LES/Filtered-Density Function simulation of turbulent combustion with detailed chemistry

By V. Raman & H. Pitsch

1. Motivation and objectives

Large-Eddy Simulation (LES) has emerged as the next generation simulation tool for modeling turbulent flows. In order to extend the applicability of LES to reacting systems, a filtered-density function approach is considered here. By directly solving for the joint PDF (called FDF in the case of LES) of the composition, the reaction source terms appear in closed form and require no modeling. However, sub-filter mixing needs to be described using a model. In order to evolve the FDF in space and time, a transport equation for the joint-scalar FDF can be constructed and solved in a coupled manner with the LES momentum equations. However, the FDF transport equation spans a high-dimensional space and cannot be solved directly using conventional Eulerian discretization schemes (Pope 1981). Typically, a Lagrangian stochastic approach is employed (Pope 1994). It is evident that due to the inherently unsteady nature of the LES methodology, the FDF scheme needs to be time-accurate, as well. The stochastic nature of the FDF algorithm poses a considerable numerical challenge, since statistical noise in mean fields can destabilize the LES solver (Raman *et al.* 2005). In addition, the Lagrangian approach is computationally expensive, and robust solvers need to be formulated. It is argued here that for this combustion model to be practically useful, the LES-FDF algorithm should be discretely consistent. It should also be demonstrably accurate and easily adaptable for handling multiple configurations. The objective of this work is to formulate such a solver and to test it with practical flame configurations.

The conditions specified above as requirements of a combustion model can be viewed as algorithmic constraints. Each of the above constraints is not naturally satisfied by currently available algorithms for the LES-FDF methodology (Colucci *et al.* 1998; Jaberi *et al.* 1999). In spite of consistency at the equation level, numerical errors as well as variations in the nature of the solvers can lead to discrete inconsistency. In this work, consistency between the two parts is established by comparing redundant fields in the LES and FDF solver. Through further comparisons with the moments of the conserved scalar, higher order equivalence is demonstrated. Using this new algorithm, a detailed-chemistry model is included, for the first time, in the LES-FDF simulation to compute two complex unsteady experimental flames. The bluff-body stabilized flame configuration is used to study the effect of particle number density on the stability and the accuracy of the algorithm. The Sandia flame configuration is then simulated using a detailed chemistry mechanism for two different flow conditions. By formulating a dynamic mixing model, the predictive capability of the LES-FDF algorithm for flames with local extinction is demonstrated. This is the first attempt to use a large chemistry mechanism in an FDF calculation. It is emphasized that such an expensive calculation will be futile if the issues regarding consistency, accuracy and stability are not addressed.

2. Background

The large-eddy simulation technique resolves or directly obtains all large scale structures of the flow and requires models for the small-scale or unresolved features of the flow. Typical combustion chemistry occurs at length scales far smaller than the resolved length scales in practical simulations (Peters 2000). This necessitates a reaction model that uses resolved scale information to predict reaction progress. Currently, several models are available for describing chemically reacting flows and have been used to simulate experimental flames with complex flow features (Peters 2000). Although the filtered-density function approach has the advantage that the chemical source terms appear closed, the molecular mixing term still needs to be modeled (Pope 2000; Fox 2003). Since LES resolves a large portion of the scalar energy spectrum, it is possible that the mixing models are less important, or the current models are accurate enough to ensure predictive accuracy. The relative effect of the mixing model compared to that in RANS based PDF simulations, in which they are critical, is yet to be studied. Before such a detailed analysis is possible, a stable, consistent and numerically accurate simulation strategy has to be established. The aim of the current work is to propose such a scheme and to test it with different complex flow problems.

