Naval Research Laboratory
Washington, DC 20375-5320

NRL/6003/MR—2022/1

Mixing Delays in Non-Equilibrium High-Speed Turbulence

JAY BORIS

Chief Scientist
Material Science and Component Technology Directorate

KEITH OBENSCHAIN

Head, Laboratory for Advanced Computational Physics
Laboratories for Computational Physics and Fluid Dynamics

February 9, 2022

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
09-02-2022	NRL Memorandum Report	09-14-2019 – 09-20-2021

4. TITLE AND SUBTITLE

Mixing Delays in Non-Equilibrium High-Speed Turbulence

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Jay Boris and Keith Obenschain

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER
4464

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--2022/1

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 third NRL Memorandum Report in a series. There will be five. The fourth is being worked on now, an extension of this CSD3 memo report.

14. ABSTRACT

This paper uses NRL's Coherent Structure Dynamics (CSD) model coupled to Surrogate Fluid Dynamics (SFD) on a triply-periodic, cubical domain to estimate the additional mixing time delays that occur in separated fuel and oxidizer systems when turbulence is quiescent initially. Mixing delays occur as the turbulent energy cascades dynamically from macroscopic driving flows to the much smaller scales where molecular mixing happens. If the turbulence is fully excited initially, the molecular mixing still takes some time but it can begin much earlier.

CSD-SFD is an incompressible model in which the computations are at least three orders of magnitude faster than in Computational Fluid Dynamics (CFD) simulations of comparable resolution. Further the SFD algorithms are implemented with no numerical diffusion and no Courant stability limit so timesteps can be appreciably longer. Added diffusion must be explicitly supplied at the grid scale to ensure fuel and oxidizer mix locally. Multiple realizations of the 3D SFD flow field can be constructed, each obeying the single time-dependent, multi-scale spectrum computed by the CSD model.

In this paper, matching pairs of 3D simulations are used, one starts with a laminar large-scale "stirring" flow where turbulence takes time to develop while the second simulation begins with a fully equilibrated turbulent spectrum. The elapsed time at which specific levels of mixing occur in each simulation is differenced to compute the non-equilibrium mixing delay. This is done for fast flows in this paper to estimate how much the reaction process will be delayed, for example,in a supersonic combustor. In a Mach 3 flow at 1 km/sec, velocities driving the turbulence of .5 km/sec, .3 km/sec, and .2 km/sec are considered. The extra distances eeded for combustion are 1.0, 1.6, and 2.4 meters respectively when the initial turbulence is far from equilibrium.

15. SUBJECT TERMS

Computational fluid dynamics; Coherent structure dynamics; Non-equilibrium turbulence; Turbulent spectrum; Surrogate fluid dynamics; High-speed mixing delay

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
					Jay Boris
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**	U	20	**19b. TELEPHONE NUMBER** *(include area code)*
U	U	U			(202) 767-3055

This page intentionally left blank.

1. Introduction to Coherent Structure Dynamics

Coherent Structure Dynamics (CSD) is a time-dependent model predicting turbulent cascade in a volume of fluid at all scales from the system size down through the viscous dissipation scale. A set of stiff ordinary differential equations is solved for the strength and number of "rotors," so named to describe the coherent, vortex-like structures that cohabit the fluid volume [Brown & Roshko, 1974]. The turbulent flow in a volume is represented as a one dimensional grid of rotor sizes, rather like the different wavelength bins of a Fourier spectrum, but here the rotor bins have logarithmically-sized scales spanning over 7 orders of magnitude. CSD is a time-dependent generalization of a "shell" model (Biferale, 2003). It uses a logarithmic grid to increase the range of space scales that can be represented to include high Reynolds Number flows. CSD is described in detail in an NRL Memorandum Report that includes the equations being solved, a discussion of the algorithms, and the results of a number of tests (Boris, 2018, "A Coherent Structure Dynamics Model for Non-Equilibrium Turbulence," reference CSD1). CSD calculates the evolution of the turbulent spectrum for minutes of real time in a few seconds of laptop computing over a much wider range of space scales than can be treated by Computational Fluid Dynamics running for days on highly parallel mainframes. CSD predicts the turbulent pre-dissipation bump (e.g., Bak and Kalmar-Nagy, 2018), sometimes called the bottleneck, as shown in Figs. 1 and 2 below, and determines the manner and time scale of the nonlinear relaxation to a turbulent equilibrium.

The CHEMEQ2 stiff chemical equation integrator (Mott and Oran, 2001) is used to solve the CSD equations. Figure 1 shows the profiles of the fluid variables for a model problem from reference CSD1. These profiles are captured at four times during the simulation using a timestep $\delta t = 0.1$ second. The fluid in this problem was driven by energizing the rotors at the 10-meter scale with a rotational velocity of 2 m/s. The initial condition in panel a has energy only in the 10-meter scale. By 5 seconds (panel b), the energy has cascaded by a factor of 100 to the 10-cm scale (lavender/purple curve = rotor velocity and red/orange curve = energy). A precursor foot has cascaded another factor of 100 to the 1 mm rotor scale. By 10 seconds (panel c), the precursor foot has dissipated but the equilibrium, power-law spectrum, marked by the straight lines on these log-log plots, is just beginning to emerge although the smaller scales are still severely under populated. By 30 seconds (panel d), the rotor velocities $V_k(t)$, the rotor energies $E_k(t)$ and the scaled Kolmogorov energy $\varepsilon_k(t)$ have closely approached the expected equilibrium spectrum. This fourth panel also clearly shows the spectrum extending into and past the viscous dissipation scales between 10^{-3} and 10^{-4} cm for this case. The Reynolds number for this case is $Re \sim 1.4 \times 10^6$.

Please notice that, after 30 seconds, the rotor packing fraction $P^f_k(t)$ is still far from its equilibrium structure as approximated by the curve of light blue data points in each of the panels. The packing fraction, which describes how close together rotors of a given size are to others of the same size in the volume, takes about 300 seconds in this particular case to reach equilibrium. The power-law Kolmogorov spectrum is dictated by a dimensional scaling argument at steady state and establishes itself very quickly. The packing fraction, however, is a non-dimensional quantity and is thus free to take a much longer time to settle down. Because CSD treats the number density of rotors at each scale as well as their energy, the model is free to exhibit more complex behaviors than dictated simply by dimensional scaling. In particular, the often-discussed pre-dissipation bump or bottleneck is predicted and is circled in red on panel d.

Manuscript approved January 28, 2022.

