

# Local Baroclinic Turbulence Models for Mixing

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*

June 29, 2022

# 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.**

<b>1. REPORT DATE (DD-MM-YYYY)</b> 29-06-2022			<b>2. REPORT TYPE</b> NRL Memorandum Report		<b>3. DATES COVERED (From - To)</b> 11-01-2021 – 04-18-2022	
<b>4. TITLE AND SUBTITLE</b>  Local Baroclinic Turbulence Models for Mixing					<b>5a. CONTRACT NUMBER</b>	
					<b>5b. GRANT NUMBER</b>	
					<b>5c. PROGRAM ELEMENT NUMBER</b>	
<b>6. AUTHOR(S)</b>  Jay Boris and Keith Obenschain					<b>5d. PROJECT NUMBER</b>	
					<b>5e. TASK NUMBER</b>	
					<b>5f. WORK UNIT NUMBER</b> 1V28	
<b>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)</b>  Naval Research Laboratory 4555 Overlook Avenue, SW Washington, DC 20375-5320					<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b>  NRL/6003/MR--2022/2	
<b>9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b>  Naval Research Laboratory 4555 Overlook Avenue, SW Washington, DC 20375-5320					<b>10. SPONSOR / MONITOR'S ACRONYM(S)</b>	
<b>11. SPONSOR / MONITOR'S REPORT NUMBER(S)</b>						
<b>12. DISTRIBUTION / AVAILABILITY STATEMENT</b>  <b>DISTRIBUTION STATEMENT A:</b> Approved for public release; distribution is unlimited.						
<b>13. SUPPLEMENTARY NOTES</b>  This is the fourth NRL Memorandum Report in a series. There will be at least five. This is an extension of the CSD#3 memo report and describes progress on new approximations estimating BaroClinic Sources (BCS) (density gradients generating vorticity) specifically for turbulent combustion.						
<b>14. ABSTRACT</b>  This progress report presents tests and simulations of algorithms approximating the effects of fuel-oxidizer reactions on the turbulent flow computed by the 3D Coherent Structure Dynamics – Surrogate Fluid Dynamics (CSD-SFD) model. Expansion from gas-phase combustion cause density gradients and these generate additional turbulence as transient baroclinic pressures act on these density gradients. This ‘baroclinic’ turbulence increases the rate of mixing and can accelerate the combustion. Direct Numerical Simulations (DNS) can resolve some details of these effects but the effective Reynolds numbers are very low and the computational costs are very high for practical situations. Here we use Coherent Structure Dynamics to provide a nonequilibrium turbulent spectrum over seven orders of magnitude in scale and use Surrogate Fluid Dynamics to simulate 3D mixing from this spectrum. CSD-SFD runs on a laptop in an hour or two for 1024-cubed grids. The price for this 100 to 1000-fold speed increase over CFD is the use of a fundamentally incompressible convection model. The compressibility effects of baroclinic turbulence must be grafted onto the SFD model. This progress report presents tests and results of two computationally efficient approaches to approximate the added mixing that results from baroclinic turbulence.						
<b>15. SUBJECT TERMS</b>  Computational fluid dynamics      Coherent structure dynamics      Non-equilibrium Turbulence Surrogate fluid dynamics      High-speed mixing for combustion						
<b>16. SECURITY CLASSIFICATION OF:</b>				<b>17. LIMITATION OF ABSTRACT</b>  U	<b>18. NUMBER OF PAGES</b>  35	<b>19a. NAME OF RESPONSIBLE PERSON</b> Jay Boris
<b>a. REPORT</b> U	<b>b. ABSTRACT</b> U	<b>c. THIS PAGE</b> U	<b>19b. TELEPHONE NUMBER (include area code)</b> (202) 767-3055			

This page intentionally left blank.

# 1. Introduction and Background

This paper reports progress in creating algorithms for baroclinic turbulence in Surrogate Fluid Dynamics (SFD) allowing high Reynold's-number computations of non-equilibrium turbulent combustion. "The baroclinic effect is due to the nonalignment of pressure and density gradients, and its result is to induce vorticity production. Because of the steep density gradients and the almost universal presence of ambient pressure signals in combustion systems, the baroclinic effect is a crucial mechanism for the production of turbulence." (Batley et al., 1996). Baroclinic turbulence is extensively studied in atmospheric and oceanographic flows (e.g. Thompson & Young, 2007; Smith, 2007; Farrell and Ioannou, 2009; Gallet & Ferrari, 2020) but is also important in combustion (Batley et al., 1994; Towery et al., 2016; Bray et al., 1985; Bray, 2016; Poludnenko and Oran, 2010; Poludnenko, 2015). However, in enhanced mixing for combustion, for which CSD and SFD were developed, the subject is seldom discussed in terms of baroclinic turbulence, perhaps because the problem with chemical reactions is generally more complicated. Baroclinic sources (BCS) for combustion are discussed further in Section 2, "Compressible Surrogate Fluid Dynamics." Here we first describe the new models developed recently at NRL to solve for turbulence and mixing in high Reynold's number flows where the turbulence is expected to be far from equilibrium, as in fast combustion.

Surrogate Fluid Dynamics (SFD) is a new, high-resolution numerical model that produces evolving realizations of a 3D flow satisfying a specified, time-dependent turbulence spectrum. The turbulence is not required to be in equilibrium or even near equilibrium but the time-varying energy in each resolved scale is embedded in the flow field produced. In this research, the evolving turbulence spectrum is provided by concurrently running the NRL 3D Coherent Structure Dynamics (CSD) model. CSD is described in the NRL Memorandum Report (ref. CSD#1) "A Coherent Structure Dynamics Model for Non-Equilibrium Turbulence" (Boris, 2018). This report includes the equations being solved, a discussion of the algorithms, and the results of a number of tests. CSD calculates the evolution of the turbulence for minutes of real time in a few seconds of laptop computing over a wider range of space scales, seven orders of magnitude, than can be treated by Computational Fluid Dynamics simulations running for days on large, highly parallel mainframes. The composite CSD-SFD model predicts the turbulent pre-dissipation bump (e.g., Biferale, 2003; Bak and Kalmar-Nagy, 2018), sometimes called the bottleneck, as shown in Figs. 1 and 2 of CSD#1, and determines the manner and time scale of the eventual nonlinear turbulent relaxation to an equilibrium spectrum.

In the mixing simulations presented in CSD#2 (Boris, 2021), "A Mixing Study Using Coherent Structure Dynamics to Drive a Surrogate Fluid Dynamics Model," Detailed Numerical Simulation (DNS) is not practical. Instead, the CSD model is used in CSD#2 to drive a Surrogate Fluid Dynamics (SFD) flow field, approximating high Reynold's-number turbulence that mixes the initially separated species. 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 captures fluid velocities at all scales driven by the time-dependent turbulent spectrum 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 numerical diffusion or diffusive regridding. Adjusting the X, Y, and Z shifts at all resolved scales simultaneously is done

in parallel for efficiency and to impress the evolving turbulence velocity spectrum in a correlated, continuously varying manner.

This SFD model 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 flow velocities in each scale bin of the CSD spectrum at the time the shift is made. In 3D the relative shear of shifts on adjacent lines generates density gradients and can be viewed as vorticity. The term ‘fine-graining’ is sometimes 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 in SFD, the fractional shifts are saved and accumulated until a later timestep can implement an integral-cell shift. Of course, there are accuracy and convergence considerations but there is no Courant stability condition and no diffusion. Therefore, shifts of several cells can be applied in each direction without stability problems. These features lead to a very efficient algorithm.

In the current version of CSD-SFD each species density is defined on a regular, triply periodic cube. This limitation is a simplification and is not required. Four grids are used to vary resolution,  $128^3$ ,  $256^3$ ,  $512^3$ , and  $1024^3$ . Two SFD species densities on the finest grid can be advanced using the CSD-based turbulent flow for hundreds of timesteps in an hour or so on a 4-core MacBook Pro. However, this gridded density field represents only the longest wavelengths of the 30 factor-of-two scale-size bins used to represent the turbulence flow. 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 macroscopic CFD scale to study the mixing in both equilibrium and non-equilibrium turbulence scenarios, as demonstrated by convergence studies in references CSD#2, CSD#3 and here.

The third NRL Memorandum Report in this series, CSD#3, “Mixing Delays in Non-Equilibrium High-Speed Turbulence,” (Boris and Obenschain, 2022) further discusses and presents results using CSD-SFD to calculate the extra time-delay before mixing in a high speed flow when the turbulence is out of equilibrium initially and must cascade to small scales before mixing can occur. In high-speed combustion applications, this extra time-delay for turbulence to build up may require adding a meter or so to the combustor length.

This report presents tests and results of algorithms approximating the effects of fuel-oxidizer reactions on the turbulent mixing in the 3D Coherent Structure Dynamics – Surrogate Fluid Dynamics (CSD-SFD) model. Heat release in gas-phase combustion causes density gradients due to expansion which generate additional turbulence from the action of transient ‘baroclinic’ pressure gradients on these density gradients. This ‘baroclinic’ source accelerates the turbulence and thus increases the rate of mixing. Direct Numerical Simulations (DNS) are capable of resolving these effects in detail but the effective Reynolds numbers are unrealistically low and the computational costs are unrealistically high for practical situations. Coherent Structure Dynamics provides the non-equilibrium turbulence spectrum spanning six orders of magnitude in scale and Surrogate Fluid Dynamics simulates 3D mixing using this time-dependent spectrum on a laptop in an hour or two on  $1024^3$  grids. The price for this 100 to 1000-fold speed increase is using a fundamentally incompressible convection model onto which the effects of baroclinic turbulence must be grafted. This report presents some of the

tests and results of two computationally efficient approaches to approximate the added mixing of baroclinic turbulence.

The next section discusses the physics of expansion-generated turbulence that underpins the models we will be testing. First, however, we present results of a couple of simulations using CSD-SFD without the combustion-related turbulence that enhances mixing. Figure 1 below shows an X-Y cross section of the fuel density on a single Z layer of the grid after 10 ms (500 timesteps). This CSD-SFD simulation used the  $1024 \times 1024 \times 1024$  grid and resolves a cube which is 30 cm on a side. The 32<sup>nd</sup> Z layer is plotted and can be compared, without interpolation, with Z layers 16, 8, and 4 on the  $512^3$ ,  $256^3$ , and  $128^3$  grids respectively. This particular  $1024^3$  run is based on the baseline scenario used to measure the mixing between two initially separate species in reports CSD#2 and CSD#3. In this scenario, no baroclinic generation of turbulence is present, either locally or non-locally, so the mixing is caused entirely by the cascade of incompressible turbulence driven by the system-scale stirring initiated at time  $t = 0$ .




Figure 1.1 Fuel density at 10 ms (500 timesteps) on layer 32 of a  $1024^3$  CSD-SFD simulation. The Mixed Volume Index,  $MVI(t)$ , in lavender, overlays the density field.  $MVI(t)$  decreases from 1.0 at  $t = 0$  ms (entirely separate fuel and oxidizer) to 0.01 on the edge of the plot after 10 ms (500 steps).

The metric for mixing is determined by a fast reaction that cancels equal amounts of fuel and oxidizer in all cells containing both species. Removing the minimum of the two species densities from both species requires almost no computation to model fast reaction. The algorithm implemented also allows the reaction to proceed more slowly to demonstrate convergence of the results vis-à-vis temporal and spatial resolution. The lavender curve beginning in the upper left of Fig. 1.1 plots the Mixed Volume Index,  $MVI(t)$ , as a function of time overlaid on the density plot. The MVI is the fraction of the total amount of fuel initially in the system that remains at time  $t$ , plotted along a horizontal axis scaled from 0.0 ms to 10.0 ms (left to right edge of the cross-section shown). The scale for  $MVI(t)$  is at the left of the figure.




Figure 1.2 Fuel density at 8 ms (400 steps) on layer 32 of the  $1024^3$  grid CSD-SFD simulation overlaid with vertical strips from  $512^3$  and  $256^3$  simulations of the identical scenario. The Mixed Volume Index,  $MVI(t)$ , is plotted over the density strips in a different color for each resolution. The white extension of the  $MVI(t)$  curve on the right extends the result to 10 ms (500 steps).

These plots use a shifting color scale for the fuel density on the cross section ranging from the maximum density (red) to the minimum density (dark blue). Black is used when the density value is 0.0 and magenta is used near the beginning of the run when the density has the initial value 1.0, as shown in Fig. 3.1 below. In the baseline run 19JanW illustrated here, the system

contained 1% more fuel than oxidizer so the Mixed Volume Index asymptotes to 0.01 at the right edge of the figure as mixing approaches completion. The small grey 'x's in the image are placed where the scaled density, minimum to maximum, is close to 0.5. The 1024 x 1024 pixels in Fig. 1.1 visualize convection of the density field at all scales down to the grid scale, here 0.03 cm. When the resolution is changed, the resolved scales of an SFD flow are the same from one resolution to another, making tests for convergence a simple graphical comparison. Movies made of an SFD flow appear realistic with motions at many scales but these flows are not solutions of the Navier-Stokes equations, but rather are based on the major features of Navier-Stokes solutions such as mass, momentum, and energy conservation, causality, continuity, proper scaling of inviscid coherent structure interactions, and consistent boundary conditions.

Figure 1.2 shows CSD-SFD simulations of the scenario in Fig. 1 at an earlier time, 8.0 ms, before the run has completed. Vertical strips from three grid resolutions, labeled along the top of the figure, are laid side by side to highlight the differences caused solely by resolution. The Mixed Volume Index is plotted on each of the slices, in lavender at 1024<sup>3</sup> resolution, in blue at 512<sup>3</sup> resolution, and in green at 256<sup>3</sup> resolution. There is essentially no displacement of the MVI(t) curve between the 1024<sup>3</sup> and 512<sup>3</sup> results, indicating a high degree of convergence for this macroscopic indicator of mixing. The composite MVI(t) curve is extended in white to 10 ms, the right edge of the plot, using the 1024<sup>3</sup> result for the last 2 ms. Comparing the colored fuel-layer 32 image from the 1024<sup>3</sup> simulation to the corresponding layer 16 pixels of the 512<sup>3</sup> simulation across the vertical lines separating the segments also shows very little difference between the different resolution computations even though much greater variability might be expected. The 256<sup>3</sup> results, though still quite similar to the abutting 512<sup>3</sup> and 1024<sup>3</sup> image strips, show somewhat greater deviations. Further, in the 256<sup>3</sup> results the pixilation itself is evident. The computations for Fig. 1 and Fig. 2 used the baseline configuration from reports CSD#2 and CSD#3. New algorithms, providing both local and nonlocal baroclinic additions to the turbulence spectrum that enhance mixing, are presented and tested in Section 3, 4, and 5.