3. Hybrid LES-FDF method

Stand-alone particle-based PDF methods have been used to model several experimental flames in the past. Numerical instability associated with particle-based schemes has led to hybrid schemes in which an Eulerian flow solver is coupled with the particle scheme (Correa & Pope 1992). Usually, the Eulerian mesh-based part computes the momentum equations while the particle-based solver evolves the joint-PDF transport equation. In the context of LES, the spatially filtered variant of the fine-grained RANS PDF is called the Filtered-Density Function (FDF) (Gao & O'Brien 1993). In the current study, a joint-composition FDF is solved using the Monte-Carlo method. The transport equation for the FDF can be written as (Colucci *et al.* 1998; Jaberi *et al.* 1999)

$$\frac{\partial F_L}{\partial t} + \nabla \cdot (\widetilde{\mathbf{u}} F_L) + \nabla \cdot \left(\widetilde{\mathbf{u}'|\boldsymbol{\psi}} F_L \right) = -\frac{\partial}{\partial \boldsymbol{\psi}} \left[\left(\frac{1}{\overline{\rho}} \overline{\nabla \cdot \rho D \nabla \boldsymbol{\phi} | \boldsymbol{\psi}} + \mathbf{S}(\boldsymbol{\psi}) \right) F_L \right], \tag{3.1}$$

where F_L is the FDF, $\mathbf{u}$ is the Favre-filtered velocity field, $\widetilde{\mathbf{u}'|\boldsymbol{\psi}}$ is the sub-filter velocity fluctuation conditioned on the scalar, $\overline{\nabla \cdot \rho D \nabla \boldsymbol{\phi} | \boldsymbol{\psi}}$ is the conditional micromixing term, and $\mathbf{S}$ is the reaction source term. As mentioned earlier, the reaction source term appears closed and requires no modeling. The conditional mixing term needs closure, and several models have been proposed (Villermaux & Falk 1994; Curl 1963; Janicka *et al.* 1970; Subramaniam & Pope 1999). Here, the Interaction by Exchange with the Mean (IEM) model (Villermaux & Falk 1994) is used:

$$\overline{\nabla \cdot \rho D \nabla \boldsymbol{\phi} | \boldsymbol{\psi}} F_L = \nabla \cdot \overline{\rho} D \nabla F_L / \overline{\rho} - \frac{C_\phi}{\tau} \left(\boldsymbol{\psi} - \widetilde{\boldsymbol{\phi}} \right) F_L, \tag{3.2}$$

where C_ϕ is the scalar-to-mechanical time-scale ratio and τ is a turbulence time scale.

The key aspect of the above equation is the high dimensionality, that essentially makes any finite-difference based solution method intractable for even small chemical mechanisms with only a few chemical species. Pope (1981) suggests a particle-based Monte-Carlo algorithm in which the computational domain is decomposed into a large number of notional particles. These particles evolve in physical as well as in compositional spaces

with time. The stochastic evolution equations for this particle-based equivalent system can then be written as (Pope 2000; Colucci *et al.* 1998)

$$d\mathbf{x}^* = \left[\widetilde{\mathbf{u}} + \frac{1}{\overline{\rho}}\nabla\overline{\rho}\,(D + D_T)\right] dt + \sqrt{2\,(D + D_T)}d\mathbf{W}, \tag{3.3}$$

where $\mathbf{x}^*$ is the instantaneous particle position, and $d\mathbf{W}$ is the Wiener diffusion term characterized by a Gaussian process with zero mean and variance of dt, where dt is the time-step of the process. All filtered variables are interpolated to particle positions. Transport in composition space occurs through mixing and reaction:

$$d\boldsymbol{\psi} = -\frac{C_\phi}{\tau}\left(\boldsymbol{\psi} - \widetilde{\boldsymbol{\phi}}\right) dt + \mathbf{S}(\boldsymbol{\psi})dt. \tag{3.4}$$

As readily observed, the stochastic evolution equations use the filtered Eulerian fields to advance the particles. In the hybrid scheme, the filtered velocity fields are provided through an external flow solver. In return, the particle-fields are used to evaluate density changes due to reaction, which are then fed back to the flow solver. The Eulerian solver then advances the flow fields taking into account the density changes. Next, the different solvers and the feedback mechanism are discussed.