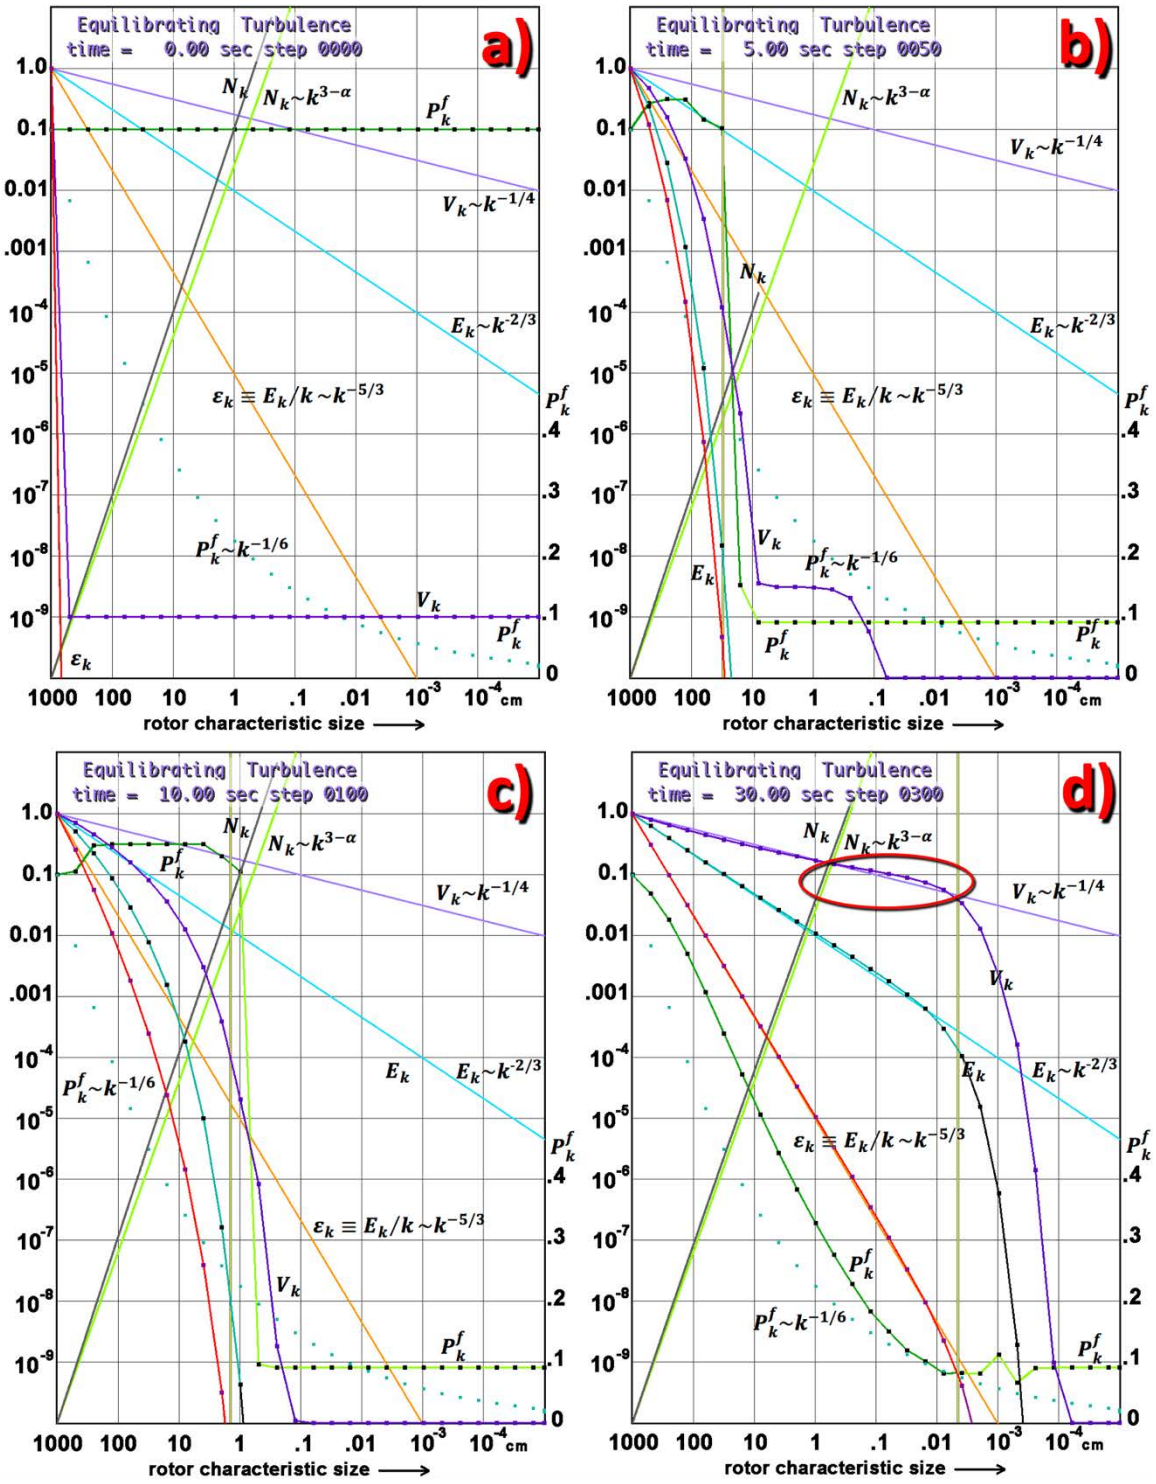

Figure 1. Spin-up of a stirred laminar flow to fully developed equilibrium turbulence in about 30 seconds. The stirring velocity, at the longest scale, 10 meters, was 2 m/s. 30 bins were used representing turbulent scale sizes spaced logarithmically by a factor of 2. Panel a) t = 0 sec; panel b) t = 5 sec; panel c) t = 10 sec; panel d) t = 30 sec, nearly equilibrium. This example is not one of the high-speed cases being analyzed in this paper.

Although the CSD model makes the nonequilibrium, time-dependent spectrum of the turbulence available as a resolved time sequence of rotor energies and number densities, we still need a model and associated algorithms to provide high-resolution, time-dependent flow fields

that conform to the given spectrum at each timestep while obeying some reasonable physical constraints such as conservation, continuity, and the physical boundary conditions. The next section presents a short overview of the Surrogate Fluid Dynamics (SFD) model, developed to provide this turbulent flow field.

CSD is a volume averaged treatment of the interaction and dynamics of these coherent structures to expose important but general behaviors but without a concrete notion of geometry. By itself, CSD does not compute one particular deterministic turbulent realization, but rather expresses the coupling of turbulent kinetic energy between the different spatial scales, i.e. "shells," spanning the turbulence spectrum. Specific spatially and temporally correlated realizations must be determined after CSD has provided the non-equilibrium, time-varying ensemble energy spectrum. The next section describes a new, Surrogate Fluid Dynamics (SFD) model to accomplish this. The composite model will be denoted as CSD-SFD. The description "surrogate" reminds us that the solutions are not detailed solutions of the Euler or Navier-Stokes equations but rather capture the main scale dependent behavior of turbulent flows necessary to approximate realistic mixing between separate chemical species.

2. Surrogate Fluid Dynamics – a Dissipation-Free Mixing Model

Turbulence continually convolutes and stretches the interface initially separating two species. However, actual mixing primarily occurs within a few molecular mean-free-paths of this evolving interface where the density gradients are largest. This means that only a small fraction of the volume becomes well mixed until the diffusing, convoluting interface has moved close to most locations in the volume. Using a grid-based flow algorithm to solve the species continuity equations requires very high resolution to keep the species as physically separated as they should be. Unless the CFD cell size is smaller than the density gradients that contribute most heavily to molecular mixing, resolution-based numerical errors will give the appearance of mixing two species prematurely that should still be relatively unmixed. Interface tracking approaches offer little help when we need to resolve an interface that becomes essentially fractal in nature with a rapidly growing area and rapidly increasing curvature. Surrogate Fluid Dynamics (SFD) is designed to provide an approximate solution to this problem.

Detailed numerical simulation would not be practical for these cases. Instead, the CSD model is being used here to drive a surrogate flow field that causes the mixing of initially separated species in the context of a high Reynold's-number, turbulent flow. To convect separated species without non-physical numerical diffusion, millions of cells would be needed in each dimension. In SFD, a 3D, axis-rectified flow model is defined that is driven by the time-dependent turbulent spectrum (rotor velocities) in each of the CSD bins. To maintain the species separation for as long as possible, the purely convective aspect of turbulent flow is represented by shifting the density profile an integral number of cells in X, Y and Z parallel to the coordinate axes of the grid at each timestep. This is mimicking convection in an idealized Lagrangian way but without any numerical diffusion or diffusive regridding. (Boris, 2021, "A Mixing Study Using Coherent Structure Dynamics to Drive a Surrogate Fluid Dynamics Model," reference CSD2). Adjusting the X, Y, and Z shifts at all of the resolved scales is done very efficiently to impress the evolving turbulence spectrum of velocities in a correlated, continuously varying manner.

This SFD is a *surrogate* convection model where the shifts parallel to the X, Y, and Z axes only move the fluid in each line an integral number of cells during each timestep. These shift

operations 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 of the CSD spectrum at the time the shift is made. In 2D this has sometimes been called a "cookie-cutter model." In 3D the relative shearing of shifts between adjacent lines generates density gradients and encompasses vorticity. 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. If a fractional cell shift is required, it is saved and accumulated until a later timestep requires a shift greater than a cell. There are, of course, accuracy and convergence concerns but there is no Courant stability condition. Therefore, shifts of several cells can be used in each direction without stability problems, making for a very efficient algorithm.

Each species density is defined on a regular, triply periodic grid. Four grids can be used to vary resolution, 128^3, 256^3, 512^3, and 1024^3. The SFD species densities at the finest grid can be advanced using the CSD-determined turbulent flow for hundreds of timesteps in an hour or so on a single MacBook Pro but this gridded density field represents only the longest wavelengths of the 30 factor-of-two CSD bins used to represent the turbulence evolution. Even the highest-resolution 1024^3 representation does not actually penetrate the viscous-dissipation range of the driving flow for Reynold's numbers greater than 10 to 100 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 scenarios, as demonstrated by convergence studies in reference CSD2.

The "Lagrangian" periodic cyclic shifts that capture the convection in SFD are performed simultaneously at many different scales down to the cell size, thus orders of magnitude in turbulence scales are being represented without numerical diffusion in this surrogate formalism. Mixing is forced to occur only at the grid scale by an explicit diffusion operation. Only this added diffusion mixes the different tracer density values representing the two species. An illustrative example of this non-diffusive convection is illustrated in the next section as the problem initial conditions are described.

In the scenarios discussed in this paper, the mixing of two species, fuel and oxidizer, is considered. The viscosity is given the value 1.18 cm^2/sec, approximating a value for high temperature air. The system size is L_{sys} = 30 cm, and the stirring velocity $V_{stir} = 5x10^4$ cm/sec. Thus, the Reynold's number is ~$1.4x10^6$. The domain of the simulations is intended as only a portion of the potentially reacting flow. Figure 2 immediately below shows the equilibrated turbulence spectrum for these high-speed conditions. This spectrum is very similar in shape to panel d) of Fig. 1 and also displays the pre-dissipation bump prominently. At longer times, such as 10 ms in the figure, the packing fraction P^f_k has had a chance to approach an approximate power-law structure more closely than in panel d) of Fig. 1. The two vertical yellow lines in the figure mark the beginning of the viscous dissipation zone. The cell size in the 1024^3 grid is about 0.03 cm and thus extends into the pre-dissipation bump but not into the viscous dissipation zone.

Below, in Section 4, a convergence study with resolution and timestep, relevant to the high-speed flow mixing of this report, is presented based on the spectrum of Fig. 2 and its time-dependent evolution leading up to the equilibrated condition shown in the figure. The bottleneck or pre-dissociative bump may actually be slightly exaggerated by parameters chosen for the CSD model (Poludnenko and Colby, 2021).

