Progress toward large-eddy simulation of turbulent reacting and non-reacting flows in complex geometries

By K. Mahesh † , G. Constantinescu, S. Apte, G. Iaccarino, F. Ham AND P. Moin

1. Motivation and objectives

The large-eddy simulation (LES) approach is used to simulate the combustor of the gas turbine engine. LES was chosen because of its demonstrated superiority over RANS in predicting mixing, which is central to combustion. The combustor simulations have two major components - gas phase and sprays. The gas phase part of the project has developed a parallel, unstructured grid LES solver which has now been completely integrated with the spray module. As discussed in previous reports (Mahesh *et al.* 2000, 2001) the gas phase solver is non-dissipative and discretely conserves energy, thus insuring both accuracy and robustness for high Reynolds number simulations in complex geometries. A dynamic procedure is used to compute the coefficients in the additional terms that are present in the filtered momentum and scalar-transport equations. The objective of this work is to develop unstructured mesh technology for LES of reacting flow (including spray physics) in realistic configurations using massively-parallel computing platforms.

2. Accomplishments

Our progress in the last year was as follows:

• The gas-phase solver was made fully implicit, resulting in a significant overall speed-up for both non-reacting and reacting flow simulations. Overall reduction in CPU time is between six and ten. Compared to the explicit code, the time step is larger by factors of $10 - 20$ while the increased overhead is $1.5 - 3.5$.

• Development of a geometric multigrid approach to solving the Poisson equation was initiated and preliminary tests were completed.

• The cold-flow simulations in the complex Pratt & Whitney combustor and the front-end validation model that were initiated last year were completed. The LES results are shown to predict experimental data considerably better than RANS.

• Extension of the gas-phase solver to the reacting, variable-density, low-Mach-number equations was completed. The progress-variable approach of Pierce & Moin (2001) is used in the reacting-flow calculations. The LES module was rewritten to reduce memory use. Also, the CPU time per time step of the reacting module was reduced from approximately 25 % of the total to about $10 - 15$ %.

• Validation of the low-Mach-number algorithm used for variable-density flow simulations was initiated. Comparison to the coaxial jet methane combustor of Spadaccini *et al.* (1976) is under way. Preliminary results show good qualitative agreement with the numerical simulations performed for the same conditions by Pierce & Moin (2001).

† Aerospace Engineering and Mechanics, University of Minnesota

• Reacting-flow simulation using JetA fuel in the complex Pratt & Whitney combustor geometry has been initiated. Extensive validation data is available for this test case.

• Validation simulations of particle-laden swirling flows in a coaxial combustor geometry (Sommerfeld & Qiu 1991) were completed, obtaining good agreement with the experimental data.

• A spray-atomization methodology for use with the Lagrangian schemes was developed. The secondary-breakup phenomenon is modeled to predict spray characteristics such as droplet distribution, penetration depth and spray angle.

• A novel hybrid approach involving tracking of individual droplets and parcels of droplets was developed to reduce the computational cost while retaining the essential dynamics of spray evolution.

• Validation of the secondary breakup model was performed and effectiveness of the hybrid scheme was demonstrated in a simplified combustor geometry.

• Validation of the spray-breakup model in Pratt & Whitney's front-end validation geometry was initiated.

• The combustor code was christened CDP in memory of the late Dr. Charles David Pierce who made several lasting contributions to the LES of reacting flows.

3. Implicit algorithm

As described in last year's progress report, the gas-phase solver used the explicit second–order Adams-Bashforth method to advance the equations in time. Algorithmic developments in the gas-phase solver emphasized spatial discretization, which resulted in the development of a non-dissipative, energy-conserving formulation in the absence of time-discretization errors. This explicit algorithm was successfully used to demonstrate the accuracy of the flow solver in both simple and exceedingly complex geometries such as the Pratt & Whitney combustor.

Last year, the gas-phase solver was sped up considerably by allowing the use of larger time steps. This was achieved by making the time advancement fully implicit. The need for implicit time-advancement was felt because numerical stability restrictions imposed by the Adams-Bashforth method were restricting the time step in the coaxial combustor simulations to be an order of magnitude less than the time step used by Pierce & Moin in their structured grid computations that treated the viscous terms implicitly. Also the simulations performed in the front-end model showed that the narrow passages in the fuel injector considerably accelerate the flow, and, as a result, the convective terms impose strict restrictions on the time step for an explicit scheme.

The second-order Crank-Nicolson scheme is used for both convection and viscous terms. The convection terms are linearized prior to solution. The implementation is such that the viscous terms alone can be implicitly advanced if so desired. At present successive over-relaxation is used to solve the implicit system. Typically 20 – 100 iterations are needed to converge the residuals. The use of multigrid techniques to solve the implicit system is under consideration. Results for some typical calculations are summarized in table 1. The savings are seen to be significant. For example, explicit calculation of the cold flow in a coaxial geometry required 320 hours $\times$ 96 processors = 30,700 CPU hours on an IBM SP3 machine. The implicit code uses about 5,000 CPU hours, which is approximately a factor of four larger than the time taken by the highly-optimized structured-grid solver of Pierce & Moin which uses the same time step. The structured solver is of course incapable of handling geometries as complex as the Pratt & Whitney combustor. The

	Grid (10^6 cvs)	Processors	Explicit (CPU hours)	Implicit (CPU hours)
Coaxial combustor (Sommerfeld)	1.6	96	30700	5000
Turbulent channel $Re_\tau = 180$	0.9	32	2240	256
Pratt & Whitney combustor	1.4	32	13500	3200

TABLE 1. CPU time in hours for the the explicit and implicit solvers.

above ratio of implicit to explicit time step also applies to the reacting-flow simulations. Reacting calculations of the Spadaccini experiment using the explicit solver would require about 600,000 CPU hours on an IBM SP3 machine, while the implicit solver requires approximately 70,000 CPU hours. This cost is comparable to that of a structured code (the physical time steps were identical and corresponded to a CFL number between 0.5 and 1.0) which required $\sim$ 50,000 CPU hours on the ASCI RED machine. Further significant reduction of the CPU time in the unstructured LES code is expected, once a multi-level multigrid solver for the pressure and momentum equations in the predictor step of the fractional-step algorithm is implemented.

The calculations in complex geometries show significant speed-up due to the implicit algorithm. The time taken for the cold-flow calculations in the full Pratt & Whitney geometry is very reasonable (3,200 CPU hours, or about 100 wall-clock hours when the job is run on 32 processors). However, computations in the front-end model are still expensive (110,000 CPU hours) in spite of the implicit algorithm speed-up by a factor of six. This is because the time step is now limited by accuracy; it is limited by the high flow speeds through the channels of the three swirlers. In normalized terms, the time step at which the calculations are being currently run are the same as that for the simple coaxial combustor. The high cost for the front-end model is therefore the price paid by unsteady simulations in general (including unsteady RANS) and is not peculiar to LES. We are presently investigating the exact requirements of grid resolution and quality inside the injector region such that the time step could be increased without compromising accuracy.

4. Combustion model implementation in the unstructured code

In this section we present the motivation behind using the flamelet / progress-variable combustion model developed by Pierce & Moin (2001), together with an overview of their method, including the equations and the algorithm to calculate the subgrid momentum and scalar transport terms and its implementation in the unstructured code. Pierce & Moin's approach is based on "quasi-steady" flamelets in which the local flame state undergoes unsteady evolution through a sequence of stationary solutions to the flamelet equations.

A single-parameter flamelet library is first developed for the given combustor conditions

by looking for stationary solutions to the one-dimensional reaction-diffusion equations. The unstable and the lower branches of the S-shaped curve are included so that the complete range of flame states, from completely extinguished (mixing without reaction) to completely reacted (equilibrium chemistry), is represented in the library. Arbitrarily complex chemical-kinetic mechanisms as well as differential-diffusion effects can be included. The result is a complete set of flame states, given in terms of mixture fraction and a single flamelet parameter, denoted by λ:

$$y_k = y_k(Z,\lambda) \ , \quad T = T(Z,\lambda) \ , \quad \rho = \rho(Z,\lambda) \ , \quad w_k = w_k(Z,\lambda) \ . \tag{4.1}$$

where y_k are the mass fractions of the chemical species, T is the temperature, ρ the density, w_k are the reaction terms in the scalar transport equation for the chemical species and Z is the mixture fraction. One of the combustion variables, or some combination of the variables that is representative of the overall flame behavior, is chosen to serve as an overall reaction progress variable.

In this model in addition to the variable-density momentum and continuity equations (4.2) and (4.3), scalar-transport equations are solved for the mixture fraction Z, which is a conserved scalar (4.4), and for the progress variable C, which is a non-conserved scalar (4.5):

$$\frac{\partial \rho \mathbf{u}}{\partial t} + \nabla \cdot (\rho \mathbf{u}\mathbf{u}) = -\nabla p + \nabla \cdot \left[2\mu (\mathbf{S} - \tfrac{1}{3}\mathbf{I}\, \nabla \cdot \mathbf{u}) \right] \ , \tag{4.2}$$

$$\nabla \cdot (\rho \mathbf{u}) = -\frac{\partial \rho}{\partial t} \ , \tag{4.3}$$

$$\frac{\partial \rho Z}{\partial t} + \nabla \cdot (\rho \mathbf{u} Z) = \nabla \cdot (\rho \alpha \nabla Z) \ , \tag{4.4}$$

$$\frac{\partial \rho C}{\partial t} + \nabla \cdot (\rho \mathbf{u} C) = \nabla \cdot (\rho \alpha \nabla C) + \rho w_C \ . \tag{4.5}$$

In the above equations, $\mathbf{S}$ is the strain-rate tensor, $\mathbf{I}$ is the identity tensor, μ is the molecular viscosity, α is the molecular diffusivity and w_C is the chemical-reaction source term.

The continuity equation acts as a constraint on the velocity field, with the time derivative of density as a source term. This constraint is enforced by the pressure, in a manner analogous to the enforcement of the incompressibility constraint for constant density flows.

Under the model assumptions, all the other fluid and flow variables (density, temperature, molecular viscosity, molecular diffusivity), chemical species and the reaction source terms in the scalar-transport equations are related to the mixture fraction and the progress variable through a flamelet library that is precalculated, given a specific fuel reaction mechanism and the flow conditions in the combustor. The only requirements for the quantity chosen to serve as progress variable are that it is representative of the overall gross flame behavior and that it varies monotonically with the flame state so that its value uniquely determines it. For instance, in the reacting-flow simulation of methane-air combustion in a coaxial jet combustor discussed below, the progress variable is chosen as the product mass fraction $C = y_{CO_2} + y_{H_2O}$ (Pierce & Moin 2001).