The mixing time-scale, τ that appears in the mixing model needs to be modeled. In addition, the scalar-to-mechanical time-scale ratio, C_ϕ needs to be specified. Here, a dynamic formulation is proposed in which the scalar mixing time scale is obtained based on a local equilibrium assumption. The details of the model will be provided elsewhere. Using this approach, the mixing time-scale, C_ϕ/τ can be specified as

$$\frac{C_\phi}{\tau} = \frac{\chi}{\widetilde{Z''^2}}, \tag{3.5}$$

where χ is the local scalar dissipation rate for the mixture-fraction. Both χ and the mixture-fraction variance $\widetilde{Z''^2}$ are obtained using dynamic models, thereby suggesting that the mixing time-scale adapts to the flow. In this work, the bluff-body stabilized flame is simulated using a conventional model for the time-scale (Raman *et al.* 2005), while the Sandia flame series is simulated using the above dynamic model.

4. Numerical implementation

The momentum equations corresponding to the filtered Navier-Stokes equations are solved in cylindrical coordinates using a low-Mach number staggered scheme (Pierce 2001; Akselvoll & Moin 1996). The sub-filter stress terms are closed using dynamic models (Moin *et al.* 1991). For the sake of comparison, Eulerian transport equations of the filtered mixture-fraction and its second moment, $\widetilde{Z^2}$, are solved. An enthalpy transport equation is also solved using the Eulerian scheme. A structured grid is used along with domain decomposition based MPI parallelism.

Although the Monte-Carlo algorithm has been used for computing several experimental flame configurations, its use in hybrid LES-FDF techniques needs further improvements. LES is inherently unsteady and requires that the stochastic algorithm be time accurate. In addition, a consistent feedback mechanism needs to be established that will preserve the accuracy of the scheme and also provide numerical stability. To resolve these issues, a consistent numerical scheme is devised and implemented in this work. Essential features of the algorithm are described here. A full description of the algorithm and implementation will be detailed elsewhere.

The Monte-Carlo algorithm for the FDF evolution uses a Lagrangian implementation with evolution equations defined by (3.3) and (3.4) and modified for a cylindrical coordinate system. Here, the particle velocities are evaluated in the Cartesian reference frame and then converted back to the cylindrical reference. The particles are initially distributed uniformly in the domain. Each particle carries a multi-dimensional vector that locates the particle in physical space (3-dimensional) and in composition space (N-dimensional, where N is the number of thermo-chemical variables). In addition, the particles also carry a weight that is initially set to be proportional to the local cell mass.

The particle evolution equations in physical space require interpolation of the filtered-velocity field onto the particle location. In this work, a trilinear interpolation algorithm is used to obtain particle values. In a finite-volume formulation, interpolation algorithms should ensure that mass-conservation is satisfied in an integral sense. The LES code prescribes a face-centered velocity field. The mass flux through the face is computed, using a one-point quadrature with uniform velocity in the face-normal direction. However, a trilinear interpolation implies a bilinear variation along the face (Zhang & Haworth 2004). To remove this inconsistency, a correction velocity is applied that adjusts the implied flux of the trilinear scheme to that of the LES scheme.

The numerical accuracy of the simulations depends to a large extent on the number of particles per computational volume. The particle weights, as noted earlier, are proportional to the local cell mass. The cylindrical coordinate system leads to very small particle mass near the centerline and large values in the outer regions. In addition, to capture the velocity gradients more accurately, grid refinement is employed in most CFD simulations. Since such gradients occur mainly near the inner fuel jet (in a set-up typical of non-premixed combustion), this leads to smaller computational cells and lower particle mass near the centerline. Such large variations in the local cell mass lead to a non-uniform distribution of the particle numbers as the simulation advances in time. To rectify this problem, a passive particle-clustering/splitting algorithm (Raman *et al.* 2005) is employed. By splitting particles that are heavier compared to the cell average and by grouping those that are lighter than the cell average, the number density can be controlled. However, such operations lead to a diffusion in composition space and are therefor employed only when the particle number density falls outside a prescribed range.