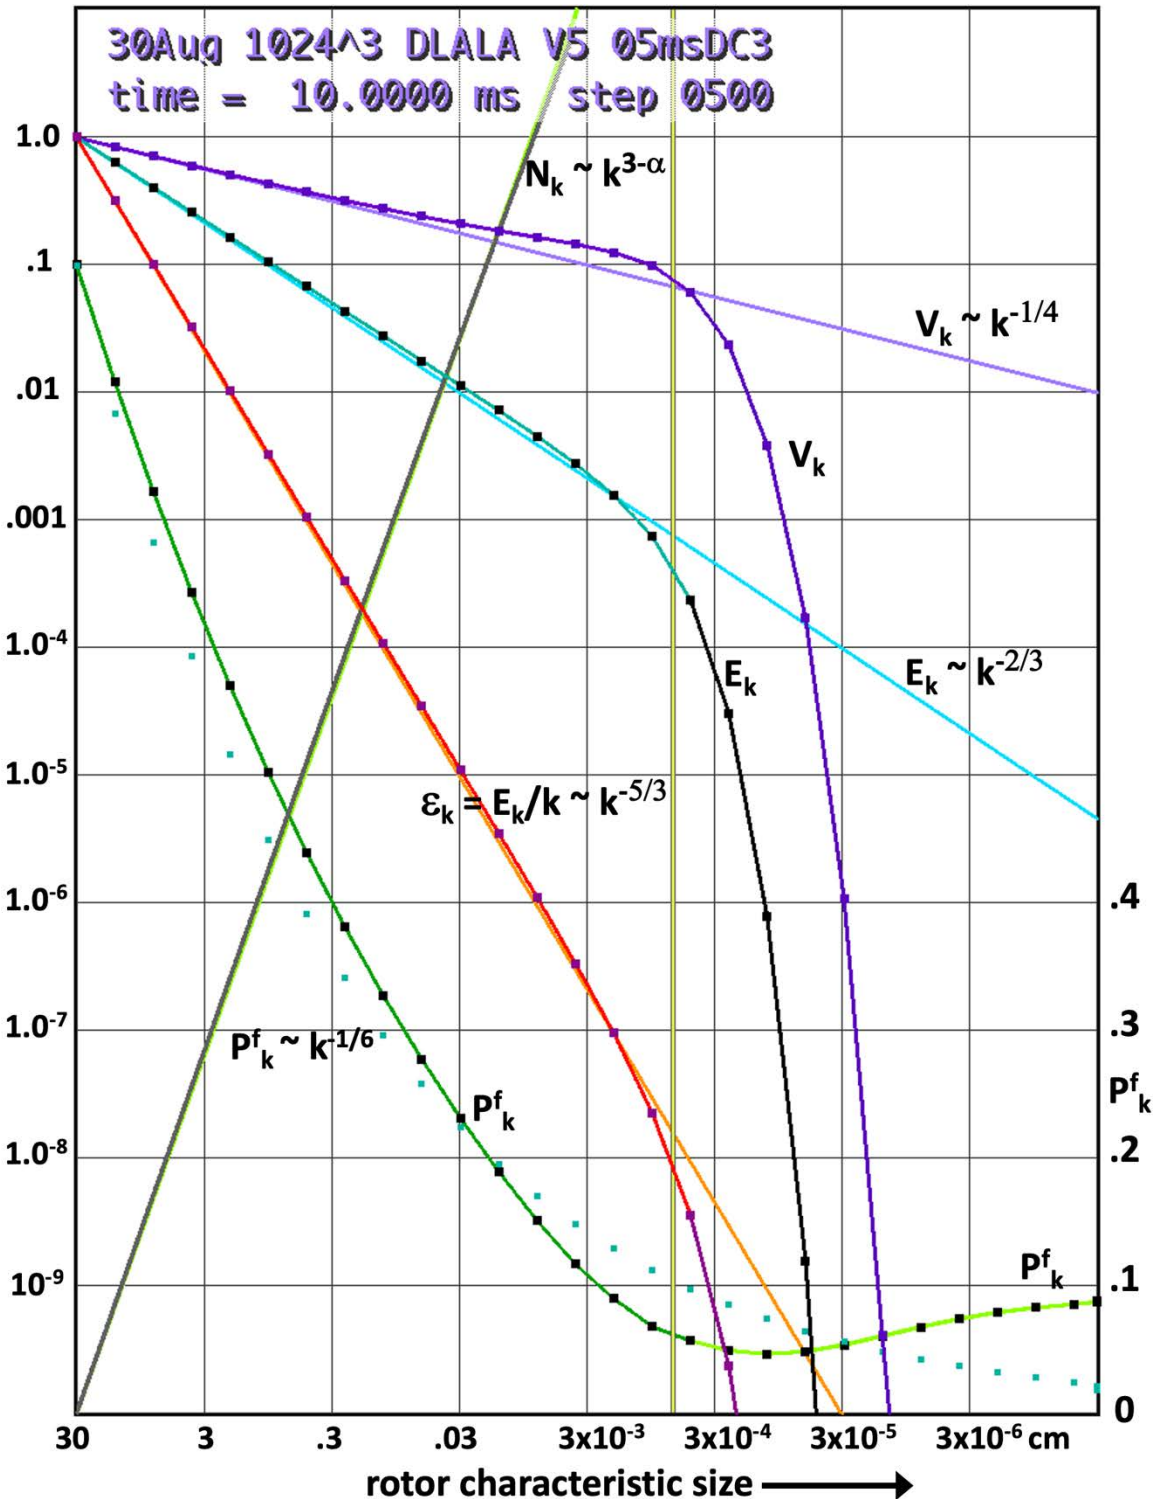

Figure 2. Spin-up of a stirred laminar flow to equilibrium turbulence in less than 10 milliseconds for parameters of this paper. The stirring velocity, at the longest scale, 30 cm, was 500 m/s. 30 CSD bins were used representing turbulent scale sizes spaced logarithmically by a factor of 2.

The following section discusses the models added to CSD-SFD to represent an idealized fuel oxidizer system and to implement the rapid reaction of the species only they are molecularly mixed.

3. High-Speed CSD-SFD Mixing with a Fuel-Oxidizer Reaction

The left-hand image of Figure 3.1 below shows the initial conditions looking down on a 1024^3 cubical domain in the -Z direction. In plots of all such cross-sections, X is to the right and Y is upward as shown in this figure. A flat layer of fuel of non-dimensional density 2.0 and Y extent 7.5 cm is shown as magenta at time t = 0.0 ms. The cube, regardless of resolution, is periodic so this fuel layer appears split in half in the figure. The initially separate oxidizer, a distinct layer extending 22.5 cm in Y in the center of the figure, is shown as black. Two different initial oxidizer values were used, 0.6 (slightly fuel rich) and 0.7 (slightly fuel lean). The left hand image plots the 1024^2 fuel density values in the single Z cross-section at k = 24. These simulations are all initialized with only the longest rotor scale energized and the remainder of the turbulent spectrum was set to zero energy. In this first run the physical diffusion is turned completely off to demonstrate the non-diffusive nature of the SFD convection algorithms.

The second cross-section plotted in the right half of Fig. 3.1, is the fuel density at each X-Y pixel (i.e. cell) location averaged over all planes k = 1 to 1024. The time of this snapshot is 0.6 ms (timestep 30) after beginning the simulation. The averaged density rapidly deviates from the initial magenta and black values because the X, Y and Z velocities vary in the Z direction starting at time t = 0. Therefore, the Z-averaged density appears to be mixing even though the physical diffusion is set exactly to zero and no actual "molecular" mixing has taken place.

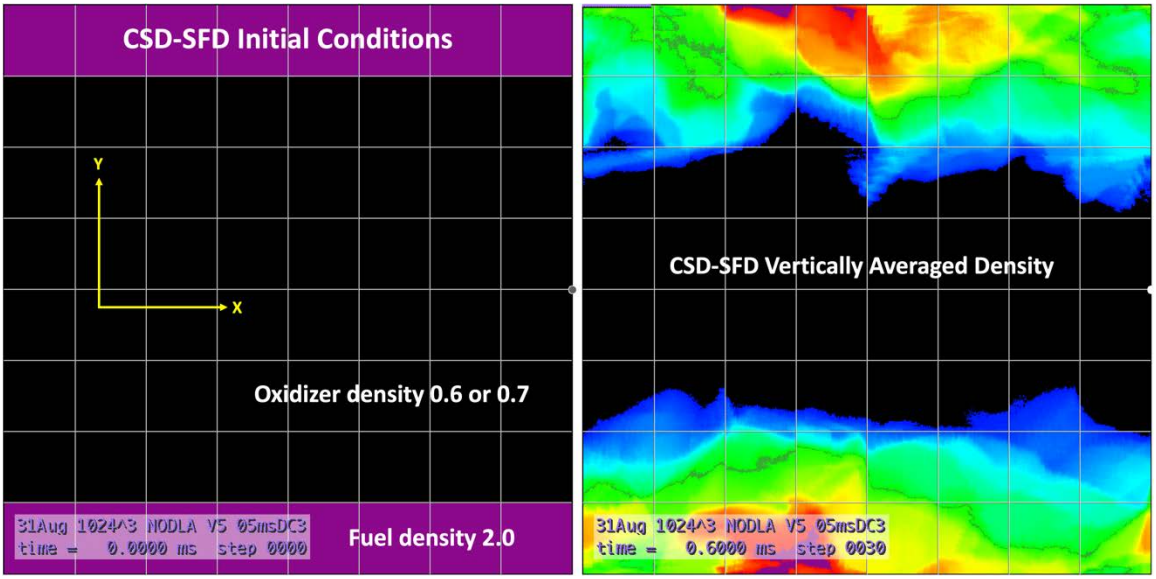

Figure 3.1. Initial conditions for CSD-SFD high-speed flow simulations. Left: Initial fuel density 2.0 (magenta) and oxidizer density 0.6 (black). Right: Fuel density averaged in Z at each X-Y cell.

A two-species SFD system can be represented by a single density array, where the two species densities are either 1.0 or 0.0, when the amount of fuel exactly equals the amount of oxidizer. In this study the fuel density is initially restricted to a quarter of the cube as shown and the oxidizer initially fills the remainder of the volume with density of 0.6 (slightly fuel rich) for most of the runs. The value 0.7 (slightly fuel lean) is used later in the paper. For these cases two density arrays are used. The left-hand image in Fig. 3.1 shows the fuel density on a single Z layer, k = 24. Layer 24, and most of the other layers, even with grid-scale diffusion turned on, show unmixed tracer densities (magenta and black) in some cells as long as 5 milliseconds after mixing begins.

Shearing flows have moved the fuel and oxidizer as much as one system length during the 0.6 milliseconds leading to the right-hand image of Fig. 3.1.

Figure 3.2 below shows the vertically averaged fuel density after 1.2 milliseconds in the left image and after 2.0 milliseconds in the right. A small black patch of unaveraged oxidizer shows in the middle left of the left-hand panel but this patch is gone by 2.0 milliseconds. Once the pure fuel and pure oxidizer are gone, the cross-section plots are rescaled to dark blue through red from the minimum density to the maximum density in the particular cross-section being plotted. Each of these cross sections has a Mixed Volume Index, MVI(t), overlaid on the density color contours. The MVI in the reacting cases of this paper is defined as the average fractional fuel remaining in the entire system. The MVI(t) curve begins at 1.0 in the upper left of each plot and the MVI value is plotted logarithmically against the Y axis while the time increases linearly to the from left to right. At the end of each run the overlaid MVI curve reaches the right edge of the plot. The maximum time is 5 milliseconds in this demonstration case and most of the runs which follow. The MVI is drawn in lavender in runs with 1024^3 cells, blue with 512^3 cells, green with 256^3 cells, and yellow with 128^3 cells. Two of the horizontal white/grey fiducial lines in the cross-sections also indicate a factor of 10 reduction of MVI(t). Thus values of MVI from 1.0 down to 10^{-4} can be plotted. However, in the demonstration run of Figs. 3.1 through 3.3, the value of MVI is pegged at 1.0, the top of the plots, because the added diffusion is rigorously set to zero.