For turbulent simulations, the governing equations (4.2) to (4.5) are filtered. A major modeling requirement is for the nonlinear density function $\overline{\rho}$ obtained by filtering the

state relation for the density ρ in (4.1). While algebraic scaling laws and scale-similarity concepts can be expected to work for quadratic nonlinearities, the only acceptable closure for arbitrary nonlinearities appears to be the probability density function (PDF) approach. Here, the presumed subgrid PDF is used:

$$\overline{\rho}^{-1} = \int \rho^{-1}(Z,\lambda)\widetilde{P}(Z,\lambda)\,dZ\,d\lambda \tag{4.6}$$

In addition, the filtered progress-variable equation contains a reaction source term $\widetilde{w}_C$ that must be closed. This is accomplished by writing

$$\widetilde{w}_C = \int w_C(Z,\lambda)\widetilde{P}(Z,\lambda)\,dZ\,d\lambda \tag{4.7}$$

and assuming that

$$\widetilde{P}(Z,\lambda) = \delta(\lambda-\lambda_0)\cdot\beta(Z,\widetilde{Z},\widetilde{Z''^2})\ . \tag{4.8}$$

That is, each subgrid state is represented by a single flamelet. For conserved scalars such as mixture fraction, the subgrid PDF is modeled using the beta distribution which is a reasonable assumption in the absence of further information about the subgrid state (Wall *et al.* 2000). The final step is to relate λ_0 to the filtered value of the progress variable, $\widetilde{C}$, that is obtained by solving the corresponding transport equation:

$$\widetilde{C} = \int C(Z,\lambda)\widetilde{P}(Z,\lambda)\,dZ\,d\lambda\ . \tag{4.9}$$

After substitution of the presumed PDF and integrating, this yields

$$\widetilde{C} = f(\widetilde{Z},\widetilde{Z''^2},\lambda_0) \tag{4.10}$$

If C is a monotonic function of λ, then the above relation can be used to eliminate λ_0 from the problem. The final result is a closed specification of the chemical system and fuel properties (molecular viscosity and molecular diffusivity) in terms of three variables $\widetilde{Z}$, $\widetilde{Z''^2}$ and $\widetilde{C}$, which are chosen as the input variables in the chemical table for turbulent calculations:

$$y_k = y_k(\widetilde{Z},\widetilde{Z''^2},\widetilde{C})\ ,\quad T = T(\widetilde{Z},\widetilde{Z''^2},\widetilde{C})\ ,\quad \rho = \rho(\widetilde{Z},\widetilde{Z''^2},\widetilde{C})\ ,\quad w_k = w_k(\widetilde{Z},\widetilde{Z''^2},\widetilde{C})\ . \tag{4.11}$$

$$\mu = \mu(\widetilde{Z},\widetilde{Z''^2},\widetilde{C})\ ,\quad \alpha = \alpha(\widetilde{Z},\widetilde{Z''^2},\widetilde{C})\ . \tag{4.12}$$

The subgrid mixture-fraction variance $\widetilde{Z''^2}$ is obtained using the method proposed by Pierce & Moin:

$$\overline{\rho}\widetilde{Z''^2} = C_Z\overline{\rho}\Delta^2|\nabla\widetilde{Z}|^2\ . \tag{4.13}$$

where the coefficient C_Z is calculated dynamically. The subgrid momentum and scalar transport terms that appear from the filtering of (4.2), (4.4) and (4.5) must be modeled. The eddy viscosity μ_t and eddy diffusivity α_t that appear in these terms are evaluated as follows:

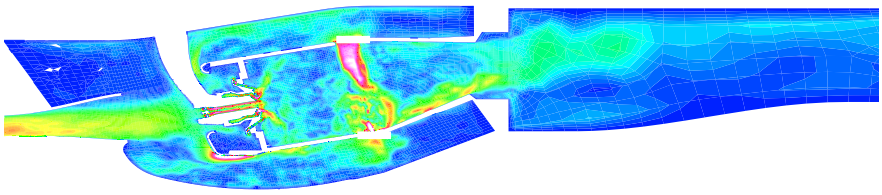

FIGURE 1. Computational domain in the $x-y$ symmetry plane, $z=0$, with contours of the instantaneous absolute value of the velocity field for the flow in the Pratt & Whitney full-combustor geometry in the same plane

$$\mu_t = C_\mu \overline{\rho} \Delta^2 |\widetilde{\mathbf{S}}| \tag{4.14}$$

$$\rho \alpha_t = C_\alpha \overline{\rho} \Delta^2 |\widetilde{\mathbf{S}}| \tag{4.15}$$

where the coefficients C_μ and C_α are calculated dynamically.

All the operators in the scalar-transport equations are discretized in a similar fashion to that used in the momentum equations.

Finally, we describe the approach for the calculation of the coefficients in the expressions of the subgrid scalar variance, eddy viscosity and eddy diffusivity in the filtered momentum and scalar transport equations. It is well known that the prediction of negative values for the dynamically-calculated coefficients (C_μ, C_α and C_Z) coupled with the long time correlation of these coefficients can cause instability of the numerical solution. Negative values of these coefficients can predicted by the dynamic procedure, especially in regions of high gradients and skewed meshes, two conditions that are often met in calculations of very complex turbulent flows using hybrid unstructured meshes. The usual way to avoid this problem is to clip these coefficients, so that the sum of the molecular and eddy viscosity (diffusivity) will remain positive. However, in our implementation we propose a more robust implementation, based on Germano's identity, in which the usual expression is used to compute the coefficients in the region where the predicted values for these coefficients are positive (e.g. $C_\mu \Delta^2 = 0.5 < L_{i,j} M_{i,j} > / < M_{k,l} M_{k,l} >$), while in regions where the least square is negative, by contracting Germano's identity with itself one obtains $C_\mu \Delta^2 = 0.5 \sqrt{< L_{i,j} L_{i,j} > / < M_{k,l} M_{k,l} >}$. The other advantage of using this approach is that it eliminates the need for averaging the terms of the form $< L_{i,j} M_{i,j} >$, an operation that is generally needed to improve robustness. In an unstructured environment this operation is not trivial, even for problems in which the flow is homogeneous in one or two directions. Moreover, we are primarily interested in calculations of complex flows in which there are no homogeneous directions. Though several methods like Lagrangian averaging (Meneveau *et al.* 1996) were proposed to address this problem, the typical solution is to apply a locally-defined filter on the field of the computed values C_μ. Using the present approach to compute the coefficients in the dynamic procedure, we found that this step was not necessary.

Location	LES Error % wrt expt.	LES Error % wrt inlet	RANS Error % wrt expt.	RANS Error % wrt inlet
OD dilution hole	3.1	0.8	11.4	3.2
ID dilution hole	3.5	0.5	7.5	1.1
Core (main swirler)	10.3	0.14	8.4	0.11
Second (OD) swirler	7.5	0.35	13.5	0.63
Third (Guide) swirler	0.4	0.02	18.9	0.84

TABLE 2. Comparison of mass-flow splits in the Pratt & Whitney full-combustor geometry

5. Cold-flow simulations in the Pratt & Whitney geometry

5.1. *Complex combustor*

Last year's report documented flow-visualization results from cold-flow calculations in the complex Pratt & Whitney combustor geometry. These computations were extended to include the effects of flow bleed and transpiration and detailed comparisons with experiment were made for mass-flow splits and pressure drop. The geometry in the symmetry plane $z = 0$ together with contours of the absolute value of the velocity field are shown in figure 1. The Reynolds number in the pre-diffuser inlet section, defined with the bulk velocity in the inlet section and a characteristic length scale $L = 1''$, is about $500,000$ while it takes a value of about $150,000$ in the main (core) swirler channel. Turbulent fluctuations from a separate calculation, in a periodic pipe sector of shape identical to the pre-diffuser inlet section, are fed in at the inlet. In the experiment, fuel is injected at the tip of the downstream part of the injector. The mass-flow rate for the fuel is matched in our simulation.

Interestingly, the RANS and LES predictions for the pressure drop across the different components of the combustor are very close (within $6,000$ Pa for 5 out of the 6 stations). However, as shown in table 2, LES does a better job overall in predicting the mass-flow splits through the swirlers and the inner and outer dilution holes. LES predictions are within 10.5% of the experimental measured values for the mass splits through the swirlers, but more importantly within 4% for the total discharges through the swirlers, inner and outer dilution holes. The errors as a percentage of the total inlet discharge are much lower.

5.2. *Front-end model*

Cold-flow simulations were performed in the Pratt & Whitney front-end test-rig geometry. This geometry has the same fuel injector and combustion chamber as the complete combustor, but air is fed to the injector through a cylindrical plenum, the inner and outer diffuser channels are absent, and the main combustor chamber does not have dilution holes on the surrounding walls. The interest in simulating the flow in this geometry is the availability of detailed LDV velocity-profile measurements in the main combustion

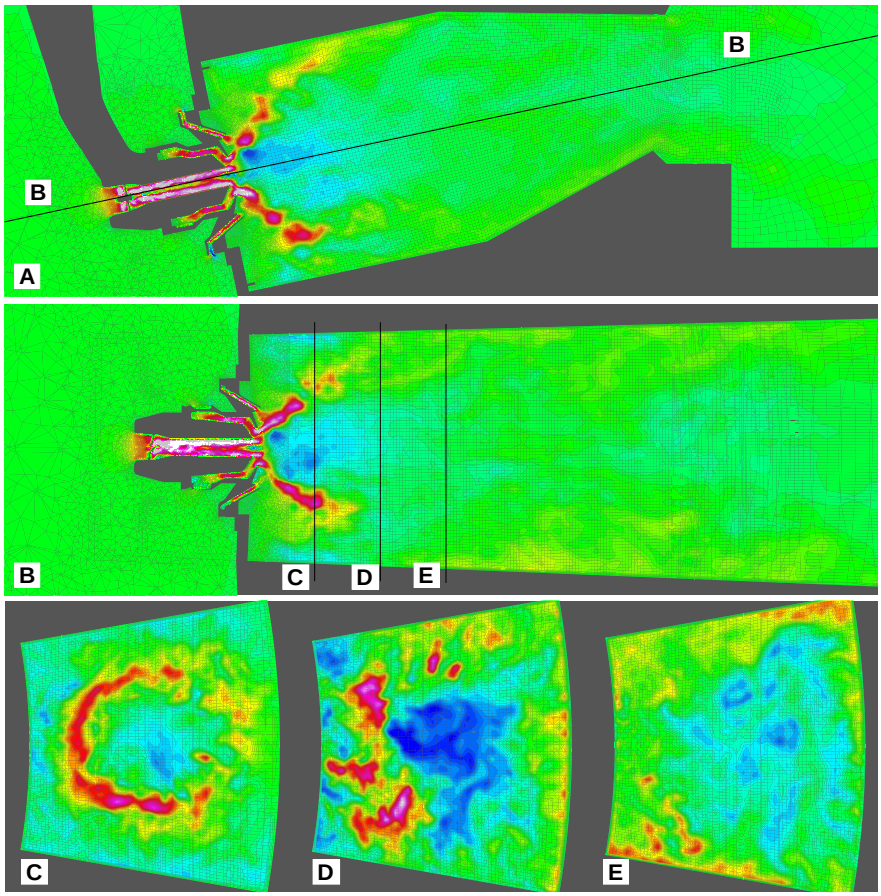

FIGURE 2. Fine-grid LES solution for the front-end rig geometry; Contours of the instantaneous streamwise velocity component (a) in the $x-y$ symmetry plane (z=0), (b) in a plane containing the main swirler symmetry axis, perpendicular on the $z = 0$ plane, and at several downstream locations (c) $x = 1.1''$, (d) $x = 2.2''$, (e) $x = 3.4''$

chamber, which can be used to fully validate the accuracy of our solver for a geometry of complexity comparable to the full combustor. The Reynolds number in the main swirler channel of the injector is close to $100,000$.