For the integration of particle equations in time, higher-order methods have been devised (Cao & Pope 2003). In this case, a modified Euler scheme (Cao & Pope 2003) is employed that yields better convergence and increased stability of the overall algorithm compared to the first-order Euler scheme. The computational requirements are comparable to the first-order scheme. A face-to-face tracking strategy is used for moving the particles in physical space. This algorithm naturally extends to parallel computations as well, and was found to scale linearly (for parallel processing) with particle numbers.

In an ideal set-up the above algorithm should be consistent and stable. However, the stochastic nature of the particle scheme leads to instabilities in practical implementations, unless a very large, computationally intractable particle number density is used. It can be shown that an initial uniform distribution of particles will remain uniform, provided that the continuity equation is satisfied (Pope 2000). However, due to finite particle numbers and errors in velocity field interpolation, the mean fields contain statistical bias and noise. Any direct feedback of the particle-based fields to the deterministic flow solver will lead to instabilities and divergence of the numerical scheme. To remedy this problem, an enthalpy transport equation is used to indirectly feed the density change due to reaction back to the finite volume solver. The equivalent enthalpy is defined in

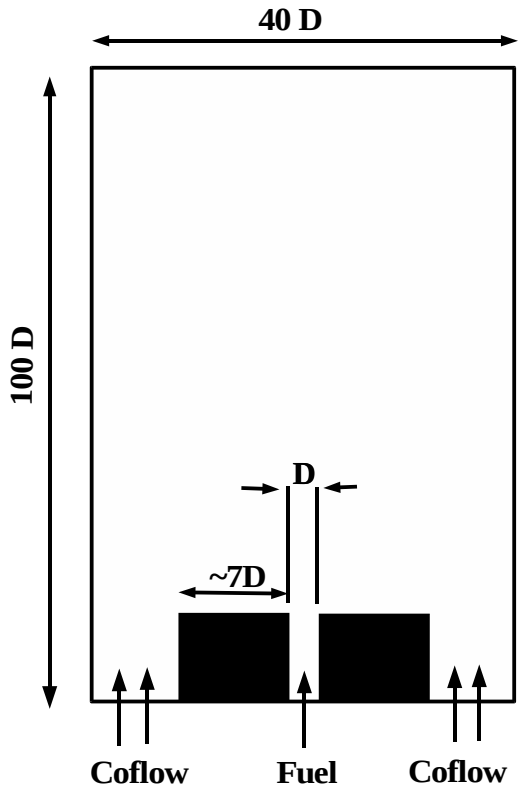

FIGURE 1. Schematic of the computational domain.

terms of the density and contains a source term proportional to the rate of change of density (Muradoglu *et al.* 1999; Raman *et al.* 2005). The particle composition is used to compute the exact source term and is then used to advance the equivalent enthalpy transport equation. The time-advanced density field for the LES scheme can then be obtained from the current enthalpy field. This method was found to provide very good stability, since the numerical diffusion in the Eulerian transport scheme smoothes out the large spikes in the density change obtained from the particles.

5. Application to a bluff-body stabilized flame

The objective of the simulations performed here is to illustrate the consistency, accuracy and feasibility of the LES-FDF algorithm. For this purpose, the hybrid LES-FDF scheme described above is used to simulate a complex non-premixed experimental configuration (Dally *et al.* 1998). The setup consists of a central fuel jet with a 1:1 volume

ratio of methane and hydrogen issuing with a bulk velocity of 118 m/s. The coflow is pure air with a speed of 40 m/s. The fuel jet and coflow are separated by an annular solid bluff-body (Fig. 1). This configuration allows for flame stabilization through enhanced mixing due to the presence of the bluff-body. It has been observed from mean velocity profiles that strong counter-rotating vortices help mix cold fuel and air with hot product mixture in order to produce a stable burning mechanism.