When the diffusion is active at the grid scale, fuel and oxidizer immediately "react" by cancelling each other out stoichiometrically and become unrepresented products wherever they occupy the same cell. Thus the fuel amount is always decreasing. This "reaction" takes the minimum of the fuel and oxidized densities in each cell and subtracts that number from both convected species, thus taking the smallest of the two numbers to zero. The overlap of fuel and oxidizer in each cell is the result of the impressed diffusion alone because the surrogate convection, applied by integral length shifts, guarantees this. Non-integral parts of a required shift along each grid line are added to the shift along that particular line during the next timestep.

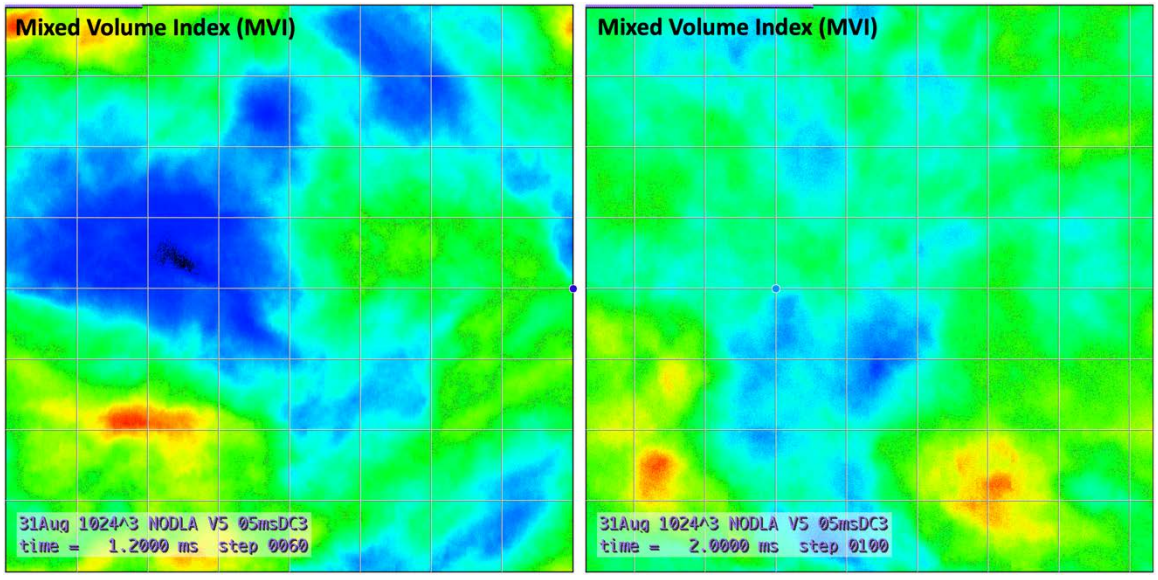

Figure 3.2. Fuel density averaged vertically in Z at each X-Y location at two times during a 5-millisecond high-speed flow simulation. Left: Timestep 60, 1.2 ms. Right: Timestep 100, 2.0 ms.

This shift model for convection provides the enabling basis of the non-diffusive SFD algorithm and has been tested extensively. This feature of the SFD convection is shown by the straight MVI

lines along the top of each image in Figs. 3.1, 3.2, and 3.3 below. Since the mixing is identically zero, the MVI cannot decrease in time as no reaction is taking place. In Fig. 3.4, however, the diffusion is turned back on for a few timesteps, as will be seen, to show the underlying coarse-grain mixing that is actually taking place during the run, even in the absence of diffusion.

Figure 3.3 below shows several more fuel cross-sections at layer k = 24 captured at later times in the diffusion-free simulation begun with the left-hand panel of Fig. 3.1. The longest scale lengths predominate up through timestep 10 (panel a) but shorter scales become evident in panels b) and c). Note that the pixels begin to move quite separately in panels c) and d) below

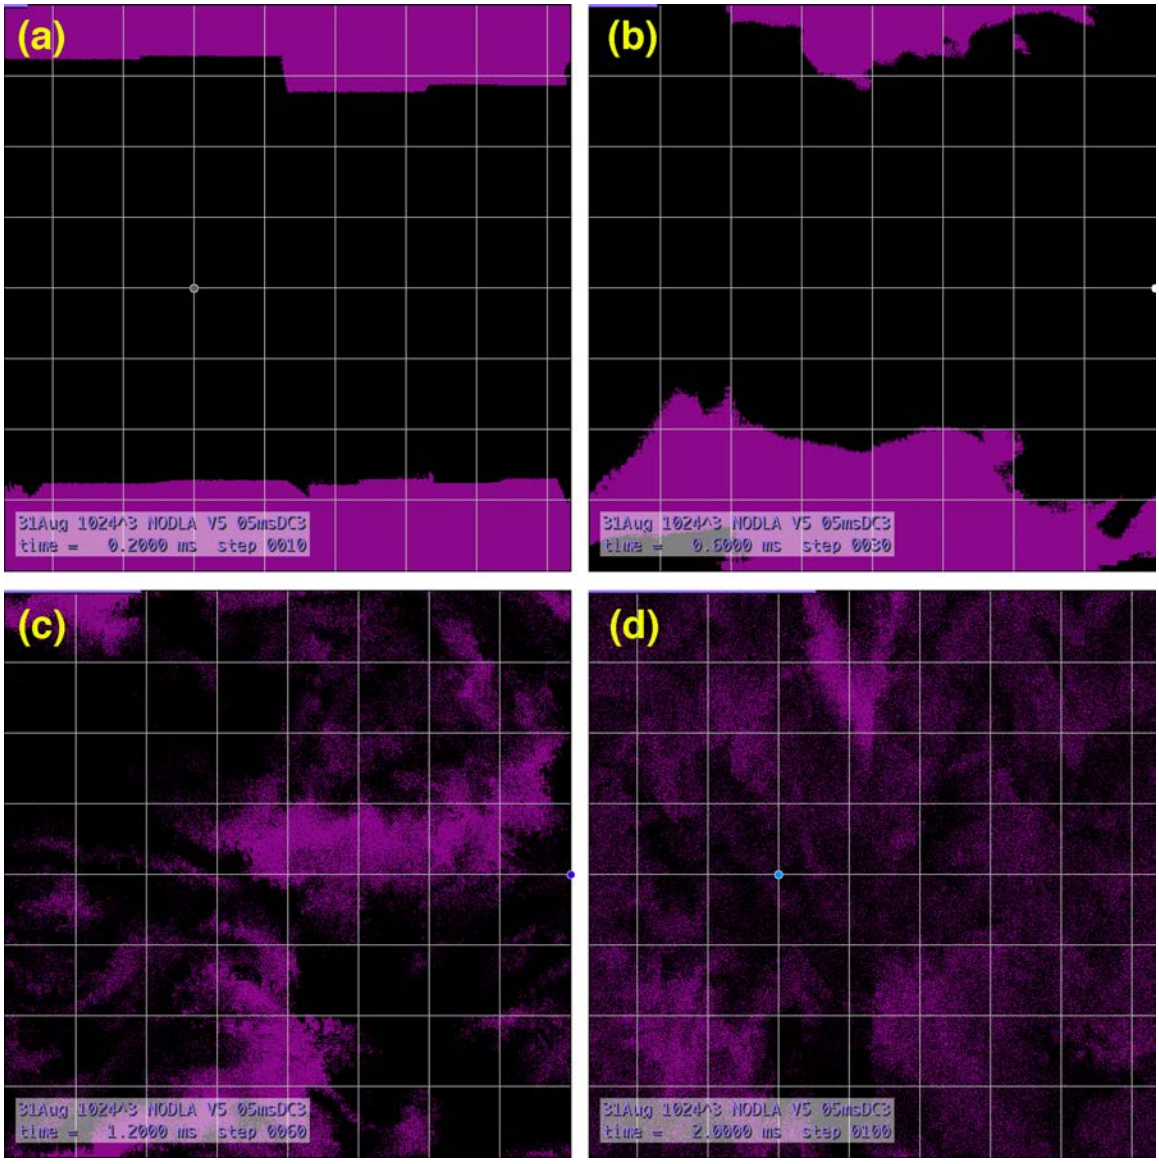

Figure 3.3. Surrogate Fluid Dynamics: high-speed convection without physical or numerical diffusion. Panel a) timestep 10 (0.2 ms). Panel b) timestep 30 (0.6 ms). Panel c) timestep 60 (1.2 ms). Panel d) timestep 100 (2.0ms). Each magenta pixel contains fuel at density 2.0. Each black pixel contains oxidizer of density 0.6.

The colors do not change from magenta and black because no microscopic mixing has actually occurred through 100 timesteps. Figure 3.4 below, taken from parts of images in reference CSD2

from a much slower problem that also has no added diffusion, extends this unique non-diffusive behavior to much later time. The fuel and oxidizer pixels on layer 24 are still clearly separate and visible in the top panel of the figure at step 720. After just two diffusion steps with dt = 1/16 second, the middle panel, has mixed the fuel and oxidizer in adjacent cells and so reveals a much more homogeneous underlying solution.