The complexity of the geometry and flow inside the test rig is illustrated in figure 2 which shows some sections of the computational domain along with contours of the streamwise velocity in the symmetry plane (z=0), in a plane containing the main swirler symmetry axis, perpendicular to the $z = 0$ plane, and at several downstream locations, $x = 1.1''$, $x = 2.2''$ and $x = 3.4''$ from the injector, for the statistically-steady solution obtained on a fine mesh. The main feature observed in these plots is the formation of a relatively large recirculation region downstream of the injector, due to the swirling jet coming out from the injector into the main combustor chamber.

Prediction of the correct dimensions of the recirculation region, together with the variation of the jet width with distance from the injector, are two of the main challenges in simulating this flow. In particular, RANS calculations of this flow, conducted both at Pratt & Whitney using an in-house $\kappa - \epsilon$ code and at Stanford using a commercial

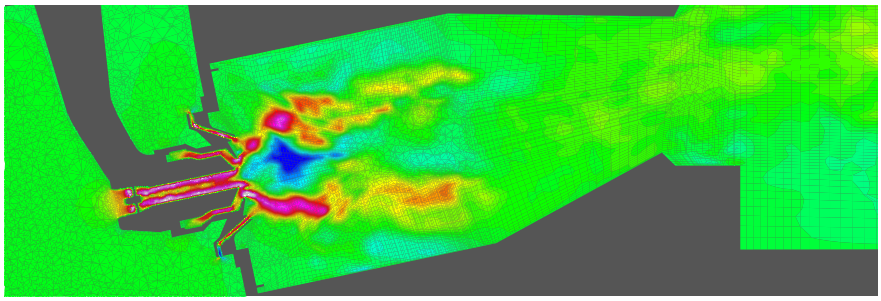

FIGURE 3. Coarse-grid LES solution; Contours of the instantaneous streamwise velocity component in the $x-y$ symmetry plane ($z=0$)

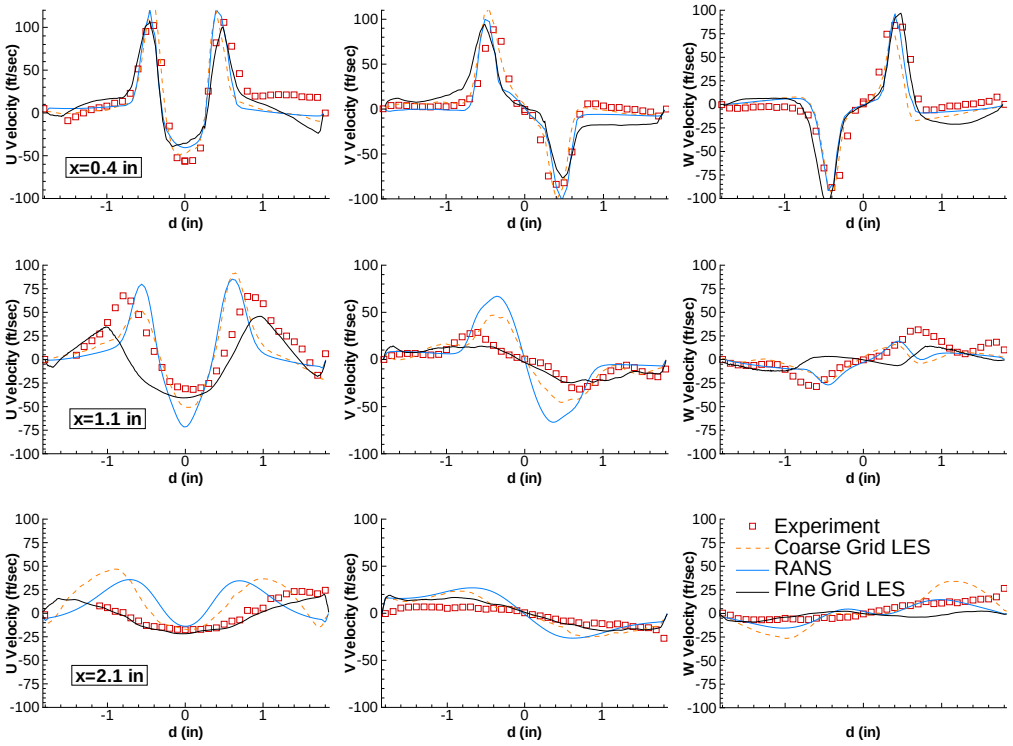

FIGURE 4. Comparison between fine-grid LES, coarse-grid LES, RANS (FLUENT) and experiment (Pratt & Whitney) for the gas-phase velocity (streamwise, radial and swirling components) in the Pratt & Whitney combustor at three stations situated at $x=0.4''$, $x=1.1''$ and $x=2.1''$ downstream of the injector

software (FLUENT), failed to predict these quantities correctly away from the injector, though they were fairly successful in predicting all three velocity components immediately downstream of the injector. This is evident from comparison of the RANS profiles with the experimental data symbols in Fig. 4. Interestingly, the Pratt & Whitney RANS simulation was able to capture the pressure drop across the injector within 2%, while FLUENT yielded a much bigger error ($\sim$ 20%). This is attributed to different grid densities inside the injector.

Two grids were generated in the present LES simulations: a coarser grid containing about 2.2 million control volumes and a finer containing about 4.5 million control volumes. Contours of the instantaneous streamwise velocity component are shown in figures 3 and 2a for the solutions obtained on the two grids at a statistically steady state. Though for the simulation on the coarser mesh the prediction of the mean pressure drop across the injector (4588 Pa) was found to be very close to the experimental value (4500 Pa), the agreement with the experimental data for the velocity profiles was not much better than the results obtained from RANS calculations. In particular, the spreading of the jet away from the injector was substantially underpredicted. This can be seen by comparing the streamwise velocity profiles at the downstream station x=2.1" from the injector in figure 4. The reason for this is that insufficient grid resolution causes the (conical) detached shear layers, which are initially shed at the correct angle from the injector (this is expected because all models do a fairly good job in predicting all the velocity components immediately downstream of the injector), to curve toward the injector centerline, due to a too-high rate of decay of azimuthal momentum inside the initial region of these layers. This results in a much smaller recirculation region compared to the fine grid solution (see figure 2) where the detached shear layers are seen to extend up to the lateral walls of the main combustor chamber. The angle between the injector axis and the conical shape corresponding to the detached layers is approximatively constant at 55°. The shedding of vortex tubes due to the Kelvin-Helmholtz instabilities in the detached shear layers is clearly observed. The fine-grid solution in figure 2 displays the right features corresponding to the flow in the test-rig geometry at the specified conditions. This also results in better quantitative predictions for the velocity profiles (see figure 4), especially away from the injector where, as clearly observed in the mean streamwise velocity profile at $x = 2.1''$ from the injector, the level of agreement between the LES fine-grid solution and the experiment is clearly superior to the one observed for the RANS or the LES coarse-grid solutions. However, some differences between the experiment and the fine-grid LES results remain (e.g. compare the streamwise velocity profiles at the $x = 1.1''$ station). These differences may occur because the flow is under-resolved in some regions. We plan to address this by performing one additional calculation on a very fine mesh (14 million cells).

6. Multigrid solution of the pressure-Poisson equation

In the present unstructured LES solver, the conjugate-gradient (cg) solution of the pressure-Poisson system is the most computationally-expensive component of the overall solution process, requiring from 50 to 80% of the total solution time. In this section we present some details of a multigrid solver for this Poisson system that will eventually replace the present cg solver. Tests using a single coarse grid (1-level multigrid) have demonstrated an overall reduction in computation time of 35 to 60% per time step. Once fully implemented, the pressure solver should become one of the least expensive components of the solution process, yielding reductions in overall computation time of 45 to 75% per time step.

6.1. *Background*

Multigrid methods are used extensively in the numerical solution of partial differential equations. They can exhibit an optimal complexity in terms of both work and storage — i.e. both work and storage scale linearly with problem size. They can also achieve very good parallel efficiency and scalability by the method of domain decomposition. For

standard multigrid methods involving point-relaxation smoothers and isotropic coarsening, however, the convergence factor is known to degenerate dramatically in the presence of coefficient anisotropy. Coefficient anisotropy can result from large cell aspect ratios, anisotropic material properties, or asymmetric operators. In the present case the anisotropy is primarily geometric, caused by the grid stretching used to resolve high-Reynolds number boundary layers.

In a effort to make multigrid methods more robust, several solutions have been proposed to this problem. Most of these solutions involve some combination of the following two ideas: 1) the use of semi-coarsening, where the coarsening is not isotropic, and 2) improvements to the smoother so that both high- and low-frequency components of the residual distribution are effectively reduced in at least one or two directions (i.e. the so-called line and plane smoothers of structured-grid methods). When the semi-coarsening and/or smoother improvements are properly matched, multigrid methods recover their optimality, even in the presence of strong coefficient anisotropies.

6.2. *Proposed multigrid method*

The multigrid method currently being integrated into the unstructured LES solver has the following characteristics:

- V-cycle multigrid with linear restriction and block correction.
- Multi-color Gauss-Seidel smoothing on all grids except the coarsest; on the coarsest grid, the residual is reduced by approximately one order of magnitude using a coarse grid version of the present cg solver.
- Coarse-grid control volumes are built by the agglomeration (sometimes called aggregation) of fine grid control volumes. The agglomeration procedure is performed once per computation, when the pressure solver is first called. Agglomeration avoids use of mesh generators to generate the coarse grids, and simplifies the restriction and prolongation operators.
- Directional agglomeration (semi-coarsening) is used to prevent stalling of the multigrid method in the presence of large cell aspect ratios and the resulting coefficient anisotropy. At present, a greedy-type algorithm is used, similar to that described by Raw (1996) and Elias *et al.* (1997).