The complex flow patterns induced by the bluff-body are highly unsteady and cannot be captured by RANS-type methods. It has been shown elsewhere that the use of LES helps to capture all mean flow profiles quite accurately (Raman *et al.* 2005). Here, several simulations are used to establish consistency, accuracy, and feasibility of the hybrid scheme for simulating such complex flows using detailed chemistry. Three different types of simulations were performed. In the first case, the Eulerian scheme and a flamelet model was used to simulate the flame (Raman & Pitsch 2005). The mixture-fraction and variance fields were obtained directly from the Eulerian transport equations for these scalars (sim A). In the second simulation, the LES-FDF scheme was used with the FDF method evolving the mixture-fraction scalar in a fully-coupled mode (sim B). Again, the flamelet model was used to describe the turbulence/chemistry interactions. In the third simulation, a fully coupled joint composition FDF was used with direct integration of the chemical source term (sim C). To the authors' knowledge, this is the first-ever variable-density full-FDF simulation of an experimental configuration. A 16-species chemistry mechanism along with the temperature was used to describe the local composition (Pepiot & Pitsch 2004). Direct integration of the chemical source term is computationally intensive and several logical conditions are used to reduce the number of such integrations. In particular, a substantial number of particles is present in the coflow, where no reactions take place. Additionally, in several cells, the sub-filter variance of the composition vector is very small and can therefore be considered as a single particle for the purpose of direct integration. Also, a neighborhood search for particles with "close" composition vectors is used to further minimize the computational time. Except for sim C, which is computationally expensive, the other simulations were repeated for several different values of parameters such as particle number density in order to analyze the effect on the final time-averaged solution.

5.1. *Bluff-body stabilized flame configuration*

Figure 1 is a ray-traced image of the bluff-body stabilized flame. It is evident that this flow configuration leads to complex unsteady effects. In particular, the outer edge of the bluff-body is known to exhibit vortex shedding that interacts with the outer shear layer in the recirculation region close to the bluff-body. Such an unsteady effect reduces the peak temperature observed in this region and affects the formation of the NO in the outer shear layer. Stream-traces of the time-averaged velocity field show the presence of two counter-rotating vortices in the recirculation region that stabilize the flame configuration (Raman & Pitsch 2005; Raman *et al.* 2005). A time-series analysis of the bluff-body flame from the simulations performed here shows that the strength and location of the vortical structures vary significantly in time. In fact, the recirculation region is continuously destroyed and re-formed due to the interaction of the vortex-shedding with the outer shear layer.

5.2. *Accuracy*

Figure 2 shows time-averaged radial profiles of the mean mixture-fraction obtained with different particle-number densities (sim B) compared with the Eulerian solution of the

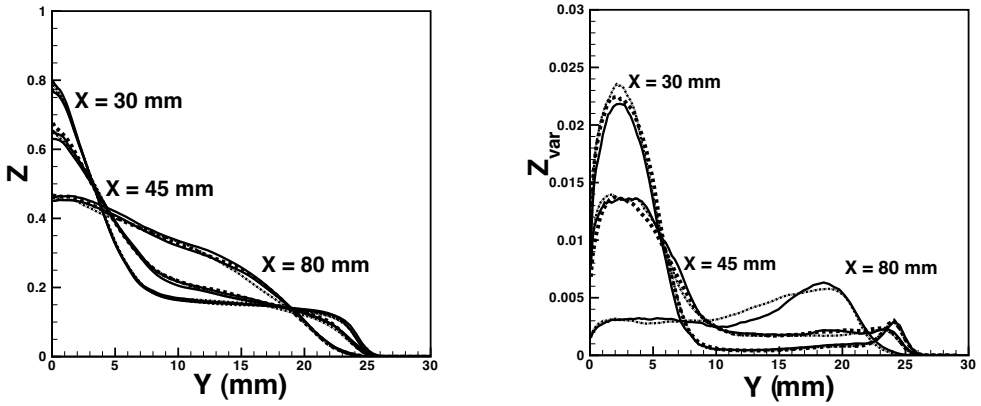

FIGURE 2. Radial profiles of the time-averaged mixture-fraction obtained from FDF with particle number densities of 20 (solid), 80 (dashed), and 120 (dashed-dotted) compared with the Eulerian solution (dotted).