Section 3 of CSD#3, figures 3.1, 3.2, and 3.3, show the initial conditions and a few frames of the graphical diagnostics that automatically accompany the simulations. The baseline problem is a cube of fluid three quarters of which is oxidizer and one quarter a layer of fuel containing 1% more material than the broader oxidizer layer. The laboratory frame is moving rapidly to the left at 1 km/sec, about Mach 3. The stirring velocity at the 30 cm scale is 0.5 km/sec but cascade to shorter wavelengths drops this to subsonic levels at the 7.5 cm scale.

## 2. Compressible Surrogate Fluid Dynamics

Computational and experimental research by Picone, Boris, Greig, Raleigh and Fernsler (Picone et al., 1981) and by Greig, Pechacek, Raleigh, and Fernsler (Greig, et al., 1983) studied the rapid cooling of hot channels produced by lasers and electric discharges, such as lightning, in ambient air. Following expansion to pressure equilibrium, these hot channels were seen to cool several orders of magnitude faster than predicted by classical thermal conduction. Since any reasonable estimate of background gas velocities falls far short of explaining the cooling rate, we require a mechanism to convert curl-free fluid expansion, which does not mix, into persistent vorticity, which can. The importance of determining the mechanisms and scaling laws governing the dynamics and cooling of such hot channels derives from the application of discharge physics to other areas, such as beam physics and the study of Nitrogen fixation in the

atmosphere by lightning. In the latter case, the cooling rate of the hot channel gas is critical to predicting the global production of nitrogen oxides. (Boris and Picone, 1982).

The present application is to extend the Surrogate Fluid Dynamics model to include the main qualitative effects that expansion, due to heat release, has on turbulent mixing. Combustion physics leads to spatially varying changes in density and pressure, driven by expansion from local heating, and culminates in the baroclinic generation of vorticity at different scales. In this work, the factor-of-two logarithmic scale-size grid of the CSD model describes the turbulence spectrum over orders-of-magnitude in scale length for regions where density is being decreased by expansion. In these regions we need to enhance the mixing rate caused by prior mixing that has allowed fresh fuel and oxidizer to react.

In the simplest approach, the convective motion of the gas that accompanies combustion-driven expansion need not be implemented at all. The distances moved in the expansion are usually small. The flow accelerates away from the center of the combusting region due to the heat-release over pressure and then decelerates as the fluid again approaches a nearly constant, equilibrated pressure situation. The displacement of the gas due to the heat release depends on the amount of energy released but not on the actual rate of energy release. The transport from the combustion that enhances mixing occurs relatively slowly as a result of the persisting vorticity and turbulence generated baroclinically.

Vorticity, energizing the baroclinic turbulence, is generated quickly by the expansion at a rather low level but persists far longer than the actual energy release and so can transport and mix distinct species significantly. To implement the turbulence enhancement in SFD, we must treat both the local effects, where enhanced mixing occurs near the baroclinically-generated turbulent vortices, and the non-local effects on the overall spectrum being predicted by CSD in lock step with the surrogate flow field. The non-local evolution of the turbulence spectrum is easily treated by adding the energy release computed from the reactions into the spectrum in the appropriate density-gradient scale bins, which may differ in size from place to place. The approach used here involves a theory developed in the early 1980s for beam-channel cooling based on experimental work in NRL's Plasma Physics Division. This research is documented in several reports and journal articles from the time. This problem was treated computationally and experimentally for the rapid cooling of laser, charged-particle beam, and lightning channels in air (J.M. Picone, et al., 1981, Greig, et al., 1983; Boris and Picone, 1982). This effort was also extended to baroclinic turbulence driven by shock waves and flames (Picone et al., 1984).

As a region of energy release expands to pressure equilibrium, any deviations from spherical symmetry will lead to misalignment between the pressure and density gradients. The equation which describes the resulting baroclinic vorticity evolution is

$$\frac{d\xi}{dt} + \xi \nabla \cdot v = \xi \cdot \nabla v + \frac{(\nabla P \times \nabla \rho)}{\rho^2} \quad \text{Eq. 2.1}$$

where vector the vorticity  $\xi(x, y, z, t) \equiv \nabla \times v(x, y, z, t)$  Eq. 2.2

and  $v$  is the fluid velocity,  $\rho$  the density and  $P$  the pressure. We can approximate the resulting flow field by vortex filament pairs of strength  $\pm\kappa$ . Here  $\kappa$ , a circulation with units of area times vorticity, satisfies the approximate equation

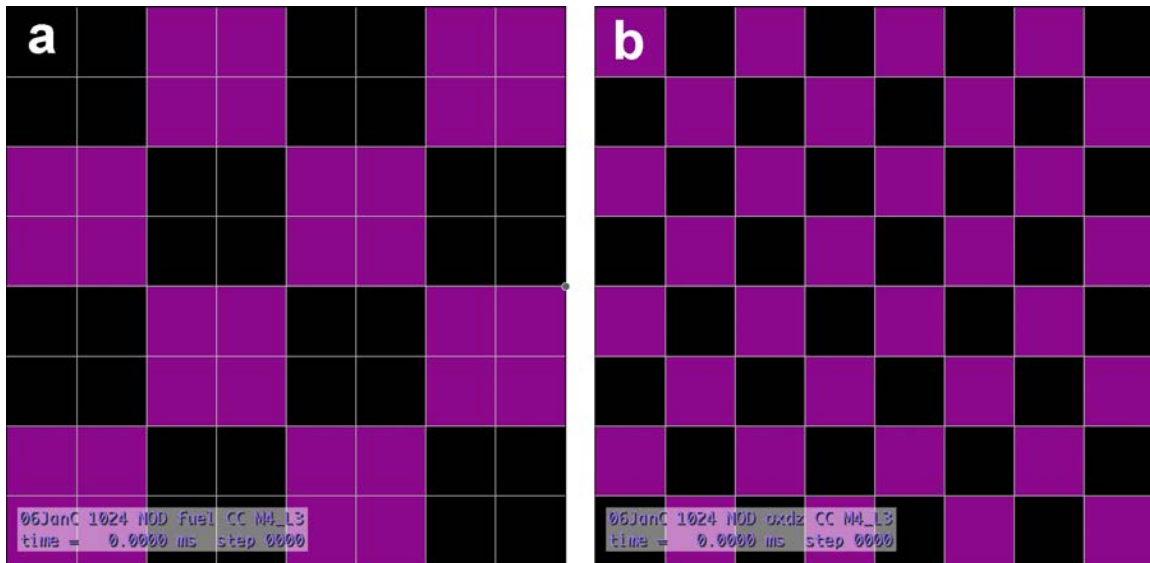
$$\kappa(U_m, \tau, \rho_\infty, \rho_0) \approx U_m^2 \frac{\tau}{2} \ln\left(\frac{\rho_\infty}{\rho_0}\right) f_{geom}(U_m, \tau, \rho_\infty, \rho_0). \quad \text{Eq. 2.3}$$

$U_m$  is the expansion velocity of the expanding-region boundary when the energy is being released,  $\rho_\infty$  and  $\rho_0$  are the ambient density and the density at the center of the reacted region just after energy release, and  $f_{geom}$  is a form factor which may be approximated by formulae derived in Section IV of Picone and Boris (1983). The quantity  $\tau$  is the time interval over which vorticity generation occurs, approximately the time during which chemical energy is being released and the hot gas expands to pressure equilibrium.

The local effect of enhanced mixing in the vicinity of the density gradients can be approximately treated by noting that  $\kappa(U_m, \tau, \rho_\infty, \rho_0)$  in Eq. 2.3 above already has the units of a diffusion coefficient. Added diffusion is typically applied to implement subgrid models of almost all types. In this case of added baroclinic ‘turbulence,’ the free numerical parameters are the strength of the diffusion and the area over which it must be applied. By enhancing the local grid-scale mixing of the model, the extra mixing that occurs can be realized in roughly the correct places at the correct times.

### 3. Tests of Two Local Baroclinic Turbulence Generation Models

Two different preliminary BaroClinic Source (BCS) algorithms are tested to approximate the local generation of baroclinic turbulence caused by expansion from chemical reactions. The first local BCS model (algorithm 1) averages the fuel and oxidizer densities in square patches near locations where significant reaction is taking place. This algorithm is similar to other subgrid models that use variable grid-scale diffusion to represent unresolved turbulent convection. The second local BCS model (algorithm 2) uses non-diffusive data exchanges swapping regions of fuel or oxidizer within a square patch to transport these species locally. This also approximates baroclinically-generated vortex pairs that move species to enhance mixing. This second BCS algorithm has advantages because both local and global conservation are maintained exactly, making the SFD implementation more efficient computationally.



**Figure 3.1** Initial conditions for testing two local baroclinic mixing models. The fuel ( panel a) and oxidizer panel (b) are non-interacting checkerboard patterns of different scale. Combustion-drive acceleration of fluid across density gradients in a flow generates ‘baroclinic’ vorticity that contributes to turbulent mixing in the vicinity of these gradients.

Figure 3.1 shows the separate 2D initial test conditions for the fuel and oxidizer species. They are treated separately in the initial tests though they are solved together for efficiency in running the full CSD-SFD model. In all cross-section figures such as 3.1, magenta regions have density 1.0 and black regions have density 0.0. The two initial density profiles in the figure are used to determine how quickly the two BCS algorithms disperse the reacting species. One goal of SFD is to treat sharp density gradients with zero numerical diffusion. Thus, the fuel and oxidizer have discontinuous patterns of two different scales initially, a) and b), to test the mixing of the two algorithms as the patch size and grid resolution are varied.

The initial implementations and tests of these two algorithms are in 2D on each X-Y plane within the 3D triply-periodic CSD-SFD grid. Both local BCS algorithms use disjoint square patches of cells with a specified power-of-two size, and these patches are laid out edge to edge to tile each X-Y plane. In both algorithms, the patches are all independent, allowing different treatments from patch to patch. Therefore the answers will be the same in each of the magenta and black regions of the checkerboard, for example. This symmetry is broken in a

turbulent flow by large scale convection that moves the species across the boundaries of the individual checkerboards and thus between the individual square patches on which the BCS algorithms are applied. In these preliminary tests run outside of the full CSD-SFD model, a constant overall shift is applied to each plane, one cell up (Y) and one cell to the right (X) every timestep on the  $1024^3$  grid. On the  $512^3$  grid, this is one cell up and to the right every other timestep. On the  $256^3$  grid, this is one cell up and to the right every fourth timestep to allow physically consistent tests. In addition a small random component is added, visible in the upper left panel of Fig. 3.2 below at iteration 001.