Other authors have solved the coefficient-anisotropy problem by using uniform coarsening, combined with line-implicit smoothing in the direction of greatest coefficient strength (Brandt *et al.* 2002, Mavriplis & Pirzadeh 1999). While effective for certain problems, line-implicit smoothing combined with uniform coarsening will not yield an optimal multigrid method in general (Montero *et al.* 2001).

- Coarse-grid coefficients are calculated using the discretized coarse-grid approximation (i.e. geometric multigrid). Algebraic multigrid with first-order restriction (insertion) of the fine-grid coefficients was tested and found to yield significantly slower convergence rates. Algebraic multigrid with higher-order restriction was not considered because of the increase in bandwidth (neighbor connections) associated with coarser grids.

6.3. *Results*

A two-grid version of this multigrid method has been added to the parallel unstructured LES solver and tested on a number of problems, including the coaxial combustor (1.1 million control volumes, 32 processors). Figure 5 compares the normalized wall-clock time per time step for this problem, breaking out the pressure solver from the other components of the solution (scalar equations, momentum equations, chemistry, etc.). For

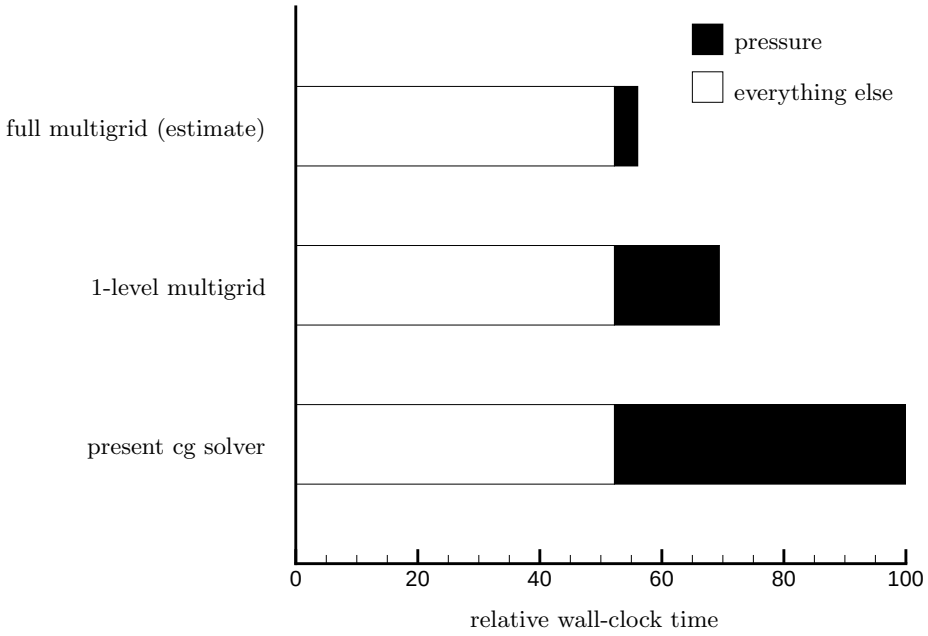

FIGURE 5. Normalized breakdown of total computation time for three different Poisson solvers: coaxial combustor problem with 1.1 million control volumes on 32 processors.

the coaxial-combustor problem, the pressure solution requires about 50% of the total solution time when the present cg solver is used.

The one-level multigrid solver is able to converge the pressure equation to the same tolerance in about 10 to 12 cycles: however each cycle is significantly more expensive than a fine-grid cg iteration, so the actual speed-up in pressure solution is about 2.7 times, corresponding to a reduction in overall solution time of about 30%. Most of the computational cost in this one-level multigrid is associated with the coarse-grid cg iterations, where 100 or more iterations may be necessary to reduce the coarse-grid residual by one order of magnitude. Extrapolating this demonstrated two-grid performance to multiple coarse grids, the speed-up in the pressure solution should be about 10 times, corresponding to a reduction in overall solution time of about 45% for this problem.

Figure 6 compares the convergence history of the the pressure solver for these three cases. The histories are plotted relative to "equivalent cg iterations", proportional to wall clock time.

7. Towards generating very large unstructured grids

In the framework of ASCI, the generation of unstructured meshes is carried out using the commercial software GAMBIT and the research code CUBIT (developed by Sandia National Laboratories); the former is usually preferred because it provides a more user-friendly interface. Grids ranging from one to five million elements have been produced for the geometries discussed above. At present, we simulate a 20° sector of the Pratt &

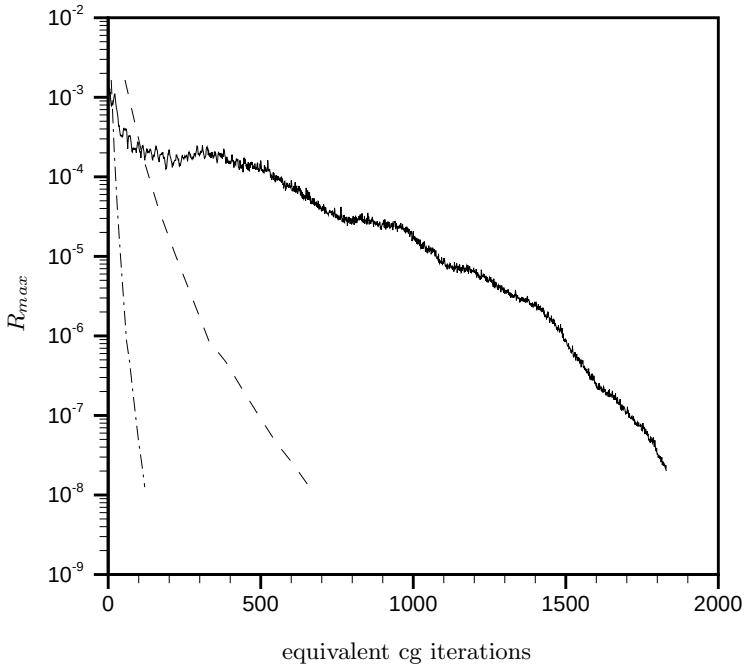

FIGURE 6. Comparison of pressure convergence history for three different Poisson solvers: ——— present cg solver; - - - - 1-level multigrid; —— — full multigrid (estimate); coaxial combustor problem with 1.1 million control volumes on 32 processors.

Whitney combustor. Future simulations are planned that incorporate the entire combustor. In order to provide sufficient resolution for a geometry as complex as the Pratt & Whitney combustor, substantially larger meshes have to be generated. Grid generation is currently carried out on SGI workstations, and the largest grid that can be handled (because of the memory limit of 8 GB on these machines) is about 8 millions elements. There are two approaches that can be followed to overcome this limitation. The first is to generate a coarse grid and then subdivide all the elements along meridian planes: from a hexahedral eight smaller homothetical hexahedrals can be constructed splitting all the faces in four (halving the edges); similarly from one tetrahedra four hometical tetrahedrals can be generated (see figure 7).

This procedure can be easily performed on a multiprocessor machine. A subset of elements is assigned to each processor and the splitting is carried out independently. The major drawback of this process is inaccuracy in the representation of the physical boundaries of the domain. Notice that the splitting of an element requires the generation of an additional node on each face and on each edge. Consider a tetrahedral grid in a spherical volume; the boundary (spherical surface) is approximated by a tessellation of (planar) triangles. The accuracy of this representation can be evaluated by considering the distance between the circumcenter of each triangle and the spherical surface; finer grids correspond to higher accuracy. If the splitting procedure discussed above is used, refined grids will have the same accuracy (in representing the spherical surface) as the initial coarse mesh. The only way to improve the accuracy would be to project the nodes

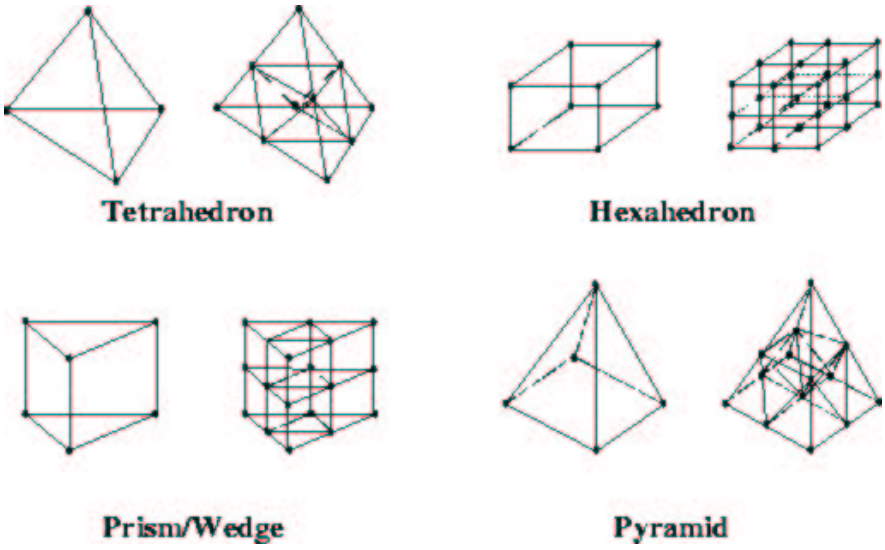

FIGURE 7. Grid generation by element splitting

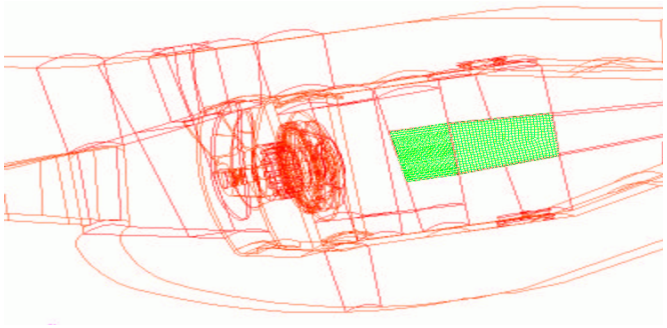

FIGURE 8. Subdomain-based grid generation

created by the splitting on the "real" boundary surface. The result is that to perform this procedure it is necessary to access the CAD representation of the computational domain and carry out projection operations for all the elements belonging to the boundaries. Therefore it is required to access to the CAD modules in GAMBIT with its related memory limitations. Note that neglecting this final boundary-fitting step produced highly distorted geometry representations in regions of high surface curvature.