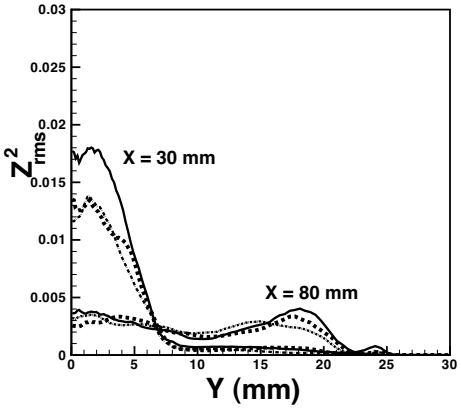

FIGURE 3. Radial profiles of the resolved mixture-fraction variance obtained from FDF with particle number densities of 20 (solid), 80 (dashed), and 120 (dashed-dotted).

mixture-fraction transport equation (sim A). It can be observed that the differences in the mean profiles are minor. Even at axial positions away from the bluff-body, the profiles agree very well, thereby indicating that a reduced number of particles is sufficient to accurately evolve the FDF transport equation, at least for a single scalar. To further investigate the accuracy, higher moments of the scalar need to be compared. In the context of LES, two different mixture-fraction variances can be defined. The sub-filter variance is evaluated based on the particle distribution within a cell and can be directly obtained from the FDF transport equation. The time-resolved variance, on the other hand, is equivalent to RANS-type variance in which the local fluctuations of the mixture-fraction are quantified. In terms of the numerics, the accuracy of the sub-filter variance is controlled by the extent of numerical diffusion in composition space. On the other hand, the accuracy of the time-resolved variance depends on the stochastic noise in the particle-based filtered mixture-fraction field. Both these quantities can therefore be used

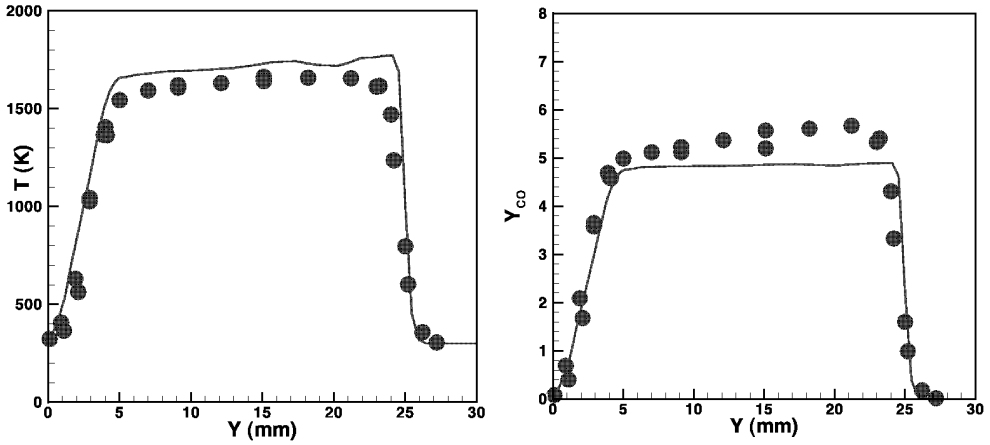

FIGURE 4. Radial profiles of temperature (left) and CO (right) at an axial position of $X = 13$ mm obtained using a 16-species finite-chemistry LES-FDF simulation. Symbols are experimental data and lines show simulation results.

to quantify the accuracy of the simulation. Figure 2 also shows the sub-filter variance obtained using different particle numbers (sim B). It is clear that the behavior is similar to the filtered mixture-fraction and shows negligible changes with particle number density.

