approach to the standard curve is slowed somewhat in panel (c), the end result is the same. The amount of diffusion has no great effect on the completion of the reaction or the end state because the diffusion really does not move the fuel or oxidizer more than a few cells over the course of the run.

Panel (d) reduces the added diffusion by a factor of five smaller than the standard case, i.e. $N_D = 25$. MVI(t) from panel (d) is the rightmost red curve in panel (b). The transition to the mixed fully reacted state is now much slower. At the end of the run, the density plot in panel (c) is much grainier than the one in panel (b) and the panel (d) snapshot is still grainier than (c). As the diffusion coefficient approaches zero, the complete separation of fuel and oxidizer shown in panel (a) can leave its mark in the density fields because adjacent cells can still have quite different values. The same behavior would appear in a Lagrangian CFD model as long as the grid did not have to be restructured. During restructuring for truncation accuracy, there would be unavoidable numerical diffusion from interpolation to a more regular grid.

10. Conclusion and Discussion

A Surrogate Fluid Dynamics (SFD) model was introduced to provide 3D realizations of idealized flow fields consistent with any given, possibly time-dependent, turbulence spectrum. In this paper the Coherent Structure Dynamics (CSD) model was used to simulate the evolution a non-equilibrium, high Reynolds-number turbulent spectrum, energized or “stirred” at the system scale, as it equilibrates to a Kolmogorov spectrum covering 7 or 8 orders-of-magnitude in scale lengths. The grids that SFD can afford, up to 1024^3 running on a laptop, reach well into the inertial range of the spectrum but still nowhere near the dissipation scale. Although a reduced order model, CSD replicates the “bottleneck”, or “pre-dissipation bump” observed in turbulence quite well (Boris, 2018) while reflecting the main conservation and fluid-dynamic scaling properties usually associated with turbulence.

SFD’s design removes numerical diffusion from the convection algorithm. A series of integral-cell, Lagrangian shifts of the density array along the coordinate axes approximates the 3D flow field without any interpolation leading to diffusion of density values. This approach has the added benefit that additional small-scale motions, if the resolution is increased, do not change the resolved-scale solution for a realization. The underlying flow is derived entirely from a higher-resolution spectrum provided by CSD, regardless of the grid resolution. Several convergence studies were conducted here to illustrate these properties of this SFD model.

SFD is fundamentally Lagrangian. Lagrangian models have been applied to complex flows for both compressible and incompressible Computational Fluid Dynamics. As long as the grid motions are entirely Lagrangian and the grid does not need to be restructured to alleviate excess distortion, the convection can be diffusion free. However, numerical diffusion is not absent from practical Lagrangian methods in multidimensions because realistic flows with vortices distort Lagrangian grids so much that remapping onto a more regular grid must be done many times in a run. Interpolating the physical properties attached to each Lagrangian node onto a remapped grid introduces diffusion. Further, even when the mass and species densities are completely

Lagrangian, the energy density in a compressible flow on the same grid cannot be Lagrangian. Energy density is convected at a different velocity from the mass and species densities and thus will have numerical diffusion.

Lagrangian fluid dynamics methods are discussed in Chapter 9 of “Numerical Simulation of Reactive Flow” (Oran and Boris, 2001). Lagrangian methods may use a grid that moves and deforms with the flow or they may be based on grid-free representations where the fluid dynamic variables are assigned to nodes that move independently in the flow. More recent reviews by Premoze, et al. (2003), Zhang and Chen (2007) and Lind, et al. (2020) broaden the discussion of Lagrangian numerical methods and provide a number of additional references. The grid-free methods, sometimes called free Lagrange methods or particle methods, would seem to offer the possibility of removing the numerical diffusion from regridding. Smoothed Particle Hydrodynamics is a well-developed approach for compressible flows (Lucy, 1977; Monaghan, 1992 & 1994). Another grid free particle method called the Moving-Particle Semi-Implicit (MPS) method was developed for incompressible fluids (Koshizuka, et al., 1996).