The second approach consists in dividing the computational domain in subdomains (blocks); the surfaces between blocks are meshed first, and then fine grids (up to the desired resolution) are constructed independently (see figure 8). A *simple* (and intrinsically parallel) code can be written to merge together all the block grids without using GAMBIT and therefore without memory limitations. Initial tests of this second approach are encouraging.

8. Reacting flow simulations in a coaxial combustor geometry

The flamelet / progress-variable model described in a previous section was tested against experimental data for a methane-fueled coaxial-jet combustor corresponding to

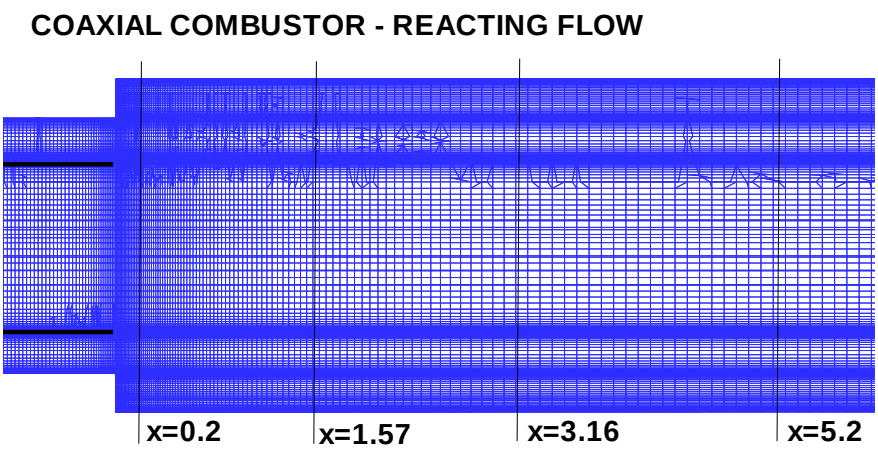

FIGURE 9. Computational domain and mesh section in an azimuthal plane

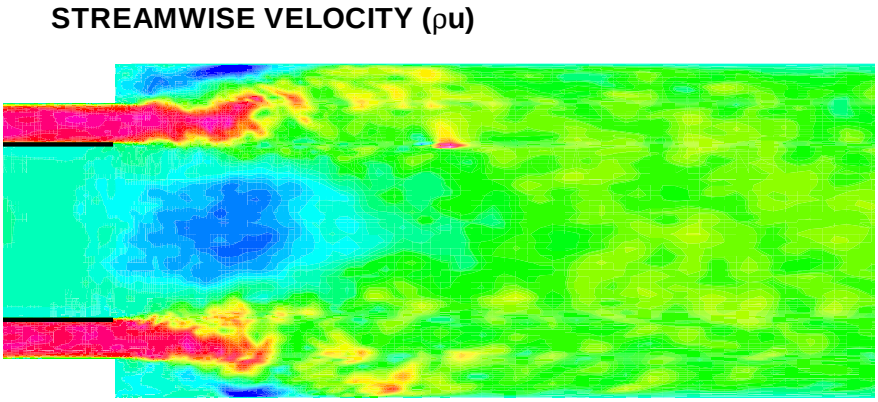

FIGURE 10. Instantaneous contours of the streamwise velocity component in an azimuthal plane

the experiment of Spadaccini *et al.* (1976). This experiment was chosen because of the relatively simple geometry, well-defined boundary conditions (confined jet simulation) coupled with a rather complex flow physics that mimics fairly well the complex-flow phenomena encountered in a realistic gas-turbine combustor. Results from a previous calculation using a structured code by Pierce & Moin (2001) are available, together with the original data of Spadaccini *et al.* (1976). This will allow a full validation of the reacting flow, including mean-velocity profiles, temperature and chemical species at several stations inside the main combustion chamber. A general view of the geometry and computational mesh in an azimuthal plane is given in figure 9. The fuel is introduced through a circular core section at a relatively low velocity, 0.93 m/s, while non-swirling air is introduced through a surrounding annulus at a much higher mean velocity of U=20.6 m/s. The mass flow rates of the fuel and air in the inlet sections are 0.0072 kg/s and 0.137 kg/s, respectively, the air temperature is 750 K, while the fuel temperature is 300 K. The mean pressure in the combustor is 3.8 atm. The fuel used in the experiment was natural

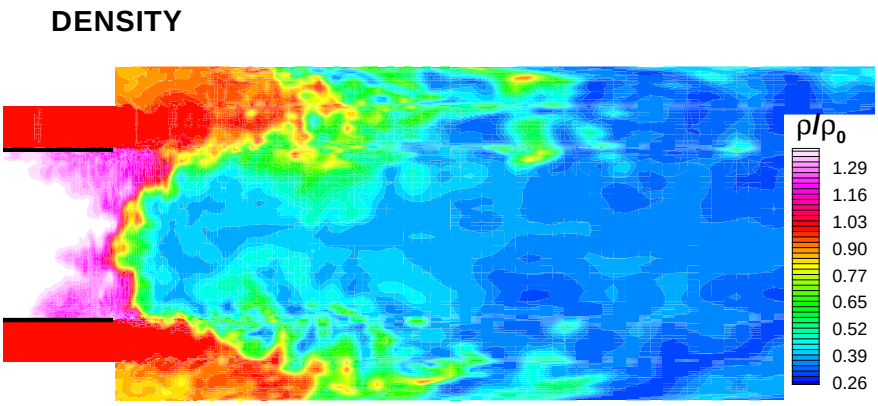

FIGURE 11. Instantaneous contours of the density in an azimuthal plane

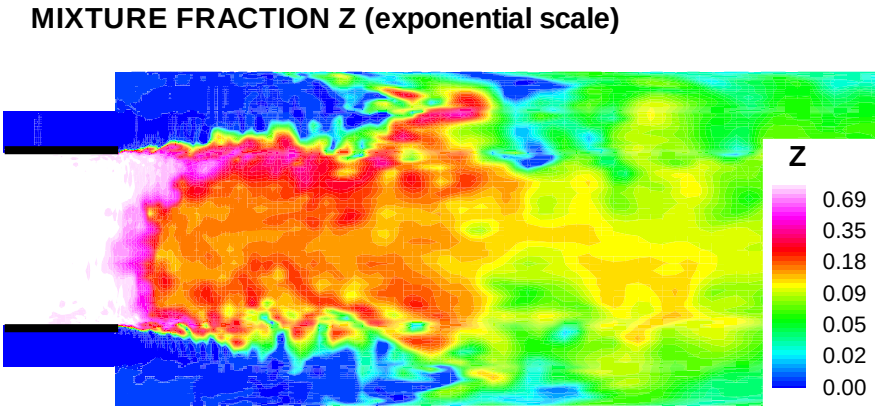

FIGURE 12. Instantaneous contours of the mixture fraction in an azimuthal plane

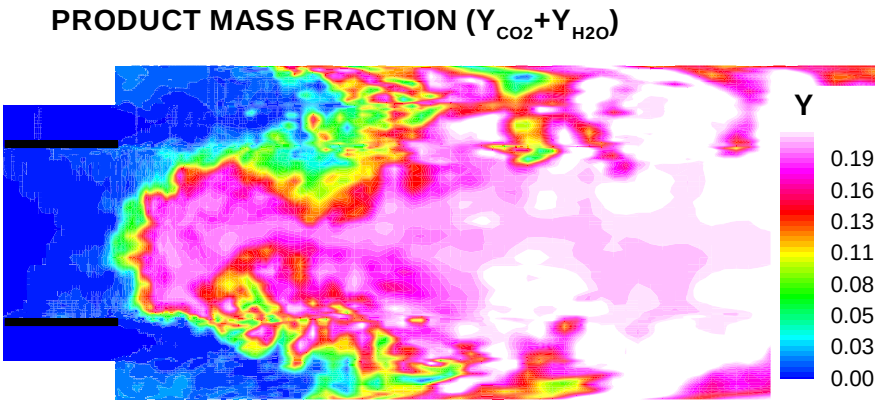

FIGURE 13. Instantaneous contours of the progress variable $y_{CO_2} + y_{H_2O}$ in an azimuthal plane

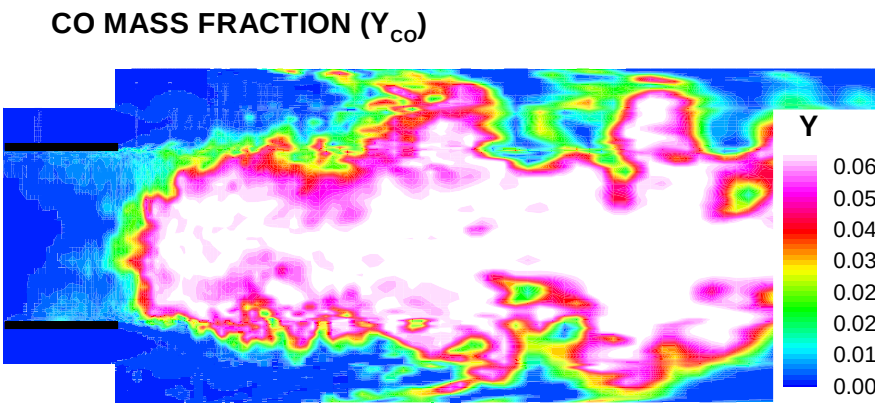

FIGURE 14. Instantaneous contours of the carbon monoxide concentration y_{CO} in an azimuthal plane

gas but in the present investigation, as well as in the numerical simulation of Pierce & Moin, it was assumed to be pure methane. Because of the high air/fuel velocity ratio, a strong central recirculation region is formed in front of the fuel port, which appears to the surrounding air stream almost as a bluff body (see also figure 10 in which contours of the streamwise velocity component are shown). The recirculating combustion products provide a continuous ignition source for the relatively cold incoming reactants, thereby stabilizing the flame. Mixture fraction in the experiment was computed based on the total carbon and hydrogen atoms mass fractions, and the product mass fraction was computed from $C = y_{CO_2} + y_{H_2O}$.