Figure 3 shows the RMS of the filtered mixture fraction or the temporal variance. This plot shows that for low particle number densities in which the stochastic noise is high, the results are drastically different from the cases with higher particle number densities. With higher particle number density, the convergence rate of the time-averaged variance is higher. This phenomenon is directly related to the intensity of turbulence or to the turbulent diffusivity experienced by the particles. The particle evolution equations show that the Wiener diffusion term depends on the local turbulent diffusivity. Since this term contributes solely to the stochastic fluctuations, regions with higher local diffusivity will exhibit larger fluctuations in mean fields. Another contributing factor is the local grid refinement, which can result in cells with disparate fluid masses being adjacent to each other. This will cause large jumps in particle weights from one time-step to the next and can lead to fluctuations in the mean fields, as well. Due to the nature of the problems considered here, cells with large diffusivity, and regions of grid refinement are both present in the same regions of the flow. The inner shear layer contains large velocity gradients requiring grid refinement. Also, the scalar gradients present in this region lead to large diffusivity based on the dynamic model. Clearly, the plot also shows that at locations away from the centerline, the variance profile becomes smooth and exhibits smaller variations with number density. Away from the centerline, the local flow field is smoother, with smaller velocity and scalar gradients thereby leading to lower statistical fluctuations. In spite of these large fluctuations, the low particle number density case still exhibits numerical stability, mostly due to the indirect feedback mechanism. Hence, for a stable simulation, it can be estimated that particle number densities in the order of 30-50 particles are sufficient. However, it should be noted that the error increases with decrease in the particle number density. Also, the effect of low particle number densities for problems with detailed chemistry could be substantially different and needs to be studied further.

5.3. *Feasibility*

Figure 4 shows time-averaged radial profiles of temperature and CO mass fraction obtained using the detail chemistry simulation with the fully coupled LES-FDF scheme (sim C). The computation was carried out on 32-64 processors for several thousand hours (total computational time across all CPUs). The particle fields were time-averaged for 1.5 residence times, with the residence time computed based on the time it takes for a particle to travel along the centerline from inlet to outlet. It was found from previous simulations (Raman & Pitsch 2005) that longer time-averaging did not alter the profiles. The comparison with experimental data shows that the agreement is quite good. More detailed analysis of the results will be discussed elsewhere.

6. Application to the Sandia flame configuration

The Sandia flame series (Barlow & Frank 1998) consists of a series of experiments using a piloted partially-premixed flame configuration. The flames are studied using increasing inlet velocity of the fuel and oxidizer jets. The increased velocities increase the probability of local extinction. The purpose of this study is to demonstrate the predictive capability of the LES-FDF algorithm. Here, the D and E flame configurations are simulated using a detailed description of the methane/air chemistry. By using a dynamic model for the mixing time-scale, the need to specify modeling constants *a priori* is removed.

The Sandia flame configuration consists of a central fuel jet with a diameter of 7.2 mm surrounded by a pilot flame with an outer diameter of 18.2 mm. A coflow of air serves as the oxidizer. The fuel jet is comprised of methane (25% by vol.) diluted in air (75% by vol.). The central jet velocity for flame D is 49.6 m/s while for flame E it is 74.4 m/s. The pilot issues at 11.4 m/s for flame D and 17.1 m/s for flame E. The coflow velocity is fixed at 0.9 m/s for all cases. The simulations performed here use a 17-species chemistry mechanism reduced from the GRI-2.11 mechanism (Pepiot & Pitsch 2004). A $256 \times 128 \times 32$ computational grid was used in the simulations. Grid clustering was used to refine the strong gradients near the inner shear layer formed between the fuel jet and the pilot, as well as near the outer shear layer formed between the pilot and the coflow. The coupled simulation was performed with 40 particles per computational cell. Direct numerical integration of the chemical source term was used. The dynamic time-scale model was employed.