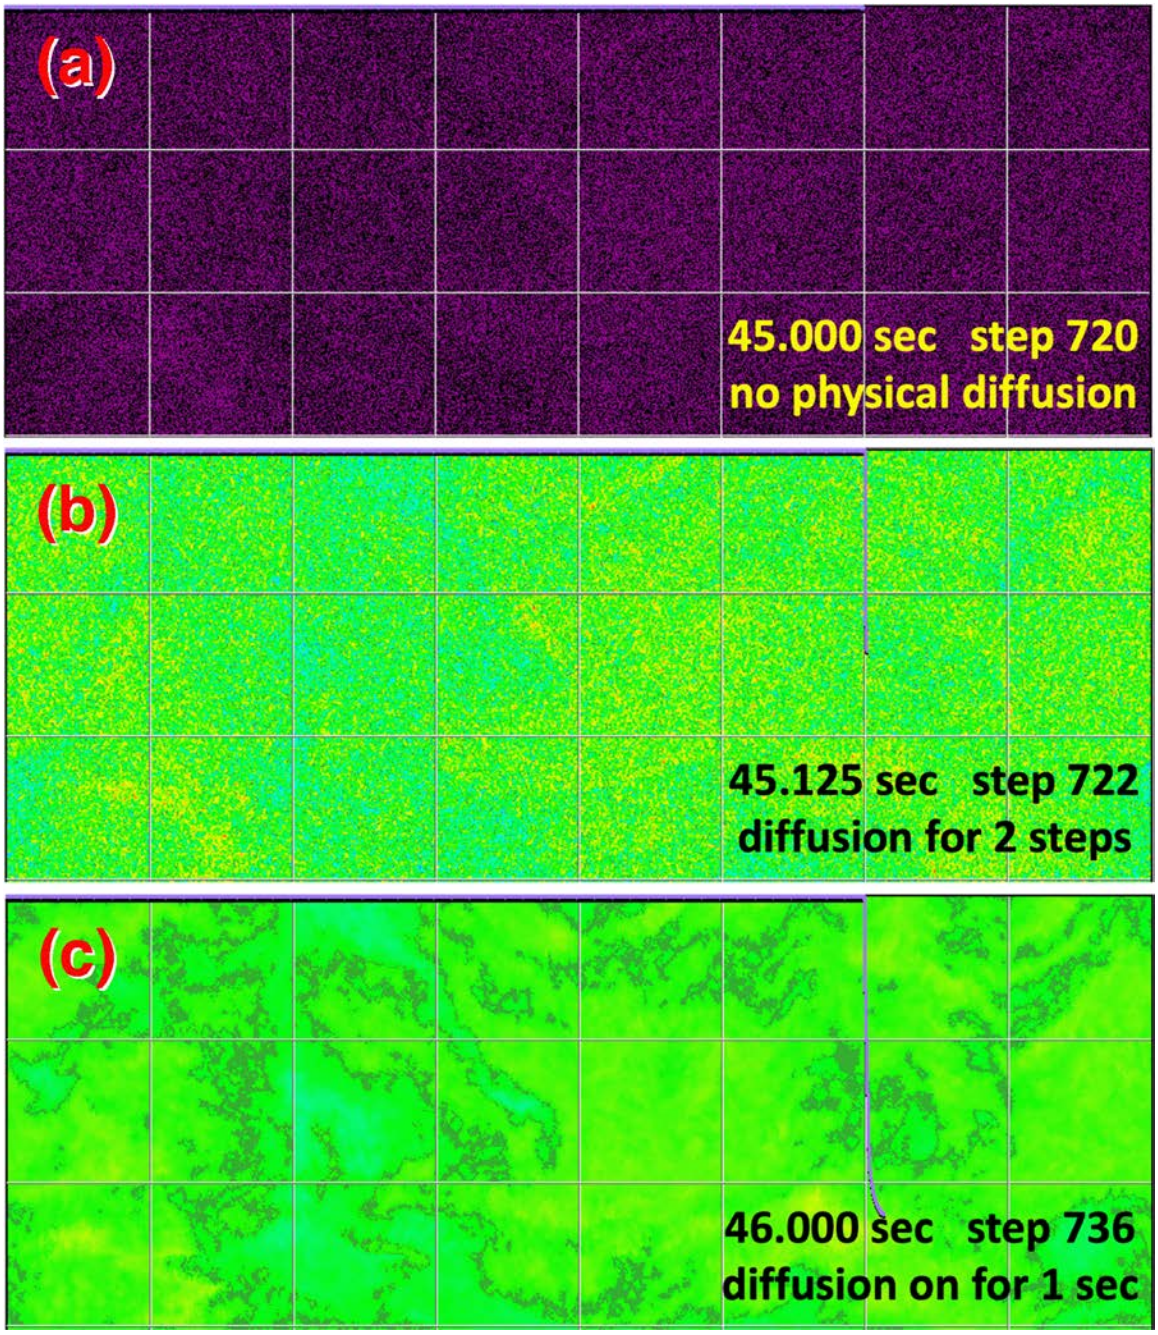

Figure 3.4. Late-time behavior of a 1024^3 non-diffusive CSD-SFD simulation (from reference CSD2). Panel a) 720 timesteps without diffusion. Panel b), grid-scale diffusion turned on for 2 steps. Panel c) diffusion turned on for 16 steps. The solution returns to the state it would have had if diffusion had been in effect for the entire run..

The bottom panel in Fig. 3.4 shows the layer 24 fuel density after 736 steps with 6 steps of grid-scale diffusion. The evidence of individual pixels, still visible in the middle panel, is gone. The

gray dithering in these density plots highlights when the averaged values are close to the average value for this problem, which is 0.5 in this case. Studies in CSD2 show that the strength of the diffusion can be varied over an appreciable range without changing the late-time solution or the time-dependent evolution of the MVI because the flow field is not being affected by the physical diffusion. Unlike a CFD model, SFD allows computing the mixing quite accurately even though the density gradients are under-resolved because the flow field is not being changed.

Mixing occurs in SFD as the overlapping and adjacent fuel and oxidizer densities are diffused each timestep by a simple finite-difference formula at the smallest resolved scale. As mixing progresses from step to step, the MVI values approach the asymptotic, fully mixed value of 0.1 when all the oxidizer has been consumed, assuming the 0.6 fuel rich condition. The detailed solution of 3D continuity equations for this mixing problem is possible in principle using Direct Numerical Simulation (DNS) but requires extremely high resolution [Woodward, et al., 2001; Grinstein, 2001; Coleman & Sandberg, 2010; Poludnenko, 2014]. The problem shown in Figs. 3.1 through 3.4 would likely require 4096^3 or 8192^3 cells to replicate a reasonably delayed onset of "molecular" diffusion to even approximate Figs. 3.3 and 3.4.

This idealized SFD mixing model has much in common with turbulent convective mixing since the material moves the correct distances at the correct velocities and interacts with different scales in about the right manner. Movies of the evolving layer density and vertically averaged density cross-sections appear very realistic despite the obvious inaccuracies of the SFD data-shift representation. The feature of zero numerical mixing is being stressed here because numerical diffusion, causing premature mixing, is a continuing obstacle to the computation of high R_e combustion, particularly in simulations beginning with unmixed reactants.

SFD is not detailed fluid dynamics, clearly deficient is some ways and clearly superior in others. Different realizations of the small scales in SFD cause otherwise identical initial conditions to begin to deviate after one or two large-eddy turnover times, exactly as happens in fluid dynamics. The phasing of even large-scale fluctuations can deviate quickly by 100% between different realizations although macroscopic averages will generally track each other for much longer, even at different spatial resolutions. The results in the following section, requiring these properties of SFD, contain the meat of this paper.

4. Delays in Mixing Due to Non-Equilibrium Turbulence

Figure 4.1 below shows the fuel density plotted on a single Z = constant plane of the fuel-oxidizer cube after 5.0 milliseconds of running time have elapsed for nine different calculations using the CSD-SFD model. These nine simulations all started from an initial condition with only the largest scale rotors, 30 cm, excited. The CSD solution is advanced in lock step with the SFD so almost no mixing occurs until turbulent cascade populates at least some of the small scales. These Surrogate Fluid-Dynamic (SFD) runs are thousands of times faster than a detailed CFD model.

The nine calculations shown in Fig. 4.1 vary the spatial grid resolution (256^3 cells top row, 512^3 cells middle row, and 1024^3 cells bottom row) on a triply period grid. The integration is conducted with three different timesteps (δt = 0.04 ms left column, 0.02 ms center column, and 0.01 ms right column). The graininess of the top row solutions shows the relatively coarse 256^3 grid. The graininess is greatly reduced in the middle row of 512^3 simulations and largely absent in the 1024^3 simulations of the bottom row. In moving from left to right in each row of the figure, the shorter timestep refines the solution but doesn't change it macroscopically. The solutions

after 125 steps (δt = 0.04 ms) , 250 steps (δt = 0.02 ms) and 500 steps (δt = 0.01 ms) are all visually the same (with the exception of the lower left image).

Just as important, the Mixed Volume Index (MVI) of the 9 different solutions can be seen to be essentially the same curve in each case. This MVI curve starts out flat for about 0.5 millisecond, and then begins to mix more and more rapidly. The rate of mixing reaches an inflection after about 2 milliseconds, after which the mixing rate decreases. Since the oxidizer initial density is 0.6 in all these cases, the MVI value drops to about 0.1 after 5 milliseconds, two fiducial lines down from the top on the right edge of each image. These fiducial lines were added to visually subdivide the density cross-sections but they serve a secondary purpose in measuring the MVI value (logarithmic) in the vertical and the MVI time (linear) in the horizontal. The vertical fiducial lines are separated by 0.625 milliseconds for the time scale in these images.