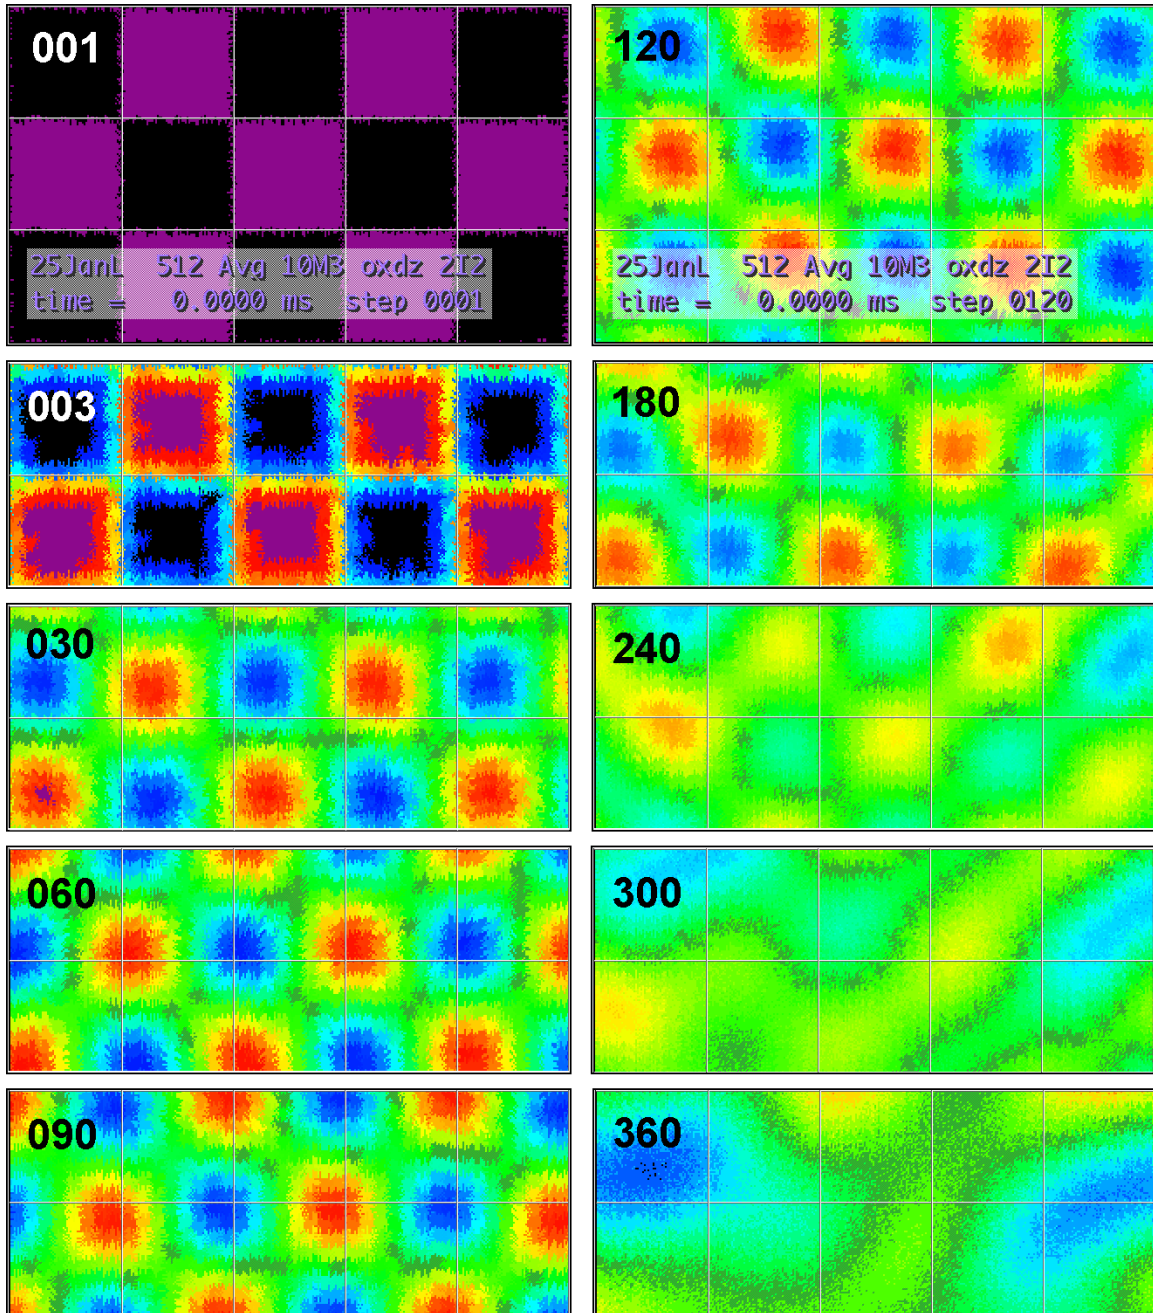


Figure 3.2 Local averaging (algorithm 1), is applied repeatedly to  $8 \times 8$ -cell patches throughout the entire 3D grid, i.e.  $M = 3$ . Oxidizer is the  $8 \times 8$  checkerboard (b) in Fig. 3.1. The panels in Fig. 3.2 display only a few of the checkerboard blocks from the lower left of Fig. 3.1 panel b. The lower 8 panels show only 10 of the 15 blocks to remove the legend.

The tests reported here, leading to Figs. 3.2 and 3.3, used the averaging local BCS algorithm. Figs. 3.5 through 3.8 were prepared using the local Data Exchange (swap) algorithm. None of these are full simulations within the CSD-SFD model but rather limited tests of the diffusion-like transport of the two underlying local BCS mixing algorithms. The additional models to determine where, when and how much BCS transport should be added to the SFD turbulent mixing and how it should be calibrated are discussed in section 6, "Discussion and Future Work."

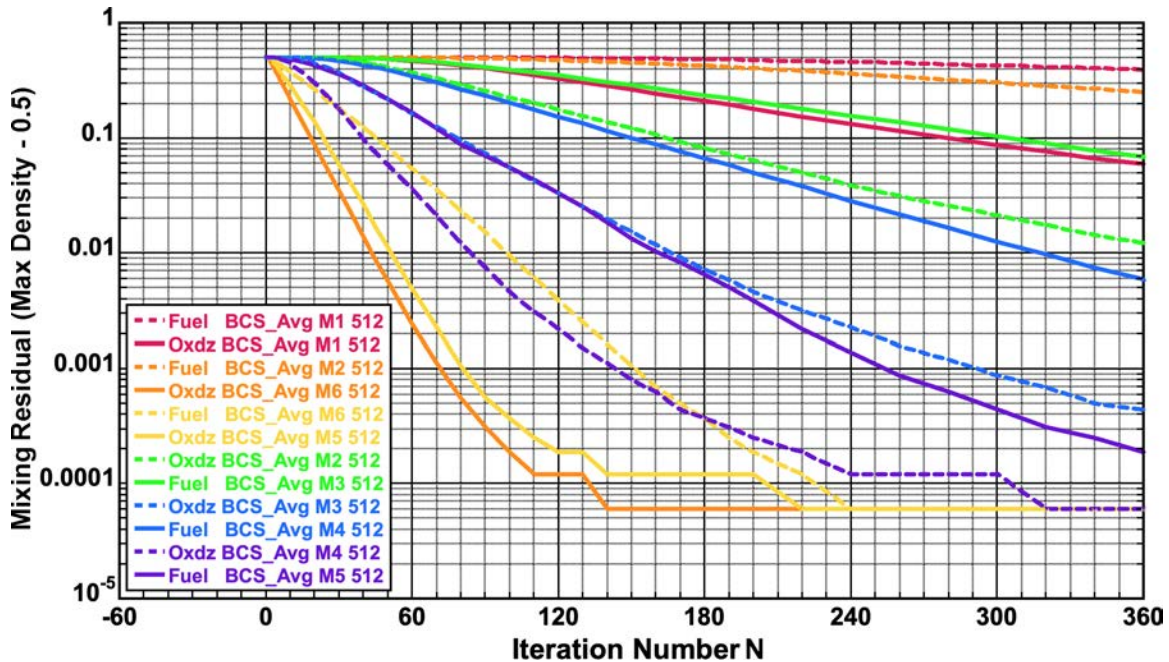
The average of the oxidizer (or fuel) values in the  $2^M$  by  $2^M$  cells of each square patch becomes the new, baroclinically mixed value across each local patch after an iteration of added BCS transport. The full 2D grid for the averaging test, labeled 25JanL in Fig. 3.2 above, is  $512^2$  cells. Simulations at  $512^3$  resolution on the 3D grid seem to be well converged in full CSD-SFD runs, as shown in Figs. 1.1 and 1.2, above and they complete in a few minutes of CPU time. The top two panels in Fig 3.2 show the averaging-algorithm oxidizer density in the lower left 15 blocks of the 2D test plane, after iterations 1 and 120. The color map in these figures ranges for red (1.0-) to dark blue (0.0+). Again, magenta and black are reserved for 1.0 and 0.0 respectively. As timestep (or iteration numbers) increase, the range between maximum (red) and minimum (blue) becomes smaller and smaller and can vary with grid resolution slightly, even in the absence of chemical reactions to deplete the species.

For each of the two algorithms being tested, the results are shown graphically by a sequence of 2D density color-contour plots of the single X-Y plane being tested. The average of the initial density for each algorithm is 0.5 because the checkerboard blocks alternate in value. The spreading of the oxidizer over the test plane shown in Fig. 3.2 for algorithm 1 varies between 0.0 to 1.0 for a few iterations (i.e. timesteps) but the oxidizer 'diffuses' sufficiently by about iteration 15 that the maximum drops below 1.0 and the minimum rises above 0.0. In Fig. 3.3 below, the quantitative mixing metric for the overall effect of the averaging BCS algorithm 1 is the density maximum at iteration number N minus the overall average value 0.5. This metric is a number that decreases from 0.5 initially and approaches zero as the density, either fuel or oxidizer, becomes more and more evenly distributed across the test plane.

The remaining eight panels of Fig. 3.2 show the progression of the averaging algorithm 1 in 10 of the 15 checkerboard blocks shown at the top, for a sequence of iterations up to 360. At iteration 1 (upper left) the edges of the magenta and black blocks of cells already show a few pixels of roughness as the conservation algorithm forces a few pixels to change to ensure rigorous conservation. These fluctuations and the progressive shifting of the solution to replicate the effects of large-scale convection on symmetry breaking in what otherwise would be a perpetually periodic solution since the densities are represented by 16 bit integers. The legend is in the lower left of the top two panels of Fig. 3.2 The notation '2I2' in the legend is a reminder that the fuel and oxidizer are represented in reduced precision as 16-bit integers (2 bytes). The notation '10M3' indicates that the patch is being fully diffused with no fractional inclusion of the timestep (or iteration) beginning conditions and that the patch size is  $8 \times 8$  (M3 =  $2^3 \times 2^3$  cells). Thus there are  $64 \times 64$  separate patches filling the  $512 \times 512$  cells in each layer of the 3D cube. '10M3' yields maximum local diffusion.

The initial conditions for fuel and oxidizer are not mirror images of each other for these isolated tests (see Fig. 3.1). Thus two separate cases of the mixing from the local BCS algorithms are evaluated without the effects of the full multi-scale SFD turbulent flow confusing

the test results. In these tests, however, as in the full CSD-SFD realized flow, the fuel and oxidizer are convected simultaneously as a single 4-byte word. The patch size for the diffusion (local ‘averaging’ algorithm 1) as well as for data exchange (local ‘swap’ algorithm 2) can be set as 2x2, 4x4, 8x8, 16x16, 32x32, or 64x64 cells. Varying the patch size is one way to vary the strength of the BCS from location to location on the species mixing. Certainly the very large patch sizes, particularly on the lower resolution spatial grids, will give an unacceptable result, just as too large a filter size can degrade some conventional CFD subgrid turbulence models. The smaller block size of the oxidizer initialization allows the action of algorithm 1 to be quantified in fewer iterations than for the fuel. Symmetry breaking can also be seen in Fig. 3.2.



**Figure 3.3** Mixing residual tests of local BCS averaging on a 512<sup>2</sup> test grid. Both fuel and oxidizer tests are shown for six square-patch sizes. BCS algorithm 1 spreads fuel and oxidizer gradients into the fluid where chemical reaction is occurring. Added numerical diffusion is part of this algorithm, above what SFD provides every other step at the finest scale as molecular mixing.

Figure 3.3 plots the residual mixing metric for algorithm 1, i.e. Density Max(N) – 0.5, for both fuel and oxidizer. All patch-size M values 1 through 6 are computed for 360 iterations, that is 2 x 2, 4 x 4, . . ., through 64 x 64 cells (M6). Uniformly, the oxidizer residual decreases faster than the fuel residual at any fixed patch size M because the oxidizer has more initial structure. Further, increasing the patch size M, and thus the averaging area, uniformly speeds the rate at which the unmixed residual decreases. Oxidizer patch sizes M=4, M=5, and M=6 all accelerate mixing enough to display roundoff limitations of the 16-bit species representations beyond 100 iterations.

For reasons of execution speed, programming issues, and adjustability in application, the use of the data exchange (DE) algorithm 2 is recommended. Thus we next switch to tests of the algorithm 2. Six possible data exchange (swap) patterns, #1 through #6, are implemented on square patches with edge lengths of 2<sup>M</sup> cells (M = 1, 2, ..., 6). These are executed in OpenMP parallel loops by calling a subroutine on disjoint square patches to ensure there are no parallel thread data conflicts. Figure 3.3 explains these 6 swap patterns.

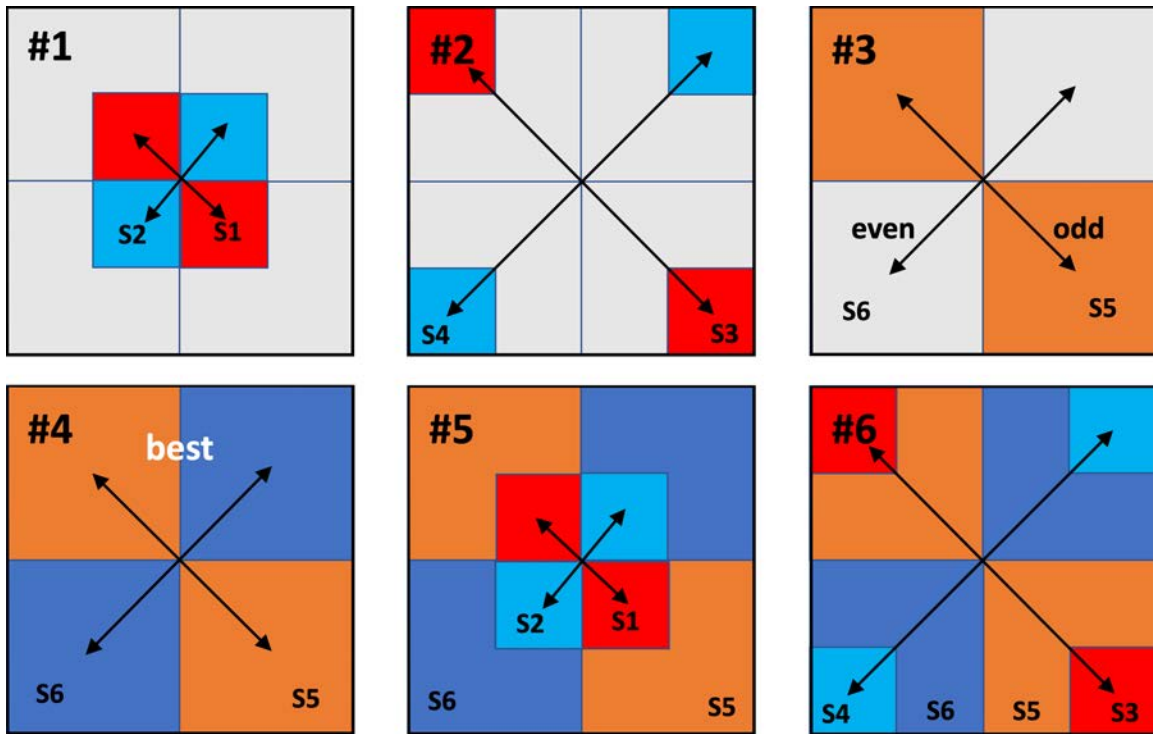


Figure 3.4 Six data exchange ‘swap’ patterns are implemented on square cell patches to represent convective transport of fuel-oxidizer gradients by surrogate baroclinic turbulence in a multispecies fluid where chemical reaction may be occurring. The swap patterns #1 to #6 suggest giving progressively more transport of the chemical species.