A computational grid containing $\sim$ 1.3 million hexahedra was generated to mesh the geometry that was used by Pierce & Moin in their simulations. Though they used approximately $\sim$ 2.5 million points in their simulations, we were able to achieve roughly the same grid density in the important flow regions (mainly the upstream part of the main combustor) where turbulent mixing and combustion-related phenomena are important. The LES model is very similar in the two simulations, except the specific treatment in the estimation of the model coefficients due to the unstructured environment of our solver. The injector radius corresponding to the annular exterior radius ($R = 4.685$D cm) and the mean air velocity in the inlet section were used for non-dimensionalization. The Reynolds number on the air side is $\sim 50,000$. The computational domain started at a distance of 1R upstream of the main combustor chamber, where fully-developed turbulent inflow conditions were specified using a precalculated database from a separate calculation for a periodic pipe and annular pipe domains at the corresponding Reynolds numbers. The computational domain was extended to a combustor length of 12R, where convective outflow boundary conditions were used to convect the turbulent eddies out of the domain. Fully-developed incompressible velocity and mixing fields were obtained before the chemistry model was turned on. The initial progress-variable scalar field was set to its maximum allowed value determined from fast chemistry, so that initially the flame was fully ignited. A chemical table corresponding to methane and the specified combustor pressure was precalculated. The product mass fraction was chosen as the progress variable. The boundary conditions for the mixture fraction were $Z = 1$ in the fuel inlet section and $Z = 0$ in the air inlet section. The progress variable C in both the fuel and air

inlet sections was set to zero, as no reaction takes place at those locations. A time step of 0.005 R/U was used in the simulation using the implicit algorithm. About 80 inner iterations were needed to reduce the momentum residuals by five orders of magnitudes, while the conjugate-gradient pressure solver (no multigrid) needed about 1,700 iterations per time step to converge the pressure. About 50% of the CPU time per iteration was spent in the pressure solver, and $\sim 15\%$ in the combustion module. Based on Pierce & Moin's estimations we expect that ~ 500 R/U time units will be needed to obtain fully converged statistics. As mentioned in a previous section we expect to use $\sim 70{,}000$ CPU hours for this computation on an IBM SP3 machine.

At this point the calculation was run for ~ 120 R/U time units and the flow seems 'statistically steady' in the upstream part of the main combustor chamber. The streamwise velocity profile and its rms fluctuation in a section situated just downstream of the entrance into the main combustor chamber show good agreement with the similar profiles computed by Pierce & Moin. The statistics are still evolving downstream in the combustor. The flame location can be inferred from the temperature or the progress-variable contours shown in figure 13, as the progress variable essentially is tracking the reaction. In the experiment, the flame location was observed to lift off the burner and reattach intermittently in a highly-unsteady asymmetric manner. Our calculations also predict a lifted flame, and an animation of the progress-variable contours shows that our simulation is able to capture this behavior. At this point good qualitative agreement is obtained for the distribution of the mixture fraction and progress variable with the calculation of Pierce & Moin. Instantaneous contours of these variables in an azimuthal plane are shown in figure 12 and 13. Also shown are similar snapshots of the density (figure 11) and carbon monoxide concentrations (figure 14).

9. Integration with spray module

The spray models described in our progress reports from 1999, 2000 and 2001 have been fully integrated with the unstructured gas-phase solver. The droplets are modeled as point particles which satisfy Lagrangian equations. They influence the gas phase through source terms in the gas-phase equations. As the particles move, their position is located and each particle is assigned to a control volume of the gas-phase grid. The gas-phase properties are interpolated to the particle location and the particle equations are advanced. The particles are then relocated, particles that cross interprocessor boundaries are duly transferred, source terms in the gas-phase equation are computed, and the computation is further advanced. Validation of the Lagrangian particle-tracking approach, application of spray simulations to the Pratt & Whitney injector, and extension of the spray module to account for evaporation are described below.

9.1. *Validation of Lagrangian Particle Tracking (LPT) Scheme in CDP*

The flow in a swirl-stabilized coaxial combustor represents an important validation case. Sommerfeld & Qiu (1991) provide detailed measurements of this flow, which tests both the gas-phase solver and the spray module. A detailed validation of the gas-phase solver in this geometry was completed last year and was presented in an earlier ASCI report. The particle-laden flowfield was computed by injecting particles with a known size distribution at the inlet of the combustor. The primary jet of the combustor is laden with glass beads whose diameter varies from 10 microns to 120 microns. The particle phase was initialized by computing over a few flow-through times. About 1.2 million particles were obtained in the combustor. The particle statistics were then computed over two flow through times

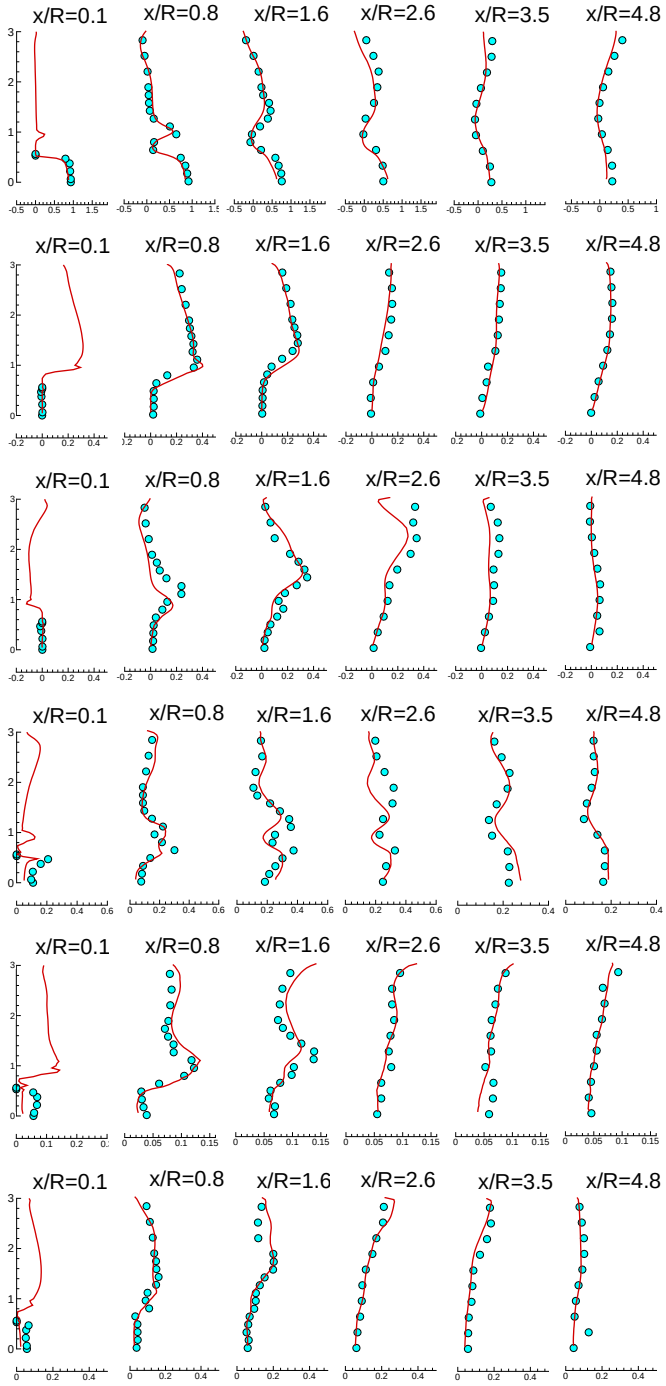

FIGURE 15. Comparison of particle phase properties between LES (———) and experiment (○ , Sommerfeld & Qiu 1991) for swirling flow in a coaxial combustor. (*a*): mean axial velocity, (*b*): mean swirl velocity, (*c*): mean radial velocity, (*d*) rms of axial velocity, (*e*) rms of swirl velocity, (*f*) rms of radial velocity.

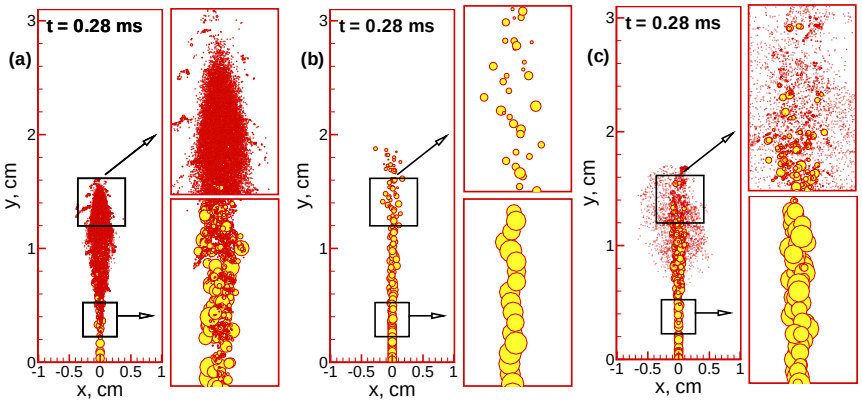

FIGURE 16. Comparison of the pure-droplet, pure-parcel, and hybrid approaches.

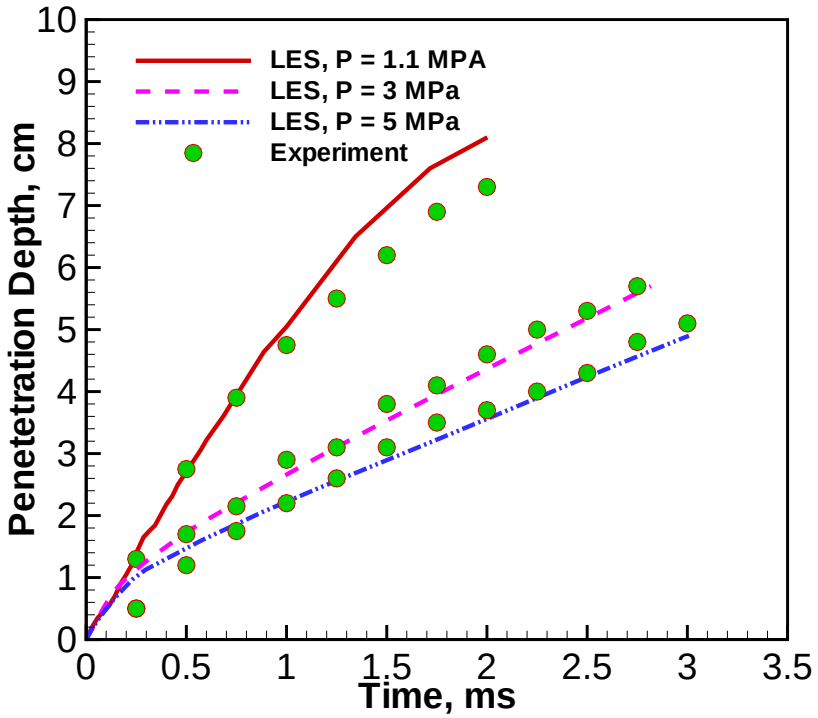

FIGURE 17. Comparison of the spray-penetration depth with experimental data.

in the region of interest. The agreement of mean and rms velocity components with the experiment is shown in figure 15. Very good agreement was also obtained for the mean and rms of diameter distribution at various axial locations. The present computation was performed on 96 processors using an explicit time-advancement procedure. The overhead due to the particle-phase was about 30%. Particle-dispersion characteristics were also obtained by tracking a large number of particles of different size classes. It was compared with the experimental data and analytical estimates to show good agreement. Details of this computation can be found in Apte *et al.* (2002). It was shown that the point-particle