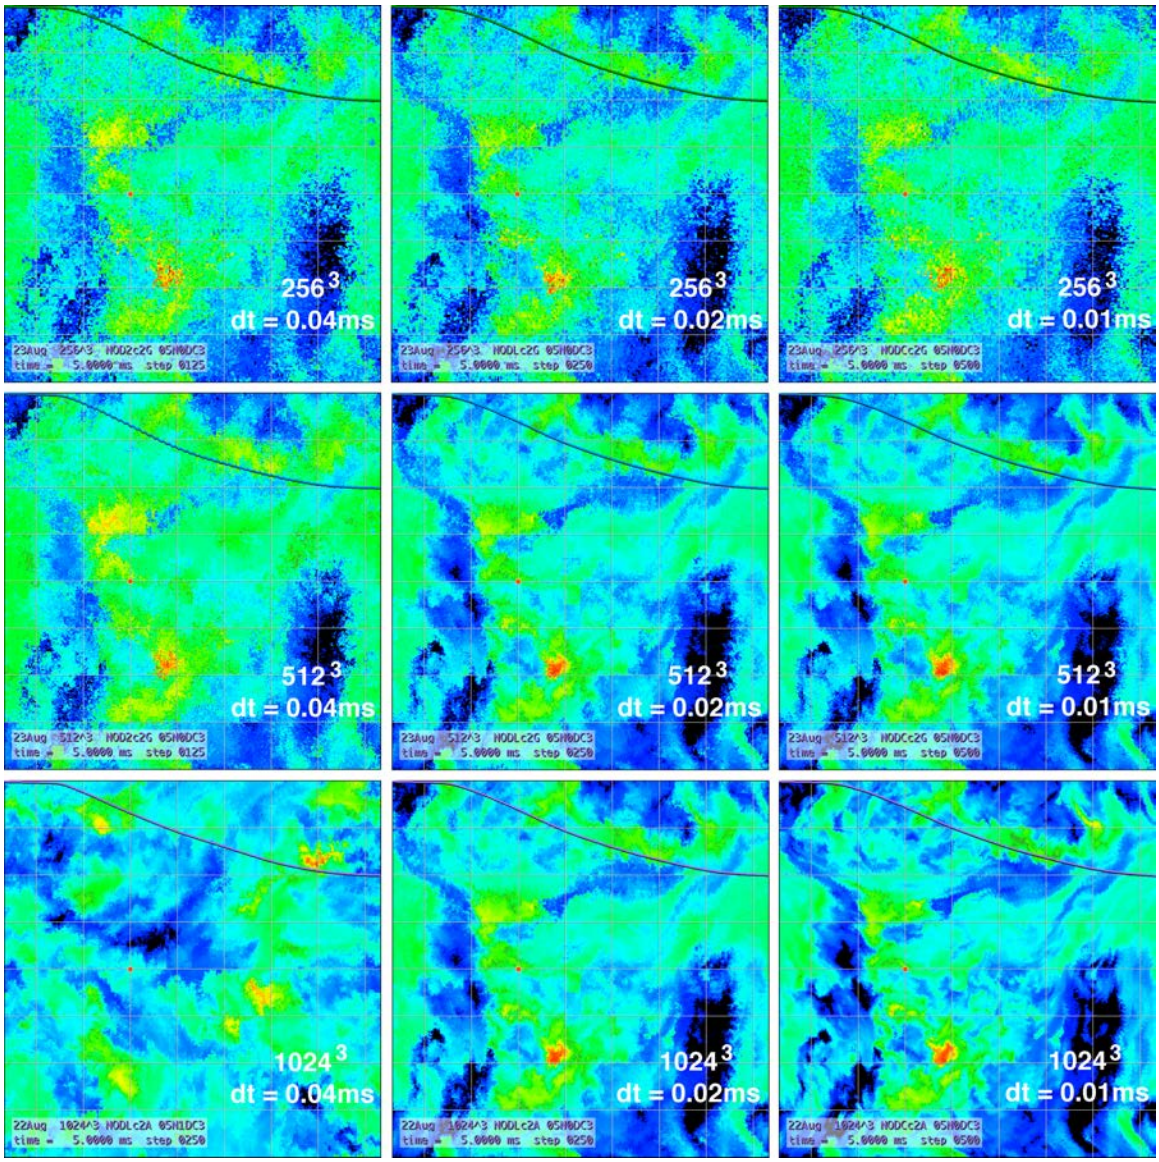

Figure 4.1. Non-equilibrium fuel–oxidizer mixing in one Z = constant layer of a 30 cm cube. Top row: 256^3 (layer 6) grid. Middle row: (b) 512^3 (layer 12) grid. Bottom row: 1024^3 (layer 24) grid. Left column: dt = 0.04 ms. Center column: dt = 0.02 ms. Right column: dt = 0.01 ms. The lower left panel records a different realization of the other eight.

The lower left image in Fig. 4.1 is visibly different from the other eight because it is a snapshot after 5.0 ms from a different realization of the SFD flow from the other eight images. That realization used the same fluid dynamic parameters and resolution and timestep, that the 1024^3 solutions to its right used but all of the random numbers, used to compute the "Monte Carlo" realization of the flow at each timestep are different in this one simulation from the common set of random numbers used for the other eight simulations in the figure. Notice, however, that the resulting MVI curve in this case is basically the same as the other eight because the statistical properties of the turbulence underlying the flow and computed by CSD have not changed.

The turbulent spectrum of the flow in Fig. 4.1 can be the same for each of the nine simulations because the spectrum has been computed at higher resolution than the highest resolution grid and extends to small scales below the Kolmogorov viscous dissipation scale. The particular turbulent flow realization used in each case is simply being evaluated at different resolutions for the different SFD grids. Having a reproduceable spectrum of turbulent flows at each scale ensures that the underlying SFD solution doesn't drift if the resolution and timestep are varied. The strengths and phasing of the rotor components at each common scale are guaranteed to identical to the other simulations but the phasing of the three directional components using in convecting the fuel and oxidizer are different in the sense of a Monte Carlo simulation for the simulation shown on the lower left.

This behavior is markedly different from a full 3D CFD simulation where each refinement of space and/or time resolution allows extra or different structure in the turbulent field being computed at all scales and thus relentlessly changes the solution macroscopically. To say this again, this is not the case in SFD because the turbulent field is fully resolved beyond the resolution of the spatial grid and does not change when the timestep or grid scale of the SFD is changed. Extra structure is available for convection and mixing on a finer grid but the phasing and strength of the larger scales in the flow are unchanged by different resolutions. This means that different solutions can be compared directly, as in Fig. 4.1, without large-scale differences appearing in the longer wavelengths.

The nine simulations leading to Fig. 4.1 are all denoted as "NOD," as seen in the labels on the images in the figure. This stands for "no delay", indicating that the SFD computations begin "without delay" in lock-step with the evolving CSD spectrum as it is being computed. These SFD simulations, as well as others, were also re-run delayed, indicated by "DLA," using the fully equilibrated spectrum shown in Fig. 2 after the spin-up equilibration. When the spectrum of turbulence does not have to be spun up to populate the shorter space scales, the mixing begins promptly so the DLA simulations always reach a given fraction of mixing and reaction earlier than the NOD simulations can reach then same level of mixing. This extra mixing delay in NOD simulations, which more closely approximate the situation existing in supersonic combustors, means that the energy release will actually require somewhat more time and space than predicted by models assuming the faster mixing from existing turbulence models. What we are trying to estimate is the extra residence time needed for fuel and oxidizer to mix in a combustor because the turbulence is initially ill-suited to fast mixing.

Figure 4.2 Overlays one MVI curve computed in an NOD simulation (blue) on top of the MVI curve from an otherwise identical DLA simulation (white) that starts with a fully developed turbulent spectrum. This NOD simulation is from the center image appearing in Fig. 4.1. The background to this Fig. 4.2 is the vertically averaged fuel density computed at 5 milliseconds for the DLA simulation. The NOD-driven mixing starts slowly because the turbulent cascade to the

smaller scales, capable of molecular-level mixing, takes a while to be established. A mixing delay of about 0.5 millisecond is illustrated by the black arrow at the top of the figure. The mixing delay increases quickly to a maximum and then decreases to near zero as the DLA solution reaches in inflection point at about 2.3 milliseconds. At this time about 30% of the fuel has been consumed, i.e. MVI ~ $\sqrt{0.1}$. The mixing delay then increases to a secondary peak between 3 and 4 milliseconds as the limiting MVI value of 0.1 is approached toward 5 milliseconds. A program was written to take the results of pairs of solutions like these two and calculate the mixing delay at each time during the simulations out to the time when the NOD solution for MVI no longer reaches as low a value as the DLA solution.