The BCS algorithm 2 is implemented by moving (swapping) the density values in square sets of cells within each  $2^M \times 2^M$  patch along the patch diagonals, as illustrated in Fig. 3.4. Each patch is divided into quarters, and when desired, into sixteenths. In swap pattern #1 (upper right) the red squares (S1) are exchanged and the blue squares (S2) are exchanged. Only 1/4 of the patch values have to be moved a short distance, that is, 1/4 the patch width. The resulting transport is correspondingly small. In swap pattern 2# (top center), only 1/4 of the patch values are again moved but the distance is three times greater, 3/4 of the patch width. Thus the transport is correspondingly larger than for swap #1 although the computational work is nominally the same. Swap pattern #3 involves the exchange of twice as many density values as swap patterns #1 and #2 and swap #3 comes in two versions, marked ‘even’ and ‘odd’ in diagram #3 above. These two versions are alternated on even and odd iterations (timesteps) to reduce the asymmetry that use of only swap #3 (even) or swap #3 (odd) produces.

Swap pattern #4 requires twice as much computation as swap #3, but also transports twice as much density. Swap patterns #5 and #6 would seem to produce more transport than swap #4 but this proves not to be the case because the additional 1/16-th exchanges in swaps #5 and #6 move some of the exchanged cell values back to where they came from. Although this extra work will increase some small-scale gradients, the amount of transport of fuel into oxidizer seems to be kept the same or even reduced. In the implementation of algorithm 2 in the full SFD flow model, only swap patterns #2, #3, and #4 are used as they provide a distinct progression from less to more transport.




Figure 3.5 BCS transport of fuel using 3 swap patterns (DE algorithm 2) and 32 x 32 cell patches. Species gradients migrate into the fluid when chemical reaction is accelerated due to surrogate baroclinic turbulence. Additional numerical diffusion is not added at this stage. The right half of swap Pattern 3 has been symmetrized.

Figure 3.5 above is a composite of progressive iterations in separate tests of swap patterns #2, #3, and #4 acting repeatedly on the large fuel-block initial condition of Fig. 3.1. These tests were performed using the  $1024^2$  planar test grid with  $32 \times 32$  BCS patches to provide maximal resolution with relatively large BCS transport while reducing the data exchange patch size to about 0.1% of the planar area. No diffusion at any scale is introduced by this BCS algorithm so the grid-scale 'molecular-diffusion surrogate' is provided by the same alternate cycle adjacent cell averaging used in the full 3D CSD-SFD model. The slight randomization of the monotonic drift algorithm as a test approximation to the impact of large-scale flow on the BCS algorithm is also included as in the averaging algorithm 1 discussed above.

The separate figure panels in Fig. 3.5 are all plotted with red as the maximum and dark blue as the minimum value to emphasize the density differences across the plane as the individual density values all approach the system average 0.5 asymptotically. The most obvious feature of Fig. 3.5 is the stark asymmetry shown in the left half of the central column for swap #3. These six blocks were computed exclusively using swap #3 (odd) from Fig. 3.4 above and this leads directly to the obvious asymmetry shown. The 6 blocks on the right half of the column were computed using the even-odd alternation of the swap #3 data exchanges shown in Fig. 3.4.

With this correction, the figure shows similar qualitative behavior for all three swap patterns with increasing iteration number, due to the continual color scaling, while later figures, 3.7 and 3.8, show up to an order of magnitude difference in the computed residuals in the separate tests. These three data exchange (swap) patterns provide a monotonic increase in transport whose effects can be hard to see by comparing the columns side by side in this figure. Iteration 10 still shows black and magenta with all three swap patterns, indicating that the local BCS transport has not yet eroded the fuel large-block initial conditions. The figure also shows the beginning of symmetry breaking as the small randomization effects of the system wide shifts, which provide a simple approximation to larger scale flow effects absent from the tests. This symmetry breaking, is most apparent in swap #4 where diagonal merging is beginning between some of the diagonally-adjacent red-yellow and blue-light blue patterns.

The symmetry breaking, which is just becoming apparent in the large-block fuel test of Fig. 3.5, is much faster and more obvious in the smaller-block oxidizer tests shown in Fig. 3.6 below. Diagonal asymmetries are beginning to creep in by iteration 200. By iteration 300 the initial eight-by-six block pattern has completely dissolved for swap patterns #3 and #4 into larger-scale blobs ranging across several of the original magenta-black initial condition blocks that are still organized as predominantly red and predominantly blue patterns symmetrically confined to the original array of square blocks. By iteration 400 even the swap pattern #2 results in the left column are showing major asymmetries. By iteration 500, the remaining symmetry in all three swap patterns tested has broken as smaller scales in the oxidizer density profiles are dissolving, leaving only longer and system wide fluctuations visible. In swap patterns #3 and #4 black and magenta have begun to reappear as the maximum and minimum values begin to merge within a couple of bits, meaning that the maximum scaled color again becomes 1 and the minimum again becomes 0.


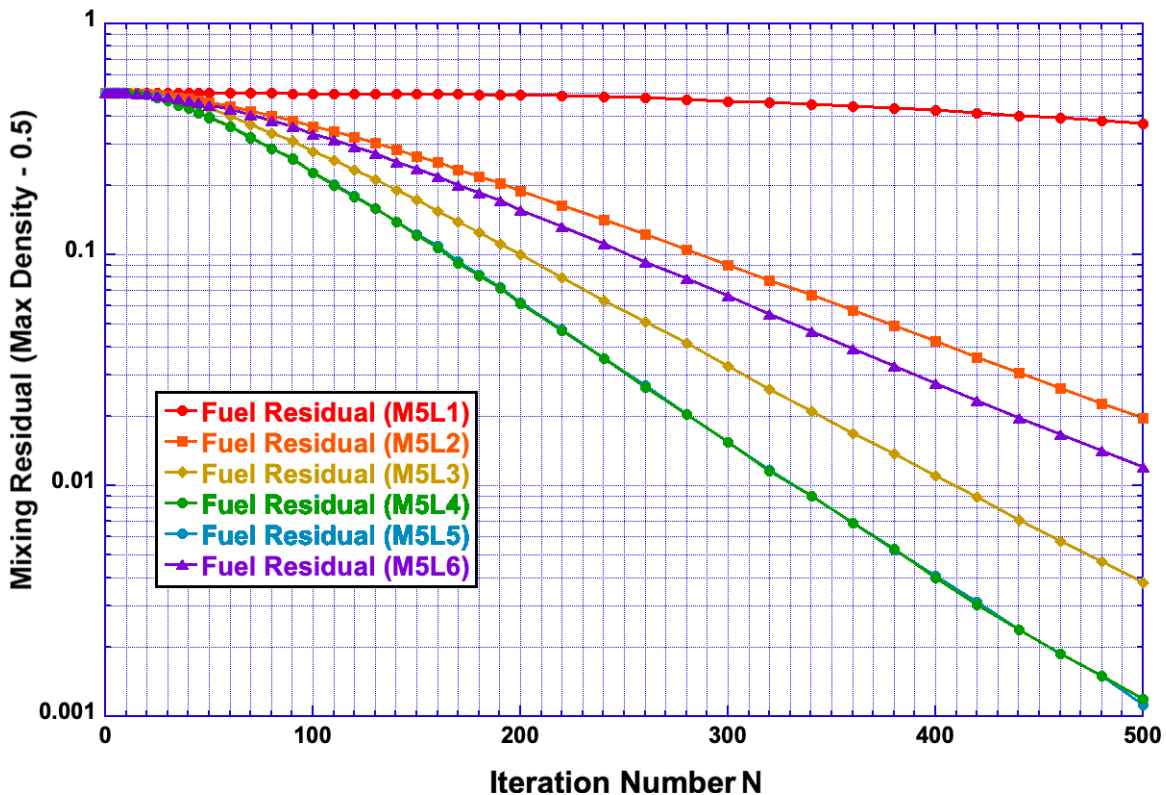


Figure 3.6 BCS transport of fuel using 3 swap patterns (DE algorithm 2) and 32 x 32 cell patches. Species gradients migrate into the fluid when chemical reaction is accelerated due to surrogate baroclinic turbulence. Additional numerical diffusion is not added at this stage. Swap Pattern 3 has been symmetrized.

The Mixing Residual metric, as computed for Fig. 3.3, was collected for both the fuel and the oxidizer tests as a function of iteration N and plotted on logarithmic scales in Figs. 3.7 and 3.8 below. The fuel underwent three decades of residual decay, as seen in Fig. 3.7, and the oxidizer, with its smaller initial block size, underwent 5 decades of decay. In both these figures, the mixing residual was calculated for all six of the diagrammed BCS swap patterns. Figure 3.7 shows that the residual for each swap pattern begins with a short region at the upper left where the residual is nearly constant at 0.5 and thus the maximum is approximately 1.0. This flat region is where the patches still have significant regions of magenta or black. Then each curve transitions to roughly an exponential decay.



**Figure 3.7** Mixing residual tests of algorithm 2, the local BCS data exchange ‘swap’ model. Fuel tests on the  $1024^2$  grid with  $32 \times 32$  patches (M5) using six swap patterns, L1 through L6. Algorithm 2 increases fuel and oxidizer gradients in the fluid where chemical reaction is occurring. Added numerical diffusion is NOT part of this algorithm.

Swap pattern #1 provides even less transport (residual decay) than #2 through #6 and this is borne out in Fig. 3.7. Swap patterns #2, #3, and #4 provide progressively more transport, as borne out in the figure. This progression provides a way to scale the extent of the BCS transport to the intensity of the local chemical reactions but additional research is needed to increase the number of data exchange patterns available in this progression and to calibrate the conditions under which each is used. It may seem surprising that swaps #4 and #5 shown nearly identical mixing residual curves although #5 shows a little more decay from iteration 480 to 500. In any case, this suggests not using swap pattern #5, which requires more computational work than swap #4. Swap pattern #6 gives a residual decay suggesting more transport than swap #2 but less than swap #3 and this is also not used in the following.

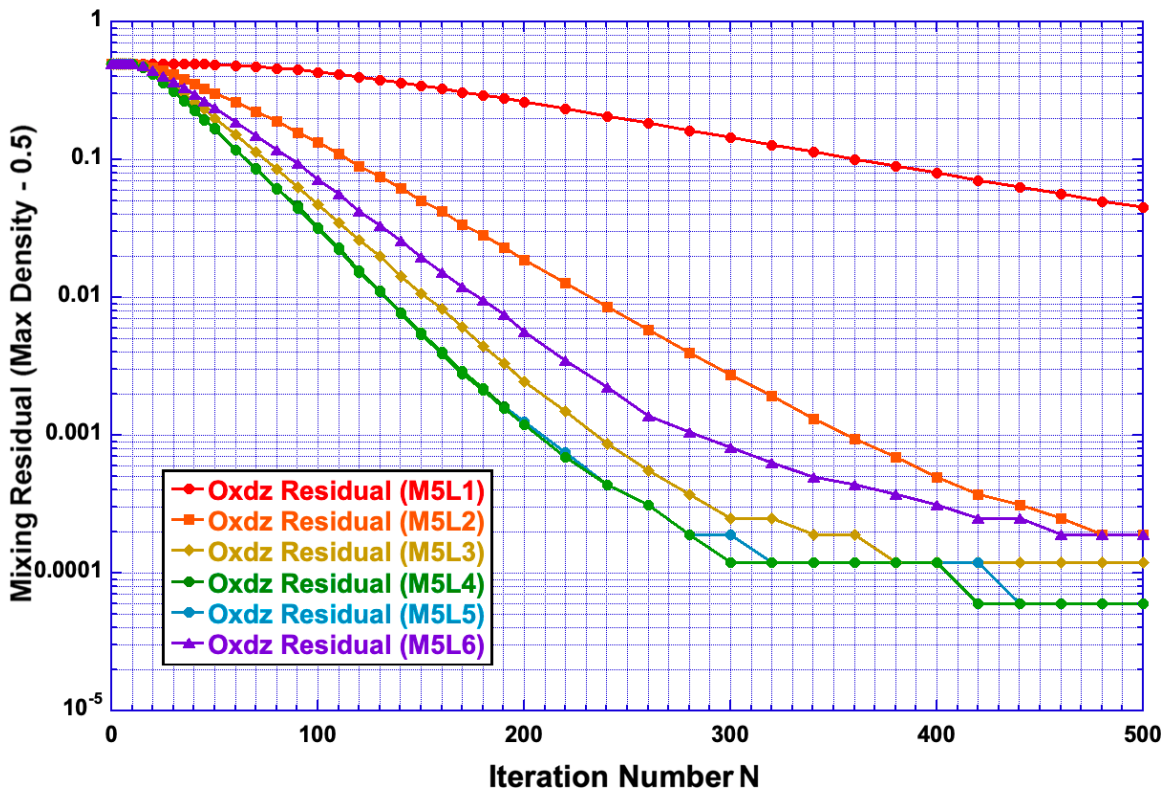


Figure 3.8 Mixing residual tests of algorithm 2, the local BCS data exchange ‘swap’ model. Oxidizer tests on the  $1024^2$  grid with  $32 \times 32$  patches (M5) using six swap patterns, L1 through L6. Algorithm 2 increases fuel and oxidizer gradients in the fluid where chemical reaction is occurring. Added numerical diffusion is NOT part of this algorithm.

The oxidizer turbulent ‘diffusion’, i.e. residual decay, recorded in Figs. 3.6 and 3.8 is much faster than for the fuel, shown in Figs. 3.5 and 3.7, because the initial-condition blocks, black and magenta, are half the size. Because of the different block scale lengths, the mixing residual becomes 100 times smaller in Fig. 3.8 than in Fig. 3.7. By iteration 300 the oxidizer residual has become so small that roundoff effects, stemming from the 16-bit precision of the densities are clearly evident. For the most transportive swap patterns, L3, L4, L5, and L6 the residuals have all approached within a small fraction of a percent of the 0.5 system-wide average. A residual of  $10^{-4}$  retains only a couple of bits of precision and roughly half this, from Fig. 3.8, seems to be the limit. Note that although the graininess of the 16-bit density representation is evident, the overall algorithm mass conservation in both species is exact.