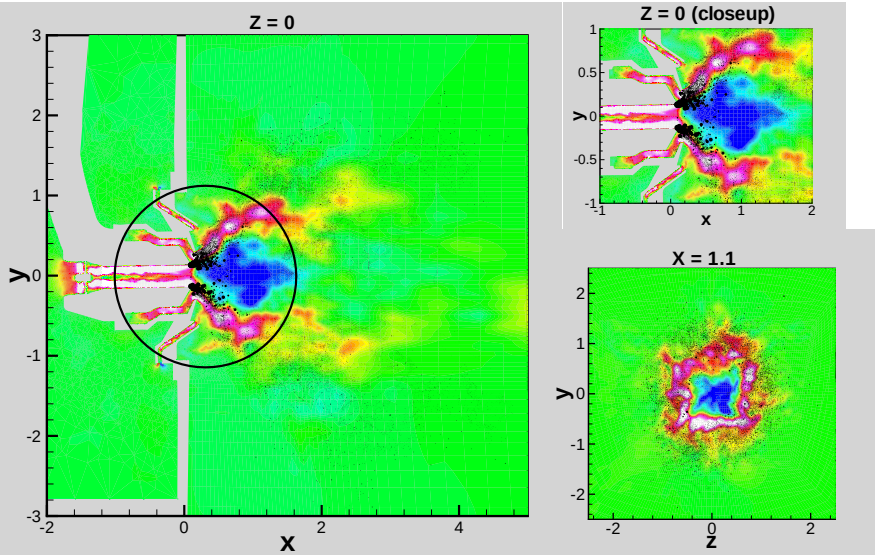

FIGURE 18. Evolution of spray from a PW injector: contours of axial velocity superimposed on particle scatter plot.

approximation of the dispersed phase was able to accurately predict the behavior of the solid phase for the particle loadings simulated. This study represents the first LES of two-phase flows in a coaxial combustor using the unstructured grid solver (CDP). After completing this comprehensive validation study, we focused on incorporating complex spray breakup and evaporation models in CDP.

10. Spray models for use in CDP

10.1. *Stochastic model for secondary breakup*

Liquid-spray atomization plays a crucial role in the combustion dynamics in gas-turbine combustors. In standard Lagrangian particle-tracking codes, emphasis is placed on obtaining the correct spray-evolution characteristics away from the injector. Only the global behavior of the primary atomization, occurring close to the injector, is considered and the details are not captured. The essential features of spray evolution, viz. droplet size distribution, spray angle, and penetration depth, are predicted away from the injector surface by secondary breakup models. Usually standard, deterministic breakup models based on Taylor Analogy Breakup (TAB) (O'Rourke & Amsden 1987) or wave (Reitz 1987) models are employed in RANS-type computations. Liquid 'blobs' with the size of the injector diameter are introduced into the combustion chamber and undergo atomization according to the balance between aerodynamic and surface-tension forces acting on the liquid phase. In the TAB model, oscillations of the parent droplet are modeled in the framework of a spring-mass system and breakup occurs when the oscillations exceed a critical value. In the wave model, new droplets are formed based on the growth rate of the fastest wave instability on the surface of the parent blob. Both models are deterministic, with 'single-scale' production of new droplets. In many combustion applications, however, injection of a liquid jet takes place at high relative velocity between the two phases (high initial Weber number). Under these conditions, intriguing processes

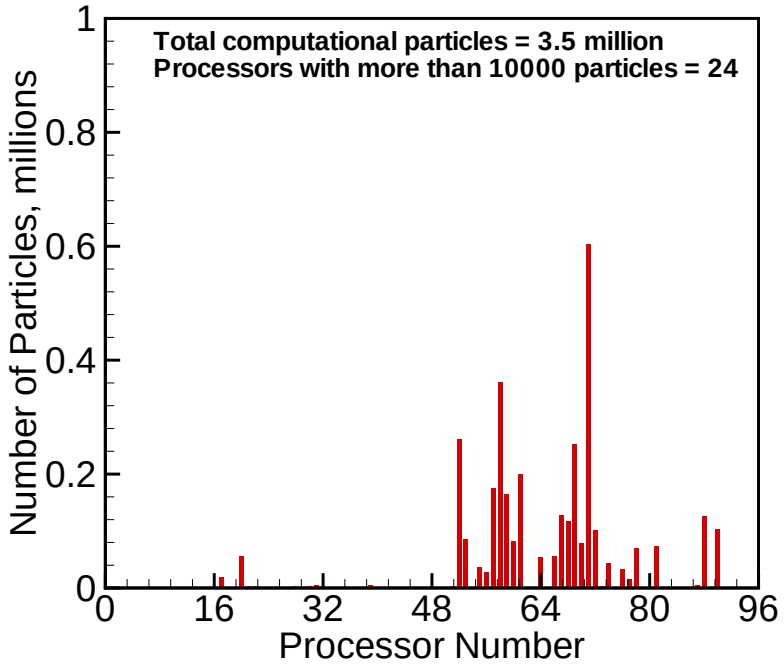

FIGURE 19. Load balance for spray-breakup simulation.

such as turbulence-induced breakup, multiple-droplet collisions in the dense-spray region, fluctuations due to cavitating flow inside the injector, etc., contribute to the process of atomization. This results in droplet formation over a large spectrum of droplet sizes and is not captured by the above models. Predicting the distribution of droplet sizes at each spray location is important for sheet-breakup modeling.

In order to predict the essential global features of these complex phenomena, a stochastic approach for droplet breakup which accounts for a range of product-droplet sizes is necessary. Specifically, for a given control volume, the characteristic radius of droplets is assumed to be a time-dependent stochastic variable with a given initial distribution function. The breakup of parent blobs into secondary droplets is viewed as the temporal and spatial evolution of this distribution function around the parent-droplet size. This distribution function follows a certain long-time behavior, which is characterized by the dominant mechanism of breakup. The size of new droplets is then sampled from the distribution function evaluated at a typical breakup time scale of the parent drop.

Owing to the complexity of the phenomenon, it is difficult to clearly identify such a dominant mechanism for breakup. Kolmogorov (1941) developed a stochastic theory for breakup of solid particles by modeling it as a discrete random process. He assumed that the probability to break each parent particle into a certain number of parts is independent of the parent-particle size. Using the central limit theorem, Kolmogorov pointed out that such a general assumption leads to a log-normal distribution of particle size in the long-time limit.

Based on Kolmogorov's hypothesis we have developed a numerical scheme for atom-

ization of liquid spray at large Weber number (Gorokhovski & Apte 2001). The discrete model of Kolmogorov is reformulated in terms of a Fokker-Planck (FP) differential equation for the evolution of the size-distribution function from a parent blob towards the log-normal law:

$$\frac{\partial T(x,t)}{\partial t} + \nu(\xi)\frac{\partial T(x,t)}{\partial x} = \frac{1}{2}\nu(\xi^2)\frac{\partial^2 T(x,t)}{\partial x^2} \ . \tag{10.1}$$

where the breakup frequency (ν) and time (t) are introduced. Here, $T(x,t)$ is the distribution function for $x = \log(r_j)$, where r_j is the droplet radius. Breakup occurs when $t > t_{breakup} = 1/\nu$. The value of the breakup frequency and the critical radius of breakup are obtained by the balance between the aerodynamic and surface tension forces. The secondary droplets are sampled from the analytical solution of equation (10.1) corresponding to the breakup time-scale. The parameters encountered in the FP equation ($\langle\xi\rangle$ and $\langle\xi^2\rangle$) are computed by relating them to the local Weber and Reynolds numbers for the parent blob, thereby accounting for the capillary forces and turbulent properties. The capillary force prescribes a lower bound limit for the produced-droplet size through the local maximum stable (or critical) radius (r_{cr}). The velocity of the produced droplets is modeled using a Monte-Carlo procedure. As new droplets are formed, parent droplets are destroyed and Lagrangian tracking in the physical space is continued until further breakup events occur. The evolution of droplet diameter is basically governed by the local relative-velocity fluctuations between the gas and liquid phases. In this respect, LES plays a key role in providing accurate, local estimates of the gas-phase turbulent quantities. Although the mesh spacing used in a typical LES computation is larger than droplet size, the superiority of LES over RANS lies in accurate predictions of mixing and momentum transport from the gas phase to the spray field. The details of this model are given by Gorokhovski & Apte (2001).

10.2. *Hybrid particle-parcel technique for spray simulations*

Performing spray-breakup computations using Lagrangian tracking of each individual droplet gives rise to a large number of droplets ($>$ 50-100 million) very close to the injector. Computing such a large number of droplet trajectories is a formidable task even with supercomputers. In parallel computation of complex flows utilizing standard domain-decomposition techniques, the load balancing per processor is achieved by distributing the number of grid cells equally among all processors. Lagrangian particle tracking, however, causes load imbalance owing to the varying number of droplets per processor.

In order to overcome the above load-balancing problem, the usual approach is to represent a group of droplets with similar characteristics (diameter, velocity, temperature etc..) by a computational particle or 'parcel'. In addition, one carries the number of droplets per parcel as a parameter to be tracked. Since, a parcel represents a group of droplets (of the order of 100-1000), the total number of computational particles or trajectories to be simulated is reduced significantly. With breakup, the diameter of the parcel is sampled according to the procedure given above and the number of droplets associated with the particles is changed in order to conserve mass. This reduces the total number of particles per processor and increases the computational overhead with sprays by about 20% depending on the number of parcels used. Each parcel has all the droplet characteristics associated with it. The parcels methodology works well for RANS-type simulations where one is interested in time- or ensemble-averaged quantities. For LES,

however, we should ideally simulate as many droplet trajectories as possible in order to obtain time-accurate results.

A hybrid scheme involving the computation of both individual droplets and parcels is proposed. The difference between droplets and parcels is simply the number of particles associated with them, N_{par}, which is unity for droplets. During injection, new particles added to the computational domain are pure drops ($N_{par} = 1$). These drops move downstream and undergo breakup according to the above breakup model and produce new droplets. This increases the number of computational particles in the domain. In the dense-spray regime, one may obtain large number of droplets in a control volume and its immediate neighbors. The basic idea behind the hybrid approach is to collect all droplets in a particular control volume, and group them into bins corresponding to their size and other properties such as velocity, temperature etc. The droplets in bins are then used to form a parcel by conserving mass, momentum and energy. The properties of the parcel are obtained by mass-weighted averaging from individual droplets in the bin. For this procedure, only those control volumes are considered for which the number of droplets increases above a certain threshold value. The number of parcels created would depend on the number of bins and the threshold value used to sample them. The parcel thus created then undergoes breakup according to the above stochastic sub-grid model, however, does not create new parcels. On the other hand, N_{par} is increased and the diameter is decreased by mass conservation.