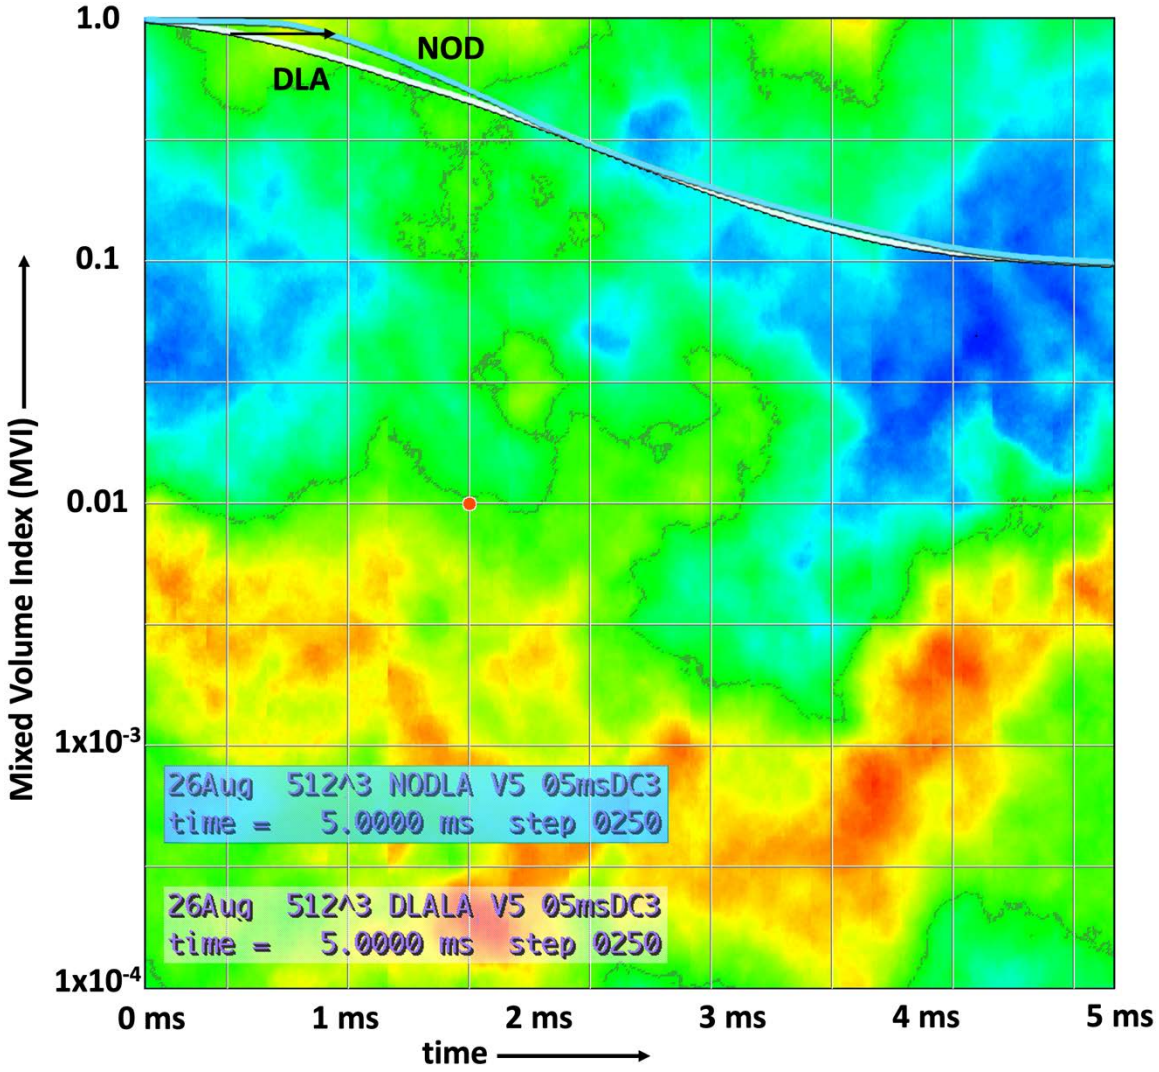

Figure 4.2. Mixed Volume Index solutions for fully equilibrated turbulence driving the mixing (DLA/white) and for turbulence spinning up before mixing (NOD/blue). These two runs used a timestep $\delta t = 0.02$ and ran to 5 milliseconds.

Ten different pairs of solutions were computed using the CSD-SFD model. These pairs of solutions are all similar to the two compared in Fig. 4.2 but used a range of run parameters. For two grid resolutions, 512^3 and 1024^3, three different timesteps, $\delta t = 0.04$ sec, 0.02 sec, $\delta t = 0.01$ sec, and three turbulence stirring speeds, 0.5 km/sec, 0.3 km/sec, and 0.2 km/sec were used in various cases. The 256^3 grid was found to be too coarse and thus was not used. Nine of these pairs were performed with a 10% fuel rich mixture with an initial oxidizer density of 0.6. For the

slower turbulence speeds, the runs had to be extended to 8 milliseconds and 10 milliseconds respectively to approach the fully mixed (reacted) condition where the oxidizer is fully expended. One final run pair was also performed with an initial oxidizer density of 0.7 (about 4% fuel lean). In this run, plotted as red in Fig. 4.3 below, the fuel is all expended and the resulting MVI value reaches .0047, over a factor of twenty more complete reaction than in the other eight simulations. These ten non-equilibrium mixing delay solutions are plotted in Fig. 4.3 below.

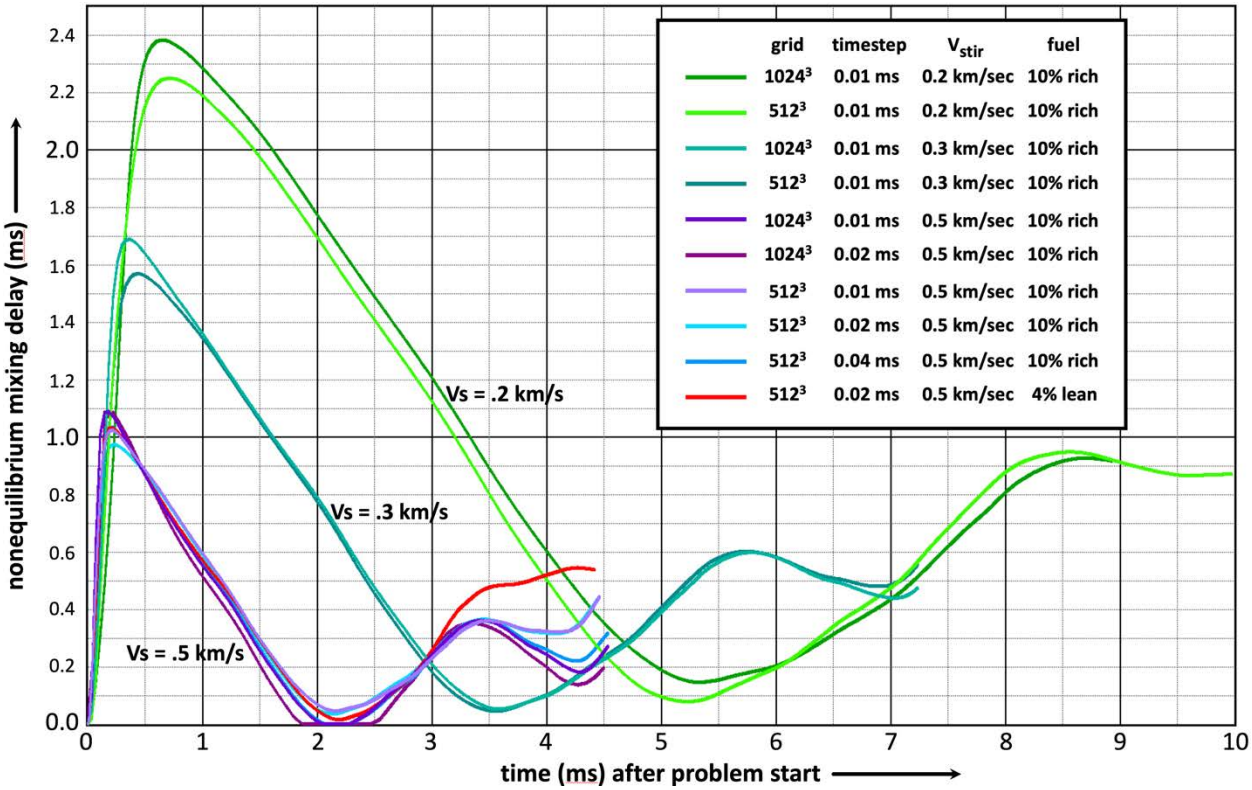

Figure 4.3. Ten solutions of the reaction mixing delay caused by non-equilibrium turbulence in a high-speed flow and plotted as a function of time.

The legend of the figure displays the key parameters of the DLA – NOD comparison solution pairs. The nonequilibrium mixing delay, shown in Fig. 4.3 for each pair, peaks shortly after the beginning of the simulations and then shows a secondary peak much later in time after the NOD solution has passed its inflection point. This rapid first peak occurs because each NOD simulation begins with a laminar flow period and thus shows essentially no fuel-oxidizer mixing until the turbulent spectrum becomes established where the Mixed Volume Index turns sharply downward from unity, which is the fully unmixed condition. MVI for the slower NOD case generally catches up to the corresponding DLA simulation at the inflection point because the large-scale interface between fuel and oxidizer in the NOD run is distorting for a couple of milliseconds before the short scales are fully populated in the spectrum. Therefore, the initial mixing with short wavelengths is active over a larger macroscopic area by about a factor of two. All the cases in Fig. 4.3 seem to be of this type. The second simulation of each pair, denoted DLA, begins with a fully developed turbulent spectrum so mixing can begin right away. It is interesting that the slightly fuel-lean run (red) is quite indistinguishable from the corresponding 0.5 km/sec fuel-rich runs (lavender, magenta, and blue) until well into the second peak at around 3.5 milliseconds.

There are three cases where grid resolution is compared for the same turbulence strength V_s and timestep. In each of these cases, the higher-resolution solution pair shows a slightly larger nonequilibrium mixing delay than the lower-resolution solutions. The 1024^3 grid delay for Vs = 0.2 km/sec (dark green curve) reaches 2.4 ms whereas the 512^3 grid delay reaches only about 2.5 ms. This is also true of the 1024^3 grid blue green curve for V_s = 0.3 km/sec and the 1024^3 grid purple curve for V_s = 0.5 km/sec. The medium blue curve for δt = 0.04 ms and the red curve for the slightly fuel lean solution, both with Vs = 0.5 km/sec, are quite similar with the other Vs = 0.5 km/sec cases for the first three or four ms despite the parameter differences.