Resolution	DE BCS time (sec)	DE BCS % time	AVG BCS time (sec)	AVG BCS % time	DE vs AVG speed up
$256^3$ Grid	16.44 s	10.5%	36.1 s	20.1%	2.2 x
$512^3$ Grid	107.7 s	11.3%	320.1 s	26.9%	2.4 x
$1024^3$ Grid	878 s	10.8%	2660 s	27.8%	3.0 x

Table 3.1 above records the run time results in computing the data for the figures and discussion in this section. Although the tests of the two local BCS algorithms, algorithm 1 (Averaging), and algorithm 2 (Data Exchange), are planar, they were performed over the entire 3D grid, replicating the planar test at each Z layer of the grid. The increase of the grid sizes accounts for the factor of 8 increase in run times. The take away from this is that the averaging algorithm is about a factor of 2 to 3 slower than the data exchange algorithm and also imposes added numerical diffusion, which costs computationally, while the data exchange algorithm does not. The next section uses these two algorithms in the baseline CSD-SFD high-speed mixing scenario from the previous report CSD#3, to see what effects enhancing the mixing with local and nonlocal baroclinic turbulence has on the overall mixing/reaction rate. The metric for Sections 4 and 5 is the decrease of the Mixed Volume Index (MVI) as a function of time.

#### 4. Simulation Campaign to study the Effects of Baroclinic Turbulence

In this section, the two local BCS algorithms, tested outside of the running SFD model in Section 3, are used here in reactive-flow mixing runs where fuel and oxidizer react instantly in grid cells that contain both. Discussion of this simple reaction approximation appears in Section 7 of the report CSD#2. The algorithm immediately reacts any fuel and oxidizer that overlap in each cell. Fuel and oxidizer can only overlap in SFD if some actual diffusion is added to the shift-based convection in SFD, by design. The algorithm used to ‘molecularly mix’ each species density  $\rho_i^n(x, y, z, t)$  on the grid scale, and implemented in the CSD#2 and CSD#3 models, is given by

$$\rho_i^n(x, y, z, t) = \rho_i^o(t) + \frac{(\rho_{i+1}^o - \rho_i^o)}{N_D} - \frac{(\rho_i^o - \rho_{i-1}^o)}{N_D} \quad \text{Eq. 4.1}$$

The present version of this surrogate ‘molecular mixing’ includes a variable factor  $N_D$ , an integer between 2 and  $\infty$ . In Eq. 4.1 superscripts o and n indicate ‘old’ (original) and ‘new’ (replaced) values of the separate species densities. When implemented with integer densities, the telescoping sum on the right ensures conservation to full round-off accuracy. Picking  $N_D = 5$  does not leave solutions on the resolved grid depending significantly on the grid resolution due to numerical error. This is shown in Section 9 of the report CSD#2 and illustrated by Fig. 9.1 there. The reaction and grid-scale mixing algorithms are spatially time split for computational efficiency and to maintain conservation to round off.

Figures 4.1 and 4.2 below treat a baseline case to which both a small nonlocal BCS spectral enhancement with  $\delta E = 4 \times 10^7$  and the two local BCS algorithms are applied for comparison. Figure 4.1 was computed entirely using the data exchange (DE) algorithm 2 on the  $1024^3$  grid with three different BCS patch sizes  $8 \times 8$  (run 11JanQ),  $16 \times 16$  (run 11JanR), and  $32 \times 32$  (run 11JanS). Each timestep that reaction takes place in a cell, the amount of fuel (and oxidizer) consumed is noted in an array and becomes the source term for the BCS approximations. Thus added transport and enhanced turbulence are scaled to the instantaneous localized source term. If the reactions were to proceed more slowly, e.g. by halving the timestep, the source terms would persist longer and the overall BCS transport would be approximately the same. The nonlocal BCS transport is implemented by augmenting the energy in each rotor bin of the resolved CSD spectrum by an amount that varies from scale to scale and is computed by the subroutine that also estimates the local time-dependent consumption of fuel and oxidizer. The algorithm ensures that the addition to the full CSD spectrum is independent of grid resolution.

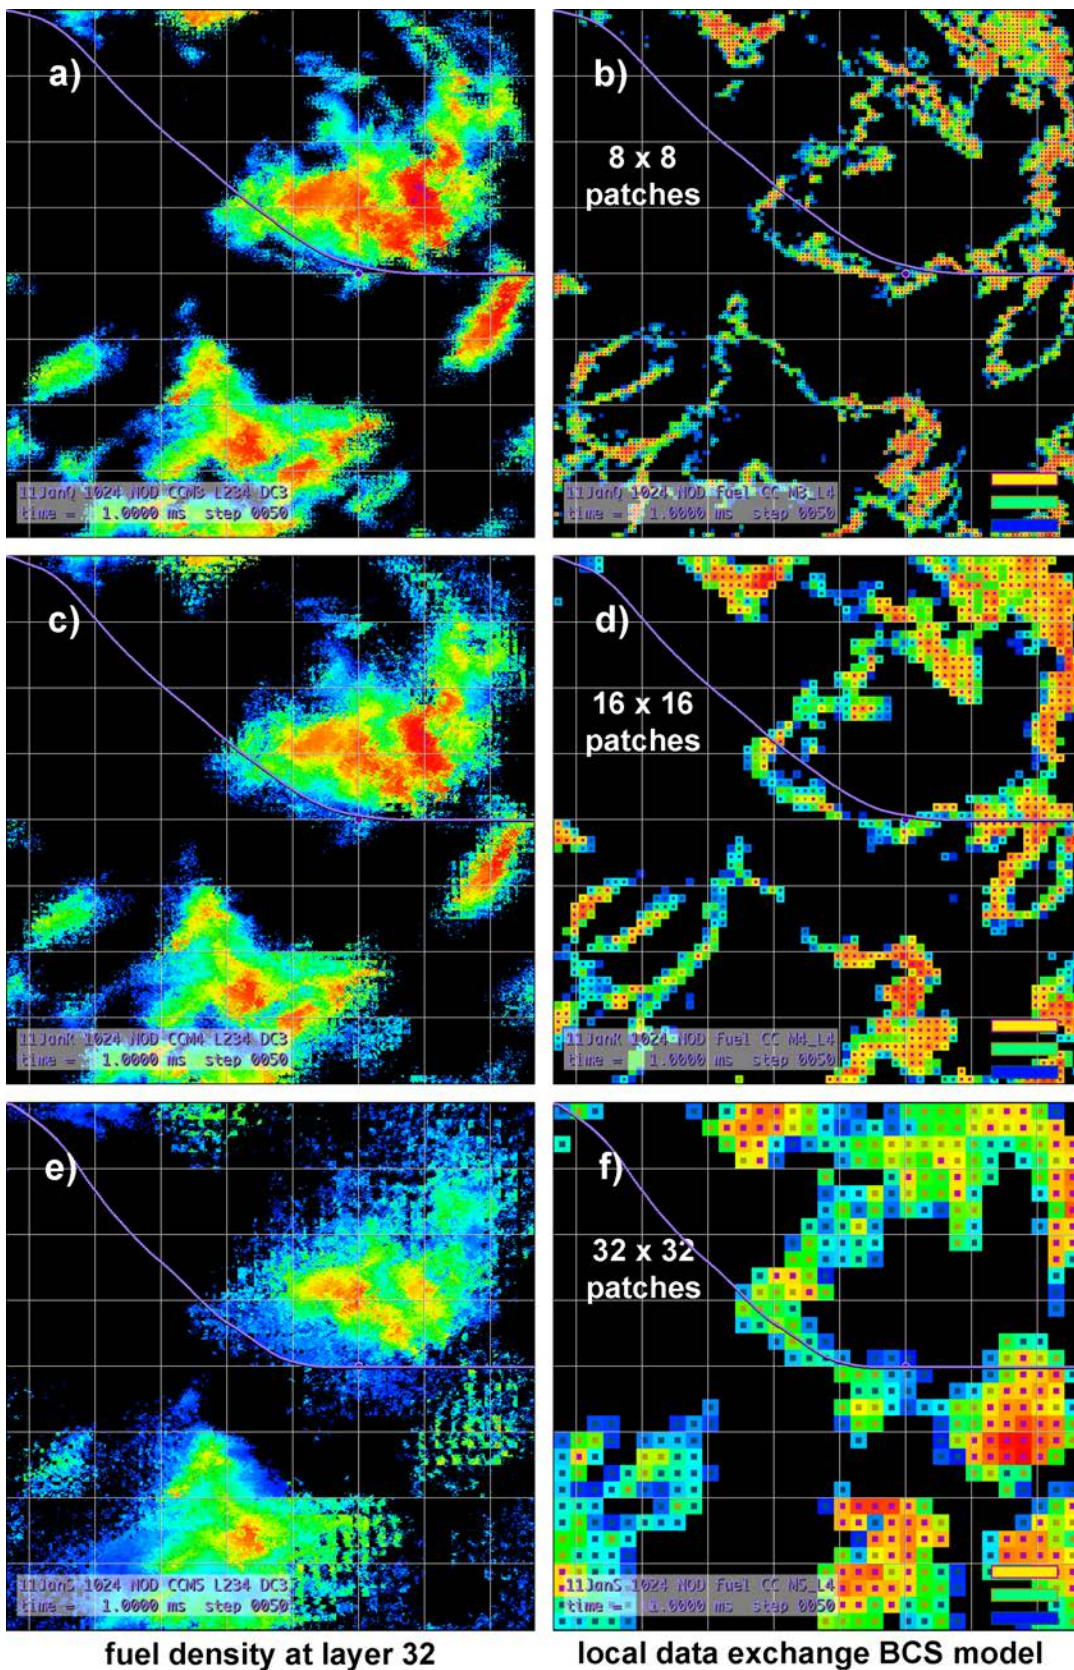


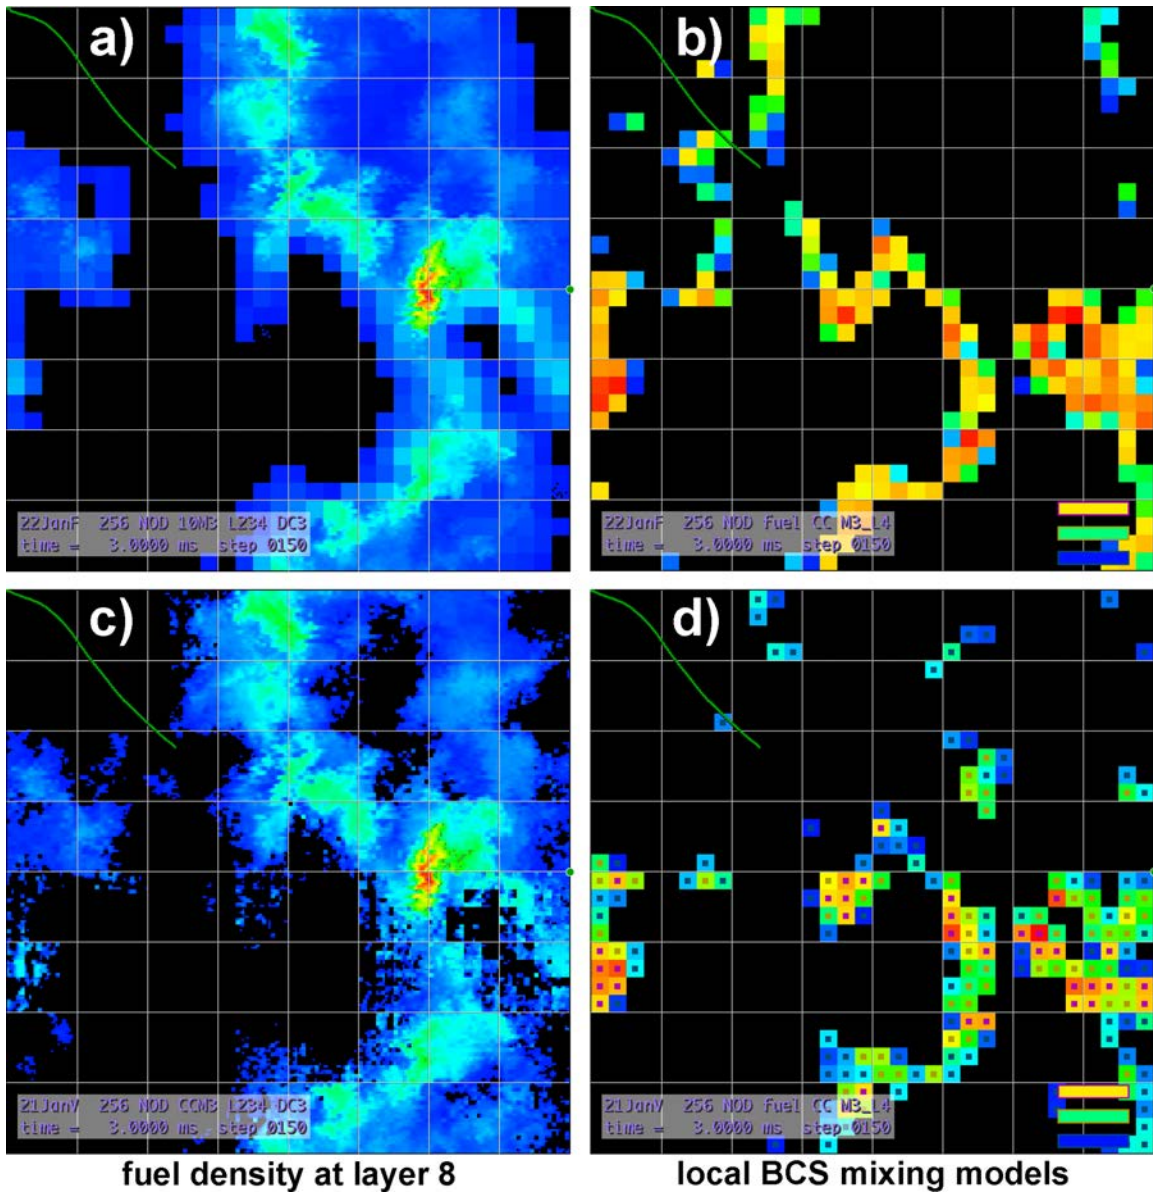
Fig 4.1 The fuel density on layer 32 density (panels a, c & e) is plotted with the corresponding local DE BCS displays (panels b, d & f) at 1.0 ms (step 50) for three patch sizes. These comparisons are computed on the  $1024^3$  grid.