The effectiveness of this hybrid approach is demonstrated in the following computations. The implementation of this method in an unstructured LES code gives us the capability of testing and validating these models in realistic industrial geometries for various combustors with multiphase flows.

10.3. *Spray-breakup simulations*

Simplified combustor geometry

In order to validate the stochastic breakup model, a standard test case for spray atomization is simulated and compared with the experimental data of Hiroyasu and Kudota (1974) in a Diesel-engine configuration. The computational domain is a closed cylinder of length 13.8 cm and diameter 5.6 cm. A liquid jet is injected through a single-hole nozzle into this constant-pressure, room-temperature nitrogen chamber. Since the chamber temperature is low, evaporation of the liquid fuel is negligible. Large blobs of diameter 300 μm, corresponding to the injector size, are injected into the combustion chamber. Initially, there is no gas-phase flow inside the chamber. Gas-phase recirculation zones are created as the spray penetrates into the combustion chamber, through momentum-transfer between the gas and liquid phases. Three cases with different chamber pressures of 1.1, 3, and 5 MPa are simulated and compared with the experimental data. The corresponding flow parameters are indicated in table 3. The number of droplets injected per iteration is determined from the droplet diameter and time step by keeping the mass-flow rate constant. The time step used in the present simulation is 1.5 ms and a uniform grid of $64 \times 65 \times 65$ cells is found to capture the spray dynamics accurately.

The computations were performed using three different approaches: tracking and creation of all droplets, tracking of parcels, and the hybrid droplet-parcel algorithm. Results for the 1.1 MPa case are qualitatively compared in figures 16a-c, respectively. The size and location of each circle corresponds to that of the actual droplets in the computational domain. Figure 16a indicates that a broad spectrum of droplet sizes is present with the co-existence of large and small droplets. It should be noted that, when all droplet

Parameters	Case 1	Case 2	Case 3
P_{liq}, MPa	10	10	10
P_{gas}, MPa	1.1	3	5
Injection diameter, μm	300	300	300
Injection time, ms	2.5	4	5
Injection velocity, m/s	102	90.3	86.4

TABLE 3. Validation cases for stochastic breakup model in Diesel-engine configuration; Hiroyasu & Kudota (1974).

trajectories are computed a large number of droplets ($\sim$ 2-3 million) is obtained even at an early stage of the simulation. This simulation, however, depicts the complex interactions between the liquid and gas phases, the momentum coupling, and spray atomization due to stripping of small droplets. Figure 16b shows a similar simulation performed using the parcels-approach normally used in RANS-type computations. An extremely coarse (global) representation of the liquid core and atomization is obtained because new droplets are not created. Figure 16c, on the other hand, indicates the effectiveness of the hybrid approach. Here, the total number of computational particles is much smaller than those obtained when all the droplet trajectories are computed.

Close to the nozzle, the liquid core shows the existence of large and small droplets. Away from the nozzle, a global representation of droplets grouped to form parcels as well as small sparse droplets is observed. The computational overload due to the hybrid approach is significantly less ($\sim$ 50 times lower) in comparison with the computation of all droplet trajectories. The essential features of the spray dynamics are captured by the hybrid approach indicating its effectiveness and applicability in Eulerian-Lagrangian formulations.

Figure 17 shows the comparison of the spray-tip penetration depth as a function of time with the experimental data for the three cases investigated. Good agreement is obtained for all the three cases using the hybrid algorithm. The penetration depth decreases with increase in pressure. This could be attributed to the decreased injection velocity as well as strong damping of the liquid momentum by the denser gas-phase at higher pressures. The liquid core lengths were less than 2 mm in the above computations. The penetration depths predicted were weakly dependent on the grid size and the resolution used was sufficient to obtain good agreement with the experiments. The ligament-like liquid structures deflected outward are clearly visible and the spray angles produced for the three cases, 20, 23, and 25 degrees, respectively, are in close agreement with the experimental observations (Hiroyasu & Kudota 1974).

PW front end validation geometry

The stochastic model together with the hybrid particle-parcel approach were used to simulate the spray evolution from the Pratt & Whitney injector. The experimental data

set was obtained by mounting the injector in a cylindrical plenum through which gas with prescribed mass-flow rate was injected. The gas goes through the main and guide swirler to inject a swirling jet into the atmosphere. Liquid film is injected through the filmer surface which forms an annular ring. The liquid mass-flow rate corresponds to certain operating conditions of the gas-turbine engine. Measurements were made of the droplet distribution and liquid mass flux in the radial direction at two different axial locations away from the injector. Gas-phase statistics for mean and rms velocities is also available at these locations. The outside air-entrainment rates were measured and prescribed as inflow conditions. A snapshot of the spray evolution in the $z = 0$ plane, together with the gas-phase axial velocity contours, is shown in figure18. The hybrid approach used herein gives a dynamical picture with correct spray angle. Preliminary results show that the liquid mass fluxes at two downstream locations are in reasonable agreement with the experimental data. However, a longer-time sample is necessary to match the computational predictions with the experiments.

This computation was performed on 96 processors. The domain decomposition is based on the optimal performance of the Eulerian gas-phase solver. Due to breakup, a large number of droplets is created in the vicinity of the injector. With the hybrid approach, the total number of computational particles tracked after 6 ms is about 3.5 million, which represents approximately 13 million droplets. This includes about 150,000 parcels. The distribution of particles on 96 processors is shown by the histogram in figure 19. This implies that less than 30% of the total number of processors contain more than 10,000 computational particles. A preliminary solution to this problem is to use more processors, which would reduce the maximum number of computational particles per processor. A better approach is to develop an alternative domain-decomposition scheme with dynamic load balancing and additional weights for CVs containing particles.

11. Future plans

Our plans for the next year are as follows:

- Complete the simulation of reacting flow in the coaxial combustor geometry corresponding to the experiment of Spadaccini *et al.*
- Complete spray breakup simulation in Pratt & Whitney front-end validation geometry.
- Address the issue of load-balancing due to spray and investigate dynamic domain-decomposition techniques for Eulerian-Lagrangian computations of spray.
- Perform validation simulation for spray evaporation corresponding to the experiment of Sommerfeld & Qiu (1998).
- Implement a multi-level version of the geometric multigrid algorithm for the pressure-equation solver and investigate the advantages of using multigrid techniques to accelerate convergence of the momentum and scalar transport equations.
- Improve the agglomeration algorithm to optimize the geometric characteristics of agglomerated coarse grid control volumes. Alternate-directional agglomeration algorithms will be investigated.
- Implement commutative filters to estimate more accurately the filtered quantities that have to be evaluated in the dynamic procedure and to implement LES models with explicit filtering in which the filter width can be specified a priori by the user instead of being dictated by the local grid spacing.
- Generate the next-generation grids for the complex Pratt & Whitney combustor.

• Calculate the reacting flow in the full Pratt & Whitney geometry and validate by comparing with the experimental data provided by Pratt & Whitney. This calculation will involve testing not only the combustion model module but also the Lagrangian particle tracking module, evaporation and breakup modules.

Acknowledgments

We would like to acknowledge the Department of Energy's ASCI program for providing financial support for this work. We are also very grateful to the Pratt & Whitney combustor group for providing detailed experimental data to validate our code.

REFERENCES

Apte, S.V., Mahesh, K., & Moin, P. 2002 LES of particle-laden swirling flow in a coaxial combustor. To be submitted to *Int. J. Multiphase Flows*.

Brandt, A., Diskin, B., & Thomas, J. L. 2002 Recent advances in achieving textbook multigrid efficiency for computational fluid dynamics simulations. *ICASE Rept. No.* 2002-16.

Elias, S. R., Stubley, G. D., & Raithby, G. D. 1997 An adaptive agglomeration method for additive correction multigrid. *Int. J. for Num. Methods in Engg.* **40**, 887–903.

Gorokhovski, M. & Apte, S.V. 2001 Stochastic sub-grid modeling of drop breakup for LES of atomizing spray. *Annual Research Briefs*, Center for Turbulence Research, NASA Ames/Stanford Univ., 169–176.

Hiroyasu & M., Kudota, T. 1974 Fuel droplet size distribution in diesel combustion chamber. *SAE Tech. Paper* 740715.

Kolmogorov, A. N. 1941 On the log-normal distribution of particle sizes during breakup process. *Dokl. Akad. Nauk SSSR* **31**, 99–101.

Mavriplis, D. J., & Pirzadeh, S. 1999 Large-scale parallel unstructured mesh computations for 3D high-lift analysis. *ICASE Rept. No.* 1999-9.

Meneveau, C., Lund, T. S. & Cabot, W.H. 1996 A Lagrangian dynamic subgrid-scale model of turbulence. *J. Fluid Mech.* **319**, 353–385.

Montero, R. S., Llorente, I. M., & Salas, M. D. 2001 Robust Multigrid Algorithms for the Navier Stokes Equations. *J. Comp. Phys.* **173**, 412–432.

O'Rourke, P. J., & Amsden, A. A. 1987 The TAB method for numerical calculations of spray droplet breakup. *SAE Tech. Paper* 872089.

Pierce, C. D. & Moin, P. 2001 Progress variable approach for large eddy simulation of turbulent combustion. *Report TF-80*, Mech. Engg. Dept., Stanford Univ.

Raw, M. 1996 Robustness of coupled algebraic multigrid for the Navier-Stokes equations. *AIAA Paper* 96-0297.

Reitz, R. D. 1987 Modeling atomization processes in high-pressure vaporizing sprays. *Atom. & Spray* **3**, 309–337.

Sommerfeld, M. & Qiu, H. H. 1991 Detailed measurements in a swirling particulate two-phase flow by a phase-Doppler anemometer. *Int. J. Heat and Fluid Flow* **12**, 20–28.

Sommerfeld, M., & Qiu, H.H. 1998 Experimental studies of spray evaporation in turbulent flow. *Int. J. Heat & Fluid Flow* **19**, 10–22.

Spadaccini, L.J., Owen, F.K. & Bowman, C.T. 1976 Influence of aerodynamic phenomena on pollutant formation in combustion. (Phase I. Gaseous fuels). *U.S. Environmental Protection Agency Rept.* EPA-600/2-76-247a.

Wall, C., Boersma, B.J. & Moin, P. 2000 An evaluation of the assumed beta probability density function subgrid-scale model for large eddy simulation of nonpremixed, turbulent combustion with heat release. *Phys. Fluids* **12**, 2522–2529.