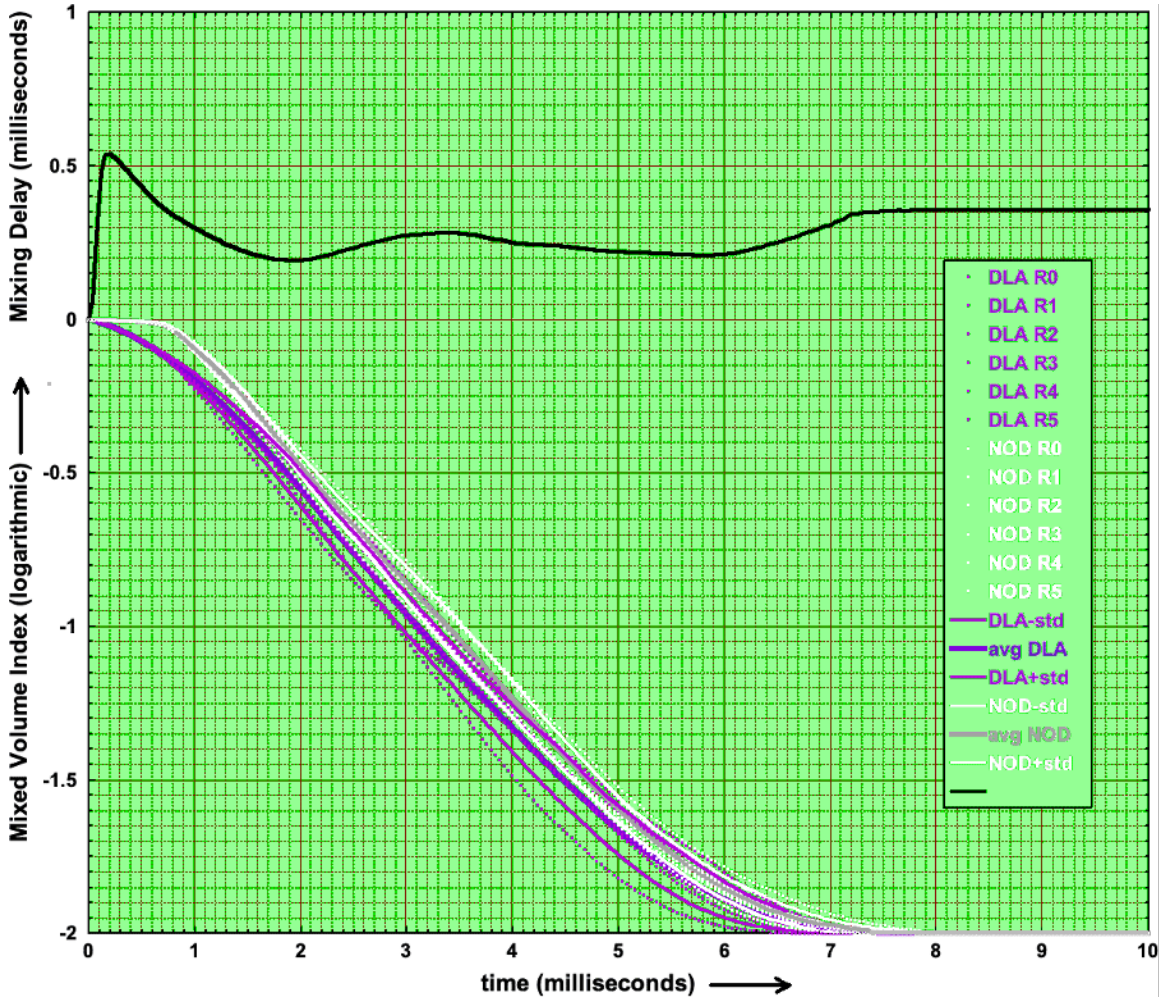

Figure 4.4. Average reaction mixing delays from 6-member ensemble. Lavender curves: average and +/- standard deviations for equilibrated turbulence. White curves: average and +/- standard deviations for non-equilibrium turbulence cascading from laminar initial condition. Black curve: Ensemble averaged mixing delay peaks at 0.55 ms.

Figure 4.4 reports the results of an ensemble study of high-speed mixing delay for a problem with only 1% excess fuel. Six equilibrated turbulence (DLA) realizations are paired with corresponding turbulence onset (NOD) realizations using a laminar stirring speed V_s = 0.5 km/sec. Lavender data points are from the six DLA realizations, R0 through R5, and white data points are from the 6 NOD realizations. The solid purple curve is the average of the 6 DLA Mixed Volume Index curves and is flanked by two thinner lavender curves separated from the average by the DLA standard deviation as a function of time. The solid grey curve is the average of the 6 NOD Mixed Volume Index curves and is flanked by two thinner white curves separated from the

average by the NOD standard deviation as a function of time. The solid black curve at the top is the mixing delay as a function of time, the lag between the average DLA curve and the slower average NOD curve.

Notice that the members of the ensembles are quite close to the average MVI(t) for both DLA and NOD realizations and thus the standard deviations are small. These results are calculated for an initial oxidizer density of 0.66 depleting all but 1% of the fuel before all of the oxidizer is gone. In Fig. 4.3 the initial oxidizer density was 0.60, only sufficient to deplete 90% of the fuel. This difference accounts for the fact that the remaining fuel, which determines the MVI(t) curves, flattens off at 10^{-2} rather than 10^{-1} as in Figs. 4.1 and 4.2. It also accounts for the fact that the mixing delay peaks at 0.55 ms in Fig. 4.4 and at twice that delay in Fig. 4.3 when $V_s = 0.5$ km/sec.

5. Conclusions

This paper used the Coherent Structure Dynamics (CSD) model coupled to the Surrogate Fluid Dynamics (SFD) model to estimate the molecular mixing time delay that occurs when turbulence in a fast flow is quiescent initially and must cascade from laminar macroscopic spatial scales to the small scales where molecular mixing enables chemical reaction, for example, in a supersonic combustor. CSD-SFD is an incompressible model. This approximation makes the computations at least three orders of magnitude faster to perform and is allowed here because the scales that actually perform the small-scale mixing are much slower than the macroscopic driving flow and thus are quite subsonic. Fig. 2, which shows the fully equilibrated turbulence spectrum for the $V_{stir} = 0.5$ km/sec runs, has the largest turbulence velocities. Nevertheless, at the beginning of the viscous dissipation zone the rotor velocities are a factor of ten to twenty slower than V_{stir}.

Although the actual mixing is a relatively slow-speed phenomenon, the non-equilibrium mixing delay can still be a significant fraction of the fuel-oxidizer passage time through small and modest-sized systems. In a frame of reference moving with the flow, assumed to be moving at about Mach 3 here (1 km/sec), the embedded turbulence, once equilibrated will be fast but quite subsonic. The fluid as a whole is moving through the system at about one meter per millisecond but the turbulence will be, at its fastest, a few tenths of that speed.

Figure 4.3 above showed that the peak mixing delay varies approximately inversely with the turbulence strength. The mixing delay is about 1 millisecond when the turbulence driving velocity $V_{stir} = 0.5$ km/sec. It increases to about 1.6 milliseconds when $V_{stir} = 0.3$ km/sec, and goes as high as 2.4 milliseconds when the turbulence driving velocity is 0.2 km/sec. This means, roughly, that the energy release from combustion will begin a meter farther downstream than might be expected from equilibrium turbulence theories when the peak turbulence speeds are 0.5 km/sec. With driving velocities of 0.3 km/sec and 0.2 km/sec, the extra distance downstream is about 1.6 meters and 2.4 meters respectively.

In a small or moderate sized combustor these mixing delay distance will change the behavior of the system significantly to drastically. If equilibrated turbulence is present before the fuel is introduced, these mixing delay distances reduce to about 0.25 meters, 0.4 meters, and 0.6 meters respectively for the three turbulence strengths. There could, however, be an important energy penalty for providing this turbulence before mixing begins. These results suggest that the introduction of strong mixing flows should be orchestrated so interpenetration of fuel and

oxidizer on scales short enough to facilitate rapid cascade of turbulence to microscopic scales without shocking up the bulk flow unacceptably.

In larger systems with combustors several meters long, the mixing delay due to non-equilibrium should be less of a problem. However, there are other situations in high-speed flight where the delay of turbulence onset could play important roles. Examples may include heat transfer and vision from optical windows. Research and development that I would like to conduct next includes refining these estimates to apply more closely to systems actually being considered, looking into these other applications of delayed turbulence onset estimates that can be made, and the addition of algorithms to the model to incorporate a treatment of energy release and approximate dilatation within the framework already established for SFD.

The heat release from combustion, once it has begun, adds another mechanism that has yet to be added to SFD: the baroclinic generation of additional turbulence at scales that enhance mixing rates once the turbulence has established. These effects, both the local baroclinic turbulence generation and the nonlocal effects that modify the entire spectrum, are currently being worked on. They may not significantly speed the onset of combustion but can likely enhance its strength once they are initiated. Surrogate models for these effects are being coded and tested in the CSD-SFD model and are the subject of a future report. The goal, as it relates to practical CFD applications, is information about the duration of the non-equilibrium delays to be expected and data for a simple reduced order approximation for incorporating these delays in more detailed (and expensive) engineering models.

Acknowledgements

The research leading to Coherent Structure Dynamics was supported by the project on “Highly Complex Fluid Dynamics” within the ONR/NRL 6.1 research program. The NRL/ONR 6.1 research program on “Advanced Computational Models that Exploit Many-Core Computer Architectures” supported the optimization of the CSD-SFD models through OpenMP and the program on "Numerical Simulation of Hypersonic Scramjet Combustion with Non-Uniform Inlet Conditions" supported the preparation of this report. My NRL collaborators, Gopal Patnaik, Gabe Goodwin, and Jesse Maxwell contributed to the progress of this work. Professor Alexei Poludnenko (University of Connecticut), Dr. Fernando Grinstein (Los Alamos National Laboratory), and Professor Paul Woodward (University of Minnesota) have contributed to the underpinnings of the Coherent Structure Dynamics Research in the last few years.

References

Bendeguz D. Bak and Tamas Kalmar-Nagy, 2018, “A Linear Model Of Turbulence: Reproducing the Kolmogorov-Spectrum,” IFAC PapersOnLine 51-2 (2018) 595-600.

Luca Biferale, 2003, “Shell Models Of Energy Cascade In Turbulence,” *Annual Reviews of Fluid Mechanics*, 35:441–68, 2003: doi: 10.1146/annurev.fluid.35.101101.161122

Jay P. Boris, 2020, “A Coherent Structure Dynamics Model for Non-Equilibrium Turbulence,” Naval Research Laboratory Memorandum Report, in process, 2020, ref. CSD1.

Jay P. Boris, 2020, “A Mixing Study Using Coherent Structure Dynamics to Drive a Surrogate Fluid Dynamics Model,” Naval Research Laboratory Memorandum Report, in process, 2020, ref. CSD2.

Garry L. Brown and Anatol Roshko, 1974, “On density effects and large structure in turbulent

mixing layers," Journal of Fluid Mechanics **64**, part 4, pp. 775-816, 1974.

O. Poludnenko and C. Colby, 2021, private communication.

M. Coantic , J.-J. Lasserre, 1999, "On pre-dissipative 'bumps' and a Reynolds-number-dependent spectral parameterization of turbulence," European Journal of Mechanics B/Fluids 18 (1999) 1027–1047

Gary N. Coleman and Richard D. Sandberg, 2010, "A Primer on Direct Numerical Simulation of Turbulence – Methods, Procedures and Guidelines," University of Southampton, Technical Report AFM-09/01a (March 2010).

Fernando F. Grinstein, 2001, "Vortex Dynamics and Entrainment in Rectangular Free Jets," *Journal of Fluid Mechanics*, **437** (2001): 69–101.

David R. Mott and Elaine 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.

Alexei Poludnenko, 2014, "Pulsating instability and self-acceleration of fast turbulent flames," Physics of Fluids **27**, 014106 (2015); https://doi.org/10.1063/1.4905298

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).