Figure 4.1 above shows results of three full CSD-SFD mixing simulations using the high resolution  $1024^3$  grid and the data exchange (DE) algorithm 2. Swap patterns #2, #3 and #4 are employed depending on the amount of fuel and oxidizer that have reacted in each BCS patch. If relatively little reaction takes place, swap pattern #2 is used. Stronger reaction activates swap #3 and still stronger local reaction activates swap #4. A new diagnostic, focusing on the baroclinic source (BCS) strength, is shown in the three cross-sections plotted for Z layer 32 (panels b, d, & f on the right in Fig. 4.1). Panels a & b in the figure show the fuel density and the local data exchange baroclinic source term respectively for  $8 \times 8$  BCS patches, panels c & d used  $16 \times 16$  patches, and panels e & f used  $32 \times 32$  patches. In these full 3D simulations each  $32 \times 32$  patch accounts for about  $1 \times 10^{-6}$  of the volume and the smaller patches account for a considerably smaller volume fraction.

The active BCS patches are illustrated in the three panels on the right by a square whose color is scaled between the maximum (red) and minimum (black) BCS strength integrated over the patch. This strength is effectively the time derivative of the densities. To indicate which of the three swap patterns is active, a smaller differently-colored square is plotted in the center of each patch depending on the swap actually used. The distribution of BCS strengths in all patches is analyzed to determine a median threshold strength  $BCS_{th}$  to control which swap pattern is activated. When  $BCS_{val} \leq 0.4 BCS_{th}$ , the patch is painted black and no data exchange is performed, a considerable computational savings. Here  $BCS_{val}$  is the baroclinic source strength integrated over the cells in a 3D patch. When  $0.4 BCS_{th} < BCS_{val} \leq 0.7 BCS_{th}$ , the small indicator square is dark blue (swap #2). When  $0.7 BCS_{th} < BCS_{val} \leq 1.0 BCS_{th}$ , the small square is gold (swap #3). When  $BCS_{val} > 1.0 BCS_{th}$ , the small square is magenta (swap #4). This new diagnostic shows both the instantaneous BCS strength as well as the swap pattern employed to cause the additional baroclinic turbulent transport.

We can see numerical artifacts from the BCS patches in the right-hand panels imprinted on the low-density regions of the fuel cross-sections on the left. The higher fuel densities, i.e. red and yellow in panels a, c, & e, are generally in the center of black regions in panels b, d, & f where the BCS fuel-oxidizer reactions are below threshold. The pattern of the square patches on the right, particularly panels b & d, forms outlines of the high-density fuel, which is where the interfaces between fuel and oxidizer allow the most rapid mixing. It stands to reason that the numerical imprint on the low densities is larger when the patches are large. It is also not surprising that the high densities, where there is relatively little overlap and reaction, show none of the imprinted artifacts. Overall, the density profiles and the patterns of where reaction is occurring are similar for all three patch sizes. However, the integrated effect of the differences is also shown in the figure.

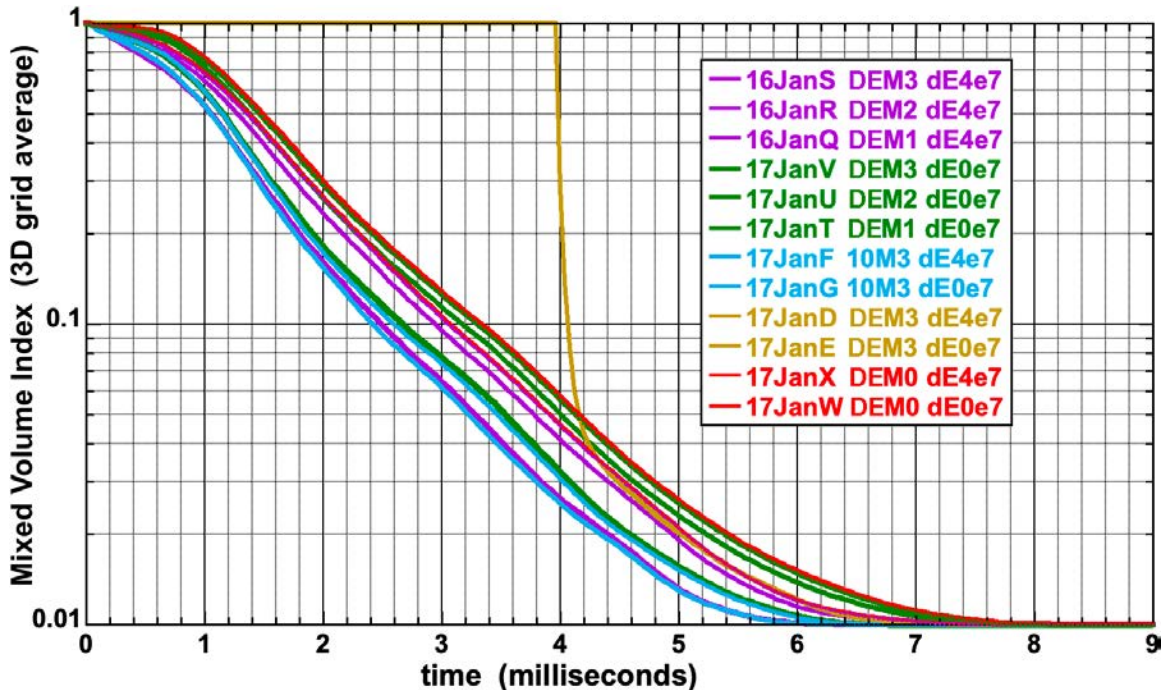
The lavender MVI(t) curves on each of the six panels extend all the way to the right edge of each panel although the cross-sections are all taken early in the run at 150 steps, 3.0 ms. MVI(t) was plotted on Fig. 4.1 from the continuation of each run to 10 ms when all of the oxidizer and all but 1% of the fuel have been consumed. Notice that MVI(t) reaches the 0.01 asymptote in panel e about 1.25 ms faster than in panel c and about 2.5 ms faster than in panel a. A small circle on the horizontal midline of each cross-section marks a fixed location in the laboratory frame of the rapidly moving fluid. The higher the MVI(t) curve, the less the amount of added mixing due to baroclinic turbulence.



**Figure 4.2** The fuel density and the fuel/oxidizer baroclinic source at 3.0 ms (150 timesteps) for two BCS algorithms at grid resolution  $256^3$ . The baroclinic source data, panels b and d, are for BCS algorithm 1 and BCS algorithm 2 respectively. The corresponding fuel densities are shown in panels c and d. Layer 8 corresponds to layer 32 on the  $1024^3$  grid.

Figure 4.2 just above compares the local baroclinic turbulence generated by the averaging algorithm (top – panels a & b) and the data exchange (swapping) algorithm (bottom – panels c & d) after 150 timesteps and 3.0 ms, as opposed to Fig. 4.1 considering only the data exchange algorithm. In Fig. 4.2, the MVI(t) curve is plotted in green, indicating the  $256^3$  grid, and only extends for the simulated 3.0 ms rather than 10.0 ms, corresponding to the right edge of each cross-section shown. For Figs. 4.2 and for 4.3 below, the grid is only  $256^3$  to magnify the small-scale differences. This resolution had been previously shown, e.g. Fig. 1.2 and previous reports, to give answers quite close to those on the higher resolution grids but the lower resolution here better exposes the numerical effects. Panel a shows the blue imprint of the BCS patches in the lower fuel densities (bordering on black) far more extensively than panel c below it.

Despite the total difference in the algorithms, the locations where BCS turbulence generation is strongest are very similar in panels b and d above. The averaging algorithm, which shows more of a source in the upper third of the panels, also shows a slightly faster decay of the MVI(t) in the two green curves in panels b and d. Figure 4.3 below continues the comparison of the AVG algorithm1 and the DE algorithm 2 on the 256<sup>3</sup> grid.



**Figure 4.3** The Mixed Volume Index MVI(t) of reacting fuel and oxidizer illustrates the effects of local and non-local baroclinic turbulence for 12 sets of model parameters. Two algorithms for the local baroclinic source of turbulence (BCS), a simple model for nonlocal BCS, and two runs where grid-scale diffusion is zeroed (gold) are generated on the 256<sup>3</sup> CSD-SFD grid.

Figure 4.3 plots the Fuel Mixed Volume Index, MVI(t), against the simulation time for 12 simulations with the CSD-SFD model. All 12 runs were done on the 256<sup>3</sup> grid to compare different scenarios for the local BaroClinic Source (BCS) algorithms and for the nonlocal BCS algorithm. When a full 3D SFD simulation is run, the Mixed Volume Index is a more appropriate metric for the entire model than the Mixing Residual, used in previous sections to summarize isolated tests of the BCS averaging (AVG) algorithm 1 and the data exchange (DE) algorithm 2. Using MVI(t) also allows comparing simulations and discussions vis-a-vis the previous three CSD-SFD reports.

The strength of the additional BCS transport applied by the DE algorithm is controlled by the thresholds for application of the swap patterns 2, 3 and 4 in each of the 8 x 8 patches (M1 runs), 16 x 16 patches (M2 runs), or 32 x 32 patches (M3 runs) in the figure. The strength of the BCS transport from the AVG algorithm in these preliminary tests is controlled by the size of the patch only. Runs were conducted with the nonlocal BCS turned off ( $E = 0.0 \times 10^7$ ) and with the added strength coefficient set to  $\delta E = 4.0 \times 10^7$  for each of the three resolutions.

Runs 16JanS, R & Q with  $\delta E = 4.0 \times 10^7$  (lavender) and runs 17JanV, U & T with  $\delta E = 0.0 \times 10^7$  (green) are versions of the baseline scenario where the patch size is varied (i.e. M = 3, M = 2 and M = 1). In these cases the larger patches uniformly result in more baroclinic

transport, making the Mixed Volume Index decrease more quickly. The three lavender curves are all lower than the corresponding green MVI(t) curves due because the nonlinear contribution to the baroclinic transport increases turbulent mixing.

Runs 17JanD & 17JanE are modifications of the baseline scenario where the grid scale diffusion is set to exact zero for the first 200 timesteps of each run and then is turned back on at the nominal level as in all other runs of the campaign. These two solutions, are plotted in gold on Fig. 4.3 and are essentially indistinguishable. Both solutions show that MVI(t) stays constant at 1.0, verifying zero interdiffusion of the fuel and oxidizer until step 200. When the limited grid-scale diffusion is turned on, the MVI(t) curves drop very rapidly to the values they would have had as if the ‘molecular’ scale mixing had been turned on for the entire run. Tests verifying that SFD, without the grid-scale effect of Eq. 4.1, is indeed diffusion-free are discussed in more detail in previous reports.

Runs 17JanX & 17JanW (red) with and without nonlocal BCS effects also have the local data exchange contributions turned off (M0). Thus run 17JanW is the baseline case on the  $256^3$  grid. Runs 17JanF & 17JanG (light blue) are runs using the averaging algorithm 1 (AVG), rather than the DE algorithm 2, but with and without the nonlocal BCS contribution. The patch size for runs 17JanF & G was  $8 \times 8$  (M3), exactly the same size as the  $32 \times 32$  patches (M5) on the  $1024^3$  grid. The notation ‘10’ in the legend for these runs means that the AVG BCS transport used the full average over the patch and was not diluted by including any of the initial condition. The freedom to blend the full average with the undiffused initial condition provides a way to make the AVG BCS contribution nearly continuous in strength. However, because the DE algorithm 2 is better for several reasons, the AVG algorithm 1 is not currently being further pursued.

Figure 4.4 below compares the Mixed Volume Index, MVI(t), for a number of runs at three different resolutions testing the data exchange model (DE algorithm 2) with two patch sizes both with and without the non-local baroclinic turbulence contribution. The patch sizes were chosen on each of the three grids to correspond to patches of the same size. For example, M3 ( $8 \times 8$ ) patches at  $256^3$  are the same physical size as M4 ( $16 \times 16$ ) patches at  $512^3$ , and as M5 ( $32 \times 32$ ) patches on the  $1024^3$  grid. A 13th run, RunW, is included with no baroclinic contribution, and thus is the baseline case at  $1024^3$  resolution. Run W is plotted as a dashed black line in the figure with a larger MVI(t) value above the other 12 solutions. Clearly the inclusion of even simple approximations to the transport arising from baroclinic turbulence generation always increases the mixing of fuel and oxidizer and speeds the chemical reaction.

In Fig. 4.4, the pairs of runs with the two darker green, darker blue, and purple lines always generate more baroclinic turbulence than the solutions with lighter green, light blue and lavender lines because the patch sizes are larger and cause more transport at each resolution. The ratio between MVI(t) on the baseline upper curve and on the fastest reacting lower curve is 1.0 both initially and finally but maximizes at about 2.8 after 4 ms. For RunS (purple), the fastest reacting case, half the fuel and oxidizer mix and react in the first millisecond while run W (no baroclinic turbulence and thus the slowest reacting baseline case), takes 1.6 ms to reach the halfway point. The fastest case takes 2.4 ms to consume 90% of the reactants while the slowest takes 3.5 ms to reach 90%. The fastest case takes 5.5 ms to reach the asymptote while the slowest takes about 8 ms. In the central region of the plot, the curves show an approximately exponential decay, echoing what can be seen in Fig. 3.7 above (Section 3) where the isolated tests of the local BCS (DE) algorithm 2 on the fuel are reported.