Despite the potential benefits of the Lagrangian representation, the real-world applications have generally been disappointing for turbulent flow simulations which must treat a number of eddy turn-over times on many overlapping spatial scales. Evaluating local values of species densities and other fluid variables to compute derivatives usually requires interpolation which introduces numerical diffusion in addition to truncation errors. Further, tracking the nearest neighbors of Lagrangian cells or nodes to evaluate derivatives is an ‘ N^2 problem’ that presents additional computational costs. Further, the non-uniformity of the neighborhoods degrades accuracy relative to structured or even uniform grids. SFD, however, limited in its fidelity to the Navier-Stokes equations, nevertheless allows study of important aspects of turbulent mixing in numerical computations including convergence and the interplay of diffusion, mixing, and chemical reaction.

Surrogate Fluid Dynamics affords the opportunity to study both an ensemble of distinct 3D flow realizations, all satisfying exactly the same turbulence spectrum time history, and the convergence of flow metrics for such an ensemble. This convergence study, however, is left to another paper.

Acknowledgements:

The research leading to Coherent Structure Dynamics and these preliminary demonstrations was supported by the projects on “Highly Complex Fluid Dynamics” and “Advanced Computational Models that Exploit Many-Core Computer Architectures” within the ONR/NRL 6.1 basic research core program. The DoD High-Performance Computing Modernization Office supported the development of the FAST3D and FAST3D-CT Large-Eddy Simulation models that provided the empirical incentive and projected need for a non-equilibrium model like CSD. My NRL collaborators, Gopal Patnaik, Keith Obenschain, Douglas Schwer, and David Mott contributed their expertise through numerous focusing conversations. Professors Elaine Oran (Texas A & M), Professor Alexei Poludnenko (University of Connecticut), Dr. Fernando Grinstein (Los Alamos National Laboratory), and Professor Paul Woodward (University of Minnesota) all have contributed significantly to the underpinnings of this work in recent years.

References:

- G.L. Brown and A. Roshko, 1974, "On density effects and large structure in turbulent mixing layers," *J. Fluid Mechanics* **64** (4), pp 775–816, 1974.
- J.P. Boris, 2018, "A Coherent Structure Dynamics Model for Non-Equilibrium Turbulence," U.S. Naval Research Laboratory Memorandum Report NRL/MR/6003--18-9815, 14 September 2018.
- L. Biferale, 2003, "Shell Models of Energy Cascade in Turbulence," *Annual Reviews of Fluid Mechanics* **35**, pp 441–68, 2003. doi: 10.1146/annurev.fluid.35.101101.161122
- F.F. Grinstein, 2001, "Vortex Dynamics and Entrainment in Rectangular Free Jets," *Journal of Fluid Mechanics* **437**, pp 69–101, 2001.
- P.R. Woodward, D.H. Porter, I. Sytine, S.E. Anderson, A.A. Mirin, B.C. Curtis, R.H. Cohen, W.P. Dannevik, A.M. Dimits, D.E. Eliason, K.-H. Winkler, and S.W. Hodson, 2001, "Very High-Resolution Simulations of Compressible Turbulent Flows," *Computational Fluid Dynamics, Proceedings of the Fourth UNAM Supercomputing Conference Mexico City, June 2000*, edited by E. Ramos, G. Cisneros, R. Fernandez-Flores, and A. Santillan-Gonzalez, pp 3–15 *World Scientific*, 2001.
- P. Moin and K. Mahesh, 1998, "DIRECT NUMERICAL SIMULATION: A Tool in Turbulence Research," *Annual Review of Fluid Mechanics* **30**, pp 539–578, 1998.
- R.D. Moser, J. Kim, and N.N. Mansour, 1999, "Direct numerical simulation of turbulent channel flow up to Re 590," *Physics of Fluids* **11-4**, pp 943–945, 1999.
- A.J. Kajzer, J. Pozorski, and K. Szewc, 2014, "Large-eddy simulations of 3D Taylor-Green vortex: comparison of Smoothed Particle Hydrodynamics, Lattice Boltzmann and Finite Volume methods," *Journal of Physics Conference Series* **530**, pp 1–8, 012019, 2014.
- O.M. Knio and A.F. Ghoniem, 1990, "Numerical Study of a Three-Dimensional Vortex Method," *Journal of Computational Physics* **86**, pp 75-106, 1990.
- D.I. Pullin and P.G. Saffman, 1998, "Vortex Dynamics in Turbulence," *Ann. Rev. Fluid Mechanics* **30**, pp 31–51, 1998.
- J.G. Brasseur and S. Corrsin, 1987, "Spectral Evolution of the Navier-Stokes Equations for Low Order Couplings of Fourier Modes," in **Advances in Turbulence**, pp 152–162, edited by G. Compte-Bellot and J. Mathieu (Springer-Verlag, Heidelberg, 1987,).
- J.G. Brasseur, 1991, "Comments on the Kolmogorov Hypothesis of Isotropy in the Small Scales", AIAA Paper 91–0230, 1991.
- J.G. Brasseur and C-H. Wei, 1994, "Interscale dynamics and local isotropy in high Reynolds number turbulence within triadic interactions," *Phys. Fluids* **6** (2), pp 842–870, February 1994.
- D.R. Mott, "New Quasi-steady-state and Partial-equilibrium Methods for Integrating Chemically Reacting Systems," PhD dissertation, University of Michigan, Ann Arbor, 1999.
- D.R. Mott and E.S. Oran, 2001, "CHEMEQ2: A Solver for the Stiff Ordinary Differential Equations of Chemical Kinetics," Naval Research Laboratory Memorandum Report 6400--01-8553, 27 July 2001.
- T.R. Young, Jr. and J.P. Boris, 1973, "A Numerical Technique for Solving Stiff Ordinary Differential Equations Associated with Reactive Flow Problems", Naval Research Laboratory Memorandum Report 2611, July 1973.
- T.R. Young, Jr. and J.P. Boris, 1977, "A Numerical Technique for Solving Ordinary Differential