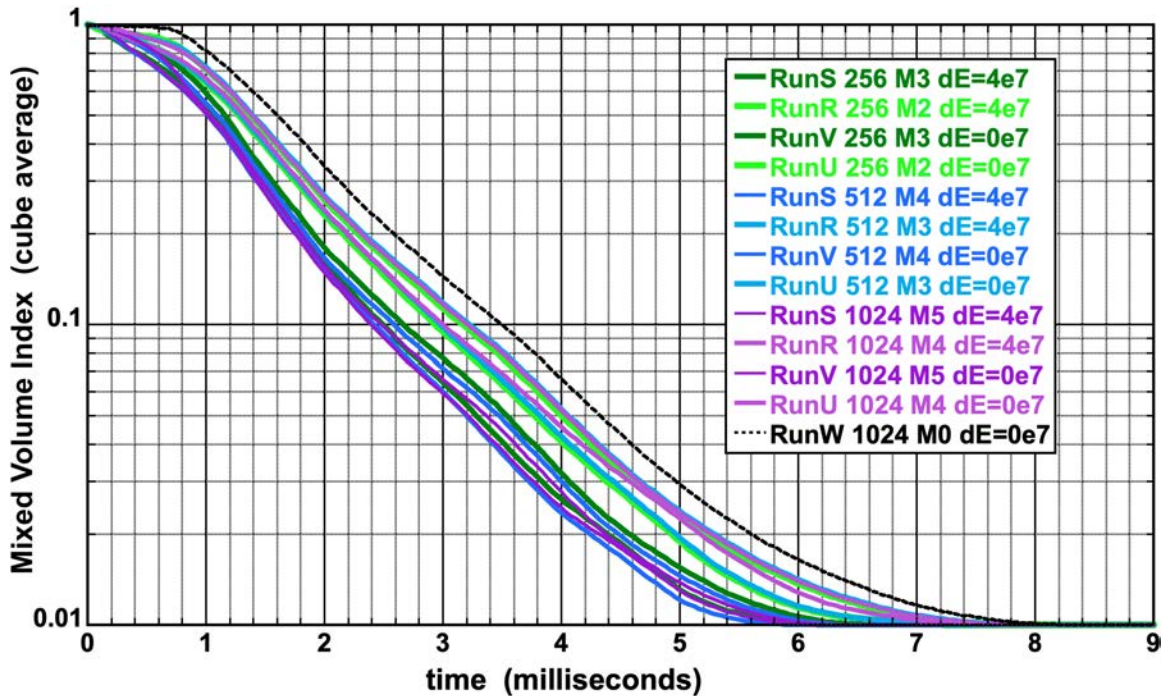


Figure 4.4 Comparing the Mixed Volume Index (residual fuel fraction) simulated at three grid resolutions using four baroclinic turbulent sources at each resolution. The non-local baroclinic source term is included when  $dE = 4e7$ . Run W (dashed black curve) is the baseline run with no BCS turbulence.

The same kinds of scenarios reported in Fig. 4.4 are considered further in Fig. 4.5 below. The MVI(t) profiles from eight CSD-SFD simulations highlight the expected increase in turbulent mixing from combustion-induced expansion. The MVI(t) in Fig. 4.5 curves all have the same general form as observed in Figs. 4.3 and 4.4 above. They start at 1.0 and then asymptote to 0.01 after 5 to 8 ms due to a 1% fuel excess in the initial condition. The vertical spread of the curves in Fig. 4.5 shows an approximate range for the effect of the BCS acceleration of mixing and the resulting reaction. Initially there is no effect but by 3.0 ms the MVI values vary by a factor of about 2.5. The unenhanced baseline turbulent mixing, shown in the RunW upper black curve, accounts for only 40% if the mixing seen in the lowest curves V (gold) and S (medium blue). By 4.0 ms, the vertical spread caused by combustion-driven baroclinic turbulence is about at its largest, a mixing factor ratio around 2.7. By 6.0 ms the mixing factor has diminished to 1.6 and reaches 1.0 again for all runs by 8.0 ms.

All runs plotted in Fig. 4.5 used the highest-resolution grid and are color coded according to the figure legend. Runs Q, R, S and X with  $\delta E = 4.0 \times 10^7$  have nonlocal BCS turbulence added to the resolved scales of the spectrum to represent the effects of baroclinic turbulence outside the regions where combustion is occurring. Runs T, U, V and W, even though the local BCS terms are present, have no added nonlocal mixing. Run W (solid black curve) and Run X (dark blue) have the data exchange local model, algorithm 2, turned off but the non-local BCS is active in Run X. The difference between them is significant. In Fig. 4.5, two sets of three runs demonstrate again the effect of larger patch sizes leading to more turbulent mixing. Run S (gold) is lower than (i.e. faster reactions) run R (orange) which is lower than run Q (red). When the nonlocal term is turned off, run V (blue) is lower than run U (light blue) which is lower than run T (green).




Figure 4.5 Enhanced turbulent mixing with different strength BCS calculations computed on the  $1024^3$  grid. Simulations of surrogate fuel-oxidizer mixing associated with chemical reactions, using 3 patch sizes, 3 data exchange patterns and including zero BCS turbulence, show enhanced mixing.

To summarize the results of this section, the nonlocal BaroClinic Source (BCS) algorithm and the two local BCS algorithms behave as expected. Some evidence has been presented, including improved execution speed and reduced numerical diffusion, for selecting the data exchange algorithm over the averaging algorithm. Increasing the patch size for application of the local BCS source increases the rate of mixing over that found in the baseline scenario by increasing the additional distance the reactants are transported in areas where the chemical reactions are occurring. The rate of mixing is quite insensitive to the resolution of the SFD grid as long as the geometrical patch size is held fixed when the grid resolution is changed. Further in all scenarios faster mixing occurs than in the baseline scenario. The goals for CSD-SFD have been met in allowing demonstration of convergence with resolution without statistical uncertainty clouding the results of usual CFD approaches, which also suffer from nonphysical numerical diffusion.


## 5. Uncertainty: Multiple Realizations of Turbulent Flow

Turbulent flows are intrinsically variable and uncertain. Aleatoric uncertainty, also known as statistical uncertainty, results from unknowns that differ each time we measure the same physical scenario. Epistemic uncertainty, also known as systematic uncertainty, results from things we might know if we worked harder. Turbulence is in the former category. To clarify issues related to numerical convergence, CSD-SFD removes some of this uncertainty by design.

This section considers aleatoric uncertainty in the turbulent mixing problem we have been considering. The physics of turbulent flow means that a statistically well specified scenario nevertheless has an infinite ensemble of possible realizations because the details of the initial conditions are uncertain. Simulations of situations such as fluctuating winds, must replicate, at least in some sense, this infinite ensemble of realizations. The alternative, if it really is one, is to solve a truncated problem usually by averaging the equations in some way that reduces the physics to a situation with a single solution. An example is seeking a time-independent steady ensemble average solution. This means somehow averaging the problem before finding the solution. When the problem is nonlinear, however, the average of the fluctuating turbulent solution is never mathematically the same as solutions of the mathematically averaged equations.

For CSD-SFD, the time-dependent turbulence spectrum from CSD is the same for each realization, regardless of the grid resolution. The realizations are distinguished in a different way, by varying the table of random numbers that determines the direction of the flows that realizes the given rotor energy in each of the logarithmic CSD energy bins. Each flow realization has the same energy as a function of time in each scale-size bin as the other realizations but they are all different because the directions of the velocities in each scale are different. This means that other space or time averages over a realization may and probably will differ from realization to realization.

Figure 5.1 below compares solutions with the 3 grid resolutions after step 250, i.e. 5.0 ms into the runs. The six simulations used in the figure were all performed with a nonlocal contribution to the BCS with strength  $\delta E = 4.0 \times 10^7$ , as used in many of the simulations of earlier sections. The local BCS is computed using the Data Exchange model with swap patterns 2, 3, and 4 activated as previously. The top row of panels shows three different resolution calculations using the particular set of random numbers defining realization 0. The middle and bottom rows of panels show the same three resolutions using another set of random numbers defining realization 5. In all, six realizations were computed at each of the three grid resolutions,  $256^3$ ,  $512^3$ , and  $1024^3$  cells. SFD calculates the deterministic turbulent spectrum and correlated random convective shifts up to the highest resolution computationally possible specifically so that shorter wavelengths appearing in the computed flow at higher resolution do not result in different flow at larger scales. Thus a set of different-resolution CSD-SFD computations of a single physical scenario all have the same long-wavelength behavior while the use of different random-number tables for each realization still allows the evaluation of imposed aleatoric uncertainty from the repeatable, deterministic computations.

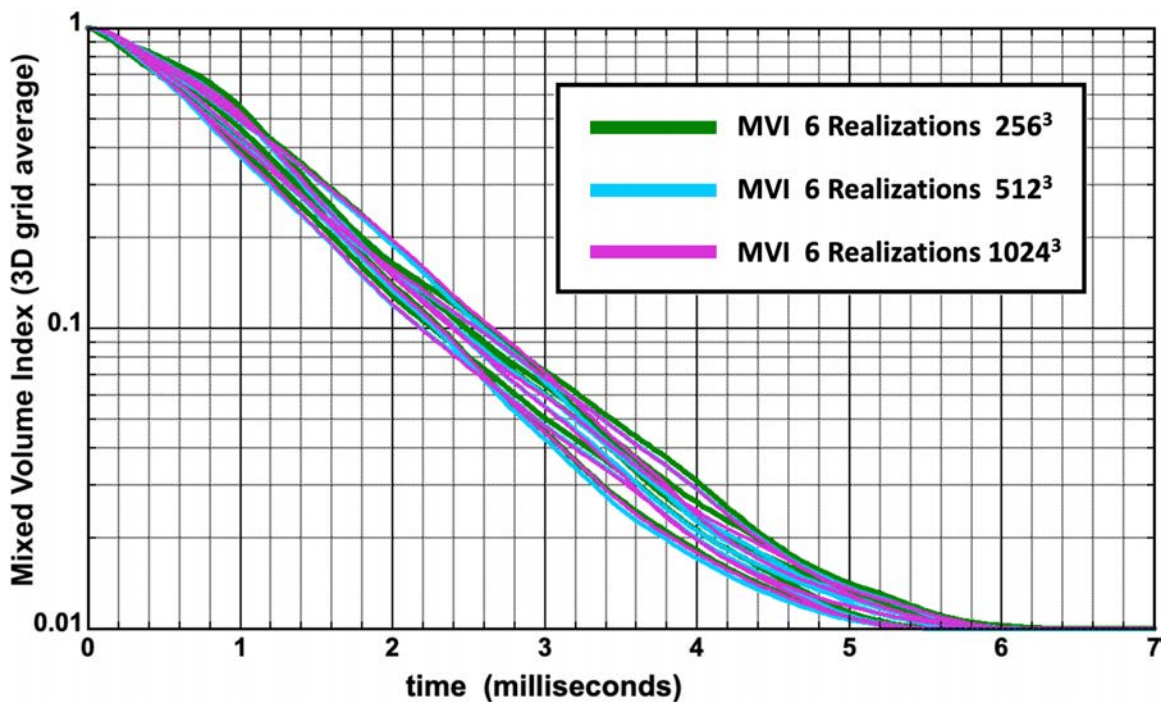


**Figure 5.1** Comparing the CSD-SFD Mixed Volume Index (residual fuel fraction) at three grid resolutions for two turbulent flow realizations after 5.0 ms. Both local and non-local baroclinic turbulence are included in each flow. Panels a, b & c: Realization 0 - fuel on Z layers 8, 16 & 32 for the  $256^3$ ,  $512^3$  &  $1024^3$  Grids. Panels d, e & f: Realization 5 - fuel. Panels g, h & i: BCS data exchange patches used for Realization 5.

Panels a, d & g on the left in Fig. 5.1 used the  $256^3$  grid with the Mixed Volume Index plotted in green. Panels b, e & h used the  $512^3$  grid with MVI(t) plotted in blue. Panels c, f & i used the  $1024^3$  grid with MVI(t) plotted in lavender. The three bottom panels in the figure plot the new BCS patch display showing the reaction-driven Data Exchange BCS patches of baroclinic turbulence where data exchange was active at step 250. Where the patches are black, no BCS data exchanges had to be computed, as in the previous section, thus saving most of the computational work. Data exchange ('swap') patterns 2, 3 & 4 were used depending on the instantaneous BCS source strength, indicated by color of each individual patch. The small

colored squares in the center of each patch indicate the data exchange ('swap') pattern chosen for that patch. This method to mark which swap pattern is active in each patch, using a dark blue, gold, or magenta square, is described above following Fig. 4.1. That swap pattern 4 provides the most transport of the swaps 2, 3, & 4 – is shown clearly in Figs. 3.7 and 3.8 in Section 3. Careful calibration with detailed experiments or very expensive detailed DNS would be needed to develop a set of conditions for choosing the different swap patterns quantitatively. That is not the goal of the development tests being reported here and may not even be a easily-realized goal of SFD developments in the near future. Rather, CSD-SFD is a surrogate model designed to dip into new parameter regimes not accessible with current CFD methodology on current computer systems.

Viewing Fig. 5.1 from an overall perspective shows that the solutions converges well with resolution since the images along each row, taken from the corresponding Z layer on each grid, are all very similar and their MVI(t) curves are basically identical. Six different realizations of the flow were then computed for each of the three grids. These 18 runs all used the same CSD turbulence spectrum and the realizations, numbered 0 through 5, all used different random-number tables for each realization, independent of the grids resolution. Figure 5.2 immediately below plots the MVI(t) curves for these 18 runs using a separate color for each of the three resolutions. There is appreciable overlap among all the realizations and the three resolutions.



**Figure 5.2 Comparison of the Mixed Volume Index (the residual fuel fraction) simulated by the CSD-SFD model at three grid resolutions for a number of separate turbulent flow realizations at each resolution. Terms accounting for the local and non-local baroclinic generation of turbulence driven by the heat-release pattern in each flow are included.**

The Mixed Volume Index, MVI(t), is computed for the entire triply-periodic grid by summing the fuel remaining in the system at time t, here 0 to 7 ms, and dividing by the total fuel initially in the system. As remarked earlier, this baseline scenario has 1% more fuel than oxidizer so MVI(t) approaches 0.01 and can drop no further for each of the 18 separate realizations recorded here. The aleotric uncertainty documented here is quite generally less

than a factor of two at any given time and seems to maximize at about 1.8 around 4 ms, as before. If the root-mean-square difference from the mean is used as the metric, the 6 realizations on the  $256^3$  grid have an average MVI of .0238 at 4.0 ms and an RMS deviation from this average of 0.0039. On the  $512^3$  grid, the average MVI is .0223 at 4.0 ms with an RMS deviation from this average of 0.0039. On the  $512^3$  grid, the average MVI is .0223 at 4.0 ms and an RMS deviation from this average of 0.0037. On the  $1024^3$  grid, the average MVI is .0236 at 4.0 ms with an RMS deviation from this average of 0.0038. At each resolution the RMS deviation from the mean, at 4.0 ms, is only  $\pm 16 - 17\%$ .

The simulations considered earlier, for example those leading to Figs. 4.4 and 4.5, range between  $MVI(t) = .023$  and  $.067$  at 4.0 ms. This is about  $\pm 50\%$  with an RMS spread of perhaps  $\pm 35\%$ . This is about twice the range of variation between different aleotoric realizations of what is physically the same scenario. The baroclinic turbulence algorithms are actually contributing in the expected directions in excess of statistical uncertainty. There is no claim, at this time, that the enhanced mixing and its contribution to the consumption of fuel and oxidizer is quantitatively correct, however, the effects have been demonstrated.

Beyond this modest conclusion, this report has also illustrated the larger issue of extracting meaningful results from standard CFD computations no matter how accurate the algorithms profess to be. In CFD, the fluid dynamics at the smallest resolved scales depends intrinsically on the grid resolution in high Reynolds-Number turbulence. It doesn't take long for 100% differences in the small scales, from one resolution to another, to percolate out into the larger scales which the two resolutions have in common. These 100% differences in the larger scales of CFD, arising from the dependence of these larger scales on the velocity field at the grid scale, swamp the 16-17% differences from the scenario aleotoric uncertainty, shown in Fig. 5.2, and the 35% changes due to the physics being varied between scenarios. This section used 6 realizations at the 3 grid resolutions to show the scale of the aleotoric uncertainty in the baseline scenario. Wind tunnel experiments, where performing many repeated measurements of fluctuating but statistically identical scenarios is possible, suggest that hundreds of realizations are necessary to begin to show convergence of important fluid dynamics quantities measured through the uncertainty (e.g. Harms et al., 2011).

SFD does not cause the inverse cascade up to larger scales in the turbulence spectrum seen in CFD attempting to reach high Reynolds numbers. That physically necessary property enters these scenarios through the Coherent Structure Dynamics (CSD) model, which solves for the turbulence spectrum over a range of scales far wider than the SFD grid scale or any CFD scale. In SFD, even coarser-grid solutions contain the inverse cascade in the same way as solutions on much finer grids. This means that the surrogate turbulence computed by SFD can converge with increasing resolution. This allows exposure of subtle effects, which are present in physical turbulence but hidden under the large-scale consequences of grid noise and statistical uncertainty in CFD.

## 6. Discussion and Future Work

The combined Coherent Structure Dynamics – Surrogate Fluid Dynamics (CSD-SFD) model, as described here and in earlier reports, now contains new algorithms to approximate the enhanced mixing and faster reaction caused by baroclinic turbulence due to expansion from the combustion of fuel and oxidizer. Several numerical approaches and their algorithms to estimate this enhanced mixing have been tested for use in SFD. Initially, two local BaroClinic Source (BCS) mixing algorithms and a nonlocal BCS algorithm were tested in standalone scenarios conducted to assess the new models without the large-scale flows in SFD obscuring their behavior. These BCS algorithms determine where, when and how much baroclinically-induced transport should be added to the existing SFD turbulent transport. The tests showed that the local BCS models displayed an extended period of exponential decay, similar to a spatial diffusion model, and that different local strengths of BCS transport can be imposed. How these baroclinic, combustion-enhanced mixing algorithms should be calibrated is a subject for future work where experiments and detailed reacting-flow CFD simulations for comparison will have to be used.

After the standalone tests of the BaroClinic Source algorithms, several computations were shown where the BCS algorithms, though still uncalibrated, were used in full CSD-SFD simulations. These 3D time-dependent simulations, presented in Section 4, all showed the expected qualitative behavior, that is enhanced mixing with faster reaction. Larger patches applying the local models caused more rapid overall mixing and reaction. Stronger nonlocal BCS transport, caused slightly more rapid mixing. Further, the ability of SFD to allow turbulence to transport species without any ‘molecular mixing’ (grid-scale diffusion) is maintained by the data exchange (DE) local model, as shown by the gold curves in Fig. 4.3. The range of full simulations, leading to Figs. 4.3, 4.4 and 4.5, all had Mixed Volume Index values between .023 and .067 after 4.0 ms of mixing. This is about  $\pm 50\%$  with an RMS spread of perhaps  $\pm 35\%$ . This is about twice the 16-17% MVI(t) variations in the statistical ensemble of multiple aleotoric realizations for the same physical scenario. The baroclinic turbulence algorithms are actually contributing in the expected directions in excess of statistical uncertainty, which in CFD would have obscured the results.

Although the physical effects of enhanced BCS mixing have been demonstrated in the context of SFD, no claim is being made that its contribution to the consumption of fuel and oxidizer is yet quantitatively correct. This paper is only a research progress report because calibration of the local data exchange model, suggested for adoption here, has yet to be performed. An obstacle to this goal is the simple, idealized configuration of the CSD-SFD triply period cube grid. An even more immediate need is to extend the local BCS model to perform a range of simple data exchanges where the amount of the baroclinic transport each timestep is a continuous monotonic function of the local BCS energy release.

The generalized approach that currently suggests itself is to exchange data for only a fraction of the  $2^M \times 2^M$  patch (2D) or  $2^M \times 2^M \times 2^M$  patch (3D). When  $M = 1$ , all the data in the patch, 8 values in 3D, have to be swapped along the diagonals. 8 or 0 is the only choice without the extra computational cost to implement fractional swaps. When  $M = 2$ , each 3D patch has 64 cells and thus there are eight possibilities beyond 0 data exchange. The 8 corner cells can be swapped. 16 cells can be swapped along diagonals. 24 cells can be swapped symmetrically adjacent to the diagonals. Or 32 cells, 40 cells, 48 cells, 54 cells, or 64 cells could be swapped

all without the swap asymmetry introduced by swap pattern #3 in Fig. 3.4. This affords a nine-level algorithm. When  $M = 3$ , the patch size in 3D is  $8 \times 8 \times 8 = 512$  cells, giving perhaps 64 different levels. The test problem in Section 3, leading to Figs. 3.2, 3.3, 3.5, 3.6, and 3.7, would have to be extended to 3D, or possibly the 2D version could be simply be implemented on every plane. The fastest, most transparent algorithm for doing this is not obvious.

There is also a more general problem because baroclinic turbulence generation also happens when passive density gradients are accelerated by large-scale rotations, by gravity, or by shocks in addition to the accelerations associated with expansion from chemical reactions in a compressible fluid. The first two of these are basically incompressible but the last two are not. The baroclinic vortex generation associate with passive density gradients is not part of the initial SFD implementation which treated an incompressible surrogate fluid where the individual components being mixed were implicitly assumed to have the same density since a local momentum equation was not being solved, momentum conservation being assured globally. This can be rectified, however, by extending the BCS source calculation to include the existing passive density gradients as well as those being generated by combustion-driven expansion. The velocity field used in SFD is 3D and time-dependent and so can be used to compute the instantaneous acceleration of the surrogate fluid for use in the baroclinic source term on the right in Eq. 2.1. This is a project not yet initiated.

The fidelity of the CSD-SFD numerical model used here merits further discussion. Although the simulations are clearly surrogates for high Reynolds-number fluid dynamics, they are not engineering-grade computations. The extensions and calibrations discussed just above have not yet been performed. Even when they are, SFD will still be a non-quantitative surrogate model. CSD-SFD scenarios are valuable because they permit study of high-speed and non-equilibrium phenomena in parameter regimes that are not being accessed by other methods.

The CSD-SFD model cannot be derived mathematically from the Navier-Stokes equations. It is constructed by stitching together a number of apparently ad hoc pieces chosen for their computational efficiency and simplicity. The anisotropy possible in detailed rotor physics is, for the present, ignored and the computational discreteness is extreme, allowing questions about the basic reality and realizability of these flow surrogate flows. However, the key principles of the physics are embedded. The underlying convection is free of numerical diffusion and entirely conservative. The approach respects causality. Any species moving from A to B must pass through locations that are in between. The SFD convection is exactly mass and momentum conservative. The CSD turbulence spectrum evolution equations are energy conservative. There is no teleportation of material and information, as in implicit algorithms, and the scale-invariance of implicit fluid dynamics is replicated in the CSD equations for evolution of the spectrum. The surrogate approach revolves entirely around respecting the boundary conditions and principles of the physics and seeing what can be predicted from that. This approach has the benefits of a much wider range of space scales than CFD and a much faster execution of the scales which are gridded.

## Acknowledgements

The initial 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 these reports. Our 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

- B.D. Bak and T. Kalmar-Nagy, 2018, “A Linear Model Of Turbulence: Reproducing the Kolmogorov-Spectrum,” *IFAC PapersOnLine* 51-2 (2018) 595-600.
- G.A. Batley, A.C. McIntosh, J. Brindley and S.A.E.G. Falle, 1994, “A numerical study of the vorticity field generated by the baroclinic effect due to the propagation of a planar pressure wave through a cylindrical premixed laminar flame,” *Journal of Fluid Mechanics* **274**, 217–237, 1994.
- G.A. Batley, A.C. McIntosh, and J. Brindley, 1996, “The baroclinic effect in combustion,” *Mathematical and Computer Modelling* **24** (8), pp. 165–176, 1996 .
- L. Biferale, 2003, “Shell Models Of Energy Cascade In Turbulence,” *Annual Reviews of Fluid Mechanics* **35**, 441–468, 2003: doi: 10.1146/annurev.fluid.35.101101.161122
- J.P. Boris, CSD#1, 2018, “A Coherent Structure Dynamics Model for Non-Equilibrium Turbulence,” Naval Research Laboratory Memorandum Report, NRL/MR/6003–18-9815, 14 September 2018.
- J.P. Boris, CSD#2, 2021, “A Mixing Study Using Coherent Structure Dynamics to Drive a Surrogate Fluid Dynamics Model,” Naval Research Laboratory Memorandum Report, NRL/6003/MR–2021/1, 29 July 2021.
- J. Boris and K. Obenschain, CSD#3, 2022, “Mixing Delays in Non-Equilibrium High-Speed Turbulence,” Naval Research Laboratory Memorandum Report, NRL/6003/MR–2022/1, 9 February 2022.
- J.P. Boris and J.M. Picone, 1982, 1983, “Vorticity Generation by Asymmetric Energy Deposition in a Gaseous Medium,” NRL Memorandum Report 4854, 16 September 1982; also J.M. Picone and J.P. Boris, *The Physics of Fluids* **26**, 365–382, 1983.
- K. Bray, A. Libby and J.B. Moss, 1985, “Unified modelling approach for premixed turbulent combustion-Part I: General formulation,” *Combustion and Flame* **61**, 87–102, 1985.
- K. Bray, 2016, “Laminar Flamelets in Turbulent Combustion Modeling,” *Combustion Science and Technology* **188**,(9), 1372–1375, 2016, <http://dx.doi.org/10.1080/00102202.2016.1195819>.
- B.F. Farrell, P. J. Ioannou, 2009, “A Theory of Baroclinic Turbulence,” *Journal of the Atmospheric Sciences* **66**, 2444–2454, 2009.

- B. Gallet and R. Ferrari, 2020, "The Vortex Gas Scaling Regime of Baroclinic Turbulence," *The Proceedings of the National Academy of Sciences (PNAS)* **117**(9), 22 October 2020.
- J.R. Greig, R.E. Pechacek, M. Raleigh, J.M. Picone and J.P. Boris, 1983, "Channel Cooling by Turbulent Convective Mixing (U)," Proceedings: DARPA Workshop on Relativistic Electron Beam Heating in Air, W. Ali (ed), 15–16 March 1983.
- F. Harms, D. Hertwig, B. Leidl, M. Schatzmann, and G. Patnaik,, 2011, "Characterization of Transient Dispersion Processes in an Urban Environment," HARMO14 – 14th International Conference on Harmonization within Atmospheric Dispersion Modelling for Regulatory Purposes, 2-6 October 2011, Kos Greece.
- J.M. Picone, J.P. Boris, J.R. Greig, M. Raleigh and R.F. Fernsler, 1981, "Convective Cooling of Lightning Channels," *Journal of the Atmospheric Sciences* **38**(9): pp. 2056–2062, September 1981.
- J.M. Picone and J.P. Boris, 1982, "Multidimensional Numerical Methods for the Simulation of Shock-Generated Turbulence," NRL Memorandum Report 4967, 30 December 1982.
- J.M. Picone and J.P. Boris, 1983, "Vorticity Generation by Asymmetric Energy Deposition in a Gaseous Medium", *Physics of Fluids*, **26** (2) 365-382, 1983.
- J.M. Picone, E.S. Oran, J.P. Boris and T.R. Young Jr., 1984, "Theory of Vorticity Generation by Shock Wave and Flame Interactions," in **Dynamics of Shock Waves, Explosions and Detonations**, edited by J.R. Bowen, N. Manson, A.K. Oppenheim, and R.I. Soloukin, (AIAA, New York, 1984), also Naval Research Laboratory Memorandum Report 5366, 5 July 1984.
- A.Y. Poludnenko and E.S. Oran, 2010, "The interaction of high-speed turbulence with flames: Global properties and internal flame structure," *Combustion and Flame* **157**(5) 995–1011, 2010.
- A.Y. Poludnenko, 2015, "Pulsating instability and self-acceleration of fast turbulent flames," *Physics of Fluids* **27**, 014106, 2015.
- K.S. Smith, 2007, "Eddy Amplitudes in Baroclinic Turbulence Driven by Nonzonal Mean Flow: Shear Dispersion of Potential Vorticity," *Journal Of Physical Oceanography*: **37**, pp 1037–1050, April 2007.
- A.F. Thompson and W.R. Young, 2007, "Two-Layer Baroclinic Eddy Heat Fluxes: Zonal Flows and Energy Balance," *Journal of the Atmospheric Sciences* **64**, pp.3214–3231, 2007.
- C.A.Z. Towery, A.Y. Poludnenko, J. Urzay, J. O'Brien, M. Ihme, and P. E. Hamlington, 2016, "Spectral energy transfer in turbulent premixed reacting flows," *Physical Review E* **93**,053115 , 23 May 2016.