Equations Associated with the Chemical Kinetics of Reactive-Flow Problems." *J. Physical Chemistry* **81**, pp 2424–2427, 1977.

T.R. Young, Jr., 1980, "CHEMEQ – A Subroutine for Solving Stiff Ordinary Differential Equations," Naval Research Laboratory Memorandum Report 4091, 1980.

Oran, E.S. and J.P. Boris, 1993, "Computing Turbulent Shear Flows – A Convenient Conspiracy," *Computers in Physics* **7** (5), pp 523–533, 1993.

E.S. Oran and J.P. Boris, 2001, **Numerical Simulation of Reactive Flow**, 2nd edition, (Cambridge University Press, New York, 2001).

S. Premože, T. Tasdizen, J. Bigler, A. Lefohn and R.T. Whitaker, 2003, "Particle-Based Simulation of Fluids", EUROGRAPHICS 2003, P. Brunet and D. Fellner (guest eds), Volume 22 Number 3, 2003.

Z. Zhang and Q. Chen, 2007, "Comparison of the Eulerian and Lagrangian methods for predicting particle transport in enclosed spaces", *Atmospheric Environment* **41**(25), pp 523–5248.

S.J. Lind, B.D. Rogers and P.K. Stansby, 2020, "Review of smoothed particle hydrodynamics: towards converged Lagrangian flow modelling," *Proc. R. Soc. A* 476: 20190801, 22 July 2020.

L.B. Lucy, 1977, "A numerical approach to the testing of the fission hypothesis," *The Astronomical Journal* **82** (12), pp 1013–1024, Dec 1977.

J. J. Monaghan (1992) "Smoothed Particle Hydrodynamics", *Ann. Rev. Astron. Astrophys.* 30(2), 543–574, 1992.

J. J. Monaghan (1994) "Simulating free surface flows with SPH," *Journal of Computational Physics* **110**, pp 399–406, 1994.

S. Koshizuka, H. Tamako, and Y. Oka, 1996, "A particle method for incompressible viscous flow with fluid fragmentation", *Comput. Fluid Dynamics J.*, **29**(4), 1996.