Figure 5 shows the instantaneous temperature profile for flames D and E. The temperature contours indicate that flame D has low levels of extinction whereas local extinction events are more common in flame E. Both these flames have relatively simple flow features with the reaction rate interlinked to the local strain in the two shear layers. Near the inlet, the presence of the pilot flame leads to enhanced stability with minimal extinction for both the flames. Further downstream, the outer coflow and the inner fuel jet interact in the high-temperature region that is present in the shear layer. This configuration resembles a simple diffusion flame and has been widely studied using flamelet models (Pitsch & Steiner 2000; Pitsch 2003). It has been observed that although the laminar flamelet model does not predict the minor species accurately, use of an unsteady flamelet model that takes into account the time-history of scalar dissipation rate fluctuations will enhance the predictive capability of the method. Here, the use of the dynamic time-scale model tends to take into account the mixing time-scale fluctuations, thereby providing a better prediction of the flame characteristics.

Figure 6 shows the conditional averages conditioned on the mixture-fraction. These

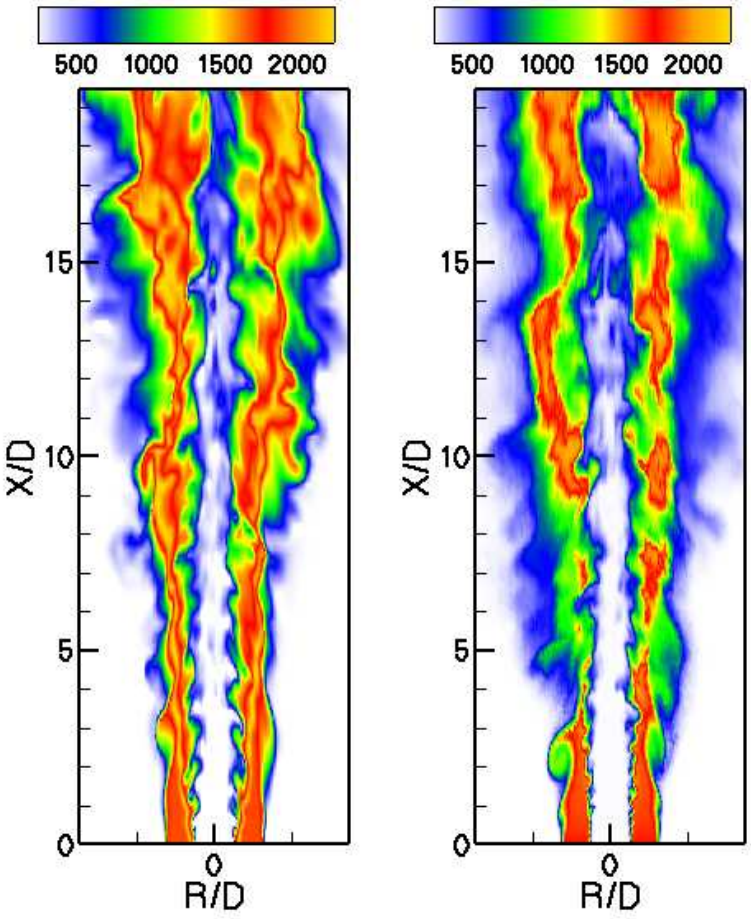

FIGURE 5. Instantaneous temperature contours for flame D (left) and flame E (right).

plots serve to illustrate the accuracy of the simulation in predicting transport in composition space. The conditional O_2 profiles for both flames D and E directly indicate the level of local extinction. When the flame is fully burning, the O_2 mass fraction should fall to very low (but non-zero) values at stoichiometric mixture-fraction value. For both flames, the O_2 level indicates some level of extinction. More noticeably, for flame E, the O_2 mass fraction at stoichiometric value is nearly twice that of flame D, indicating higher levels of extinction. The simulations tend to under-predict extinction in flame D, but to correctly predict extinction levels for flame E. Consequently, the conditional mean of the temperature is over-predicted for flame D while it is accurately predicted for flame E. It is clear that the use of the dynamic time-scale model adapts the mixing model to the flow configuration, thereby providing a quantitatively better prediction.

7. Conclusions

A consistent, stable and numerically accurate LES-FDF scheme has been formulated for turbulent non-premixed combustion. Using appropriate correction schemes, an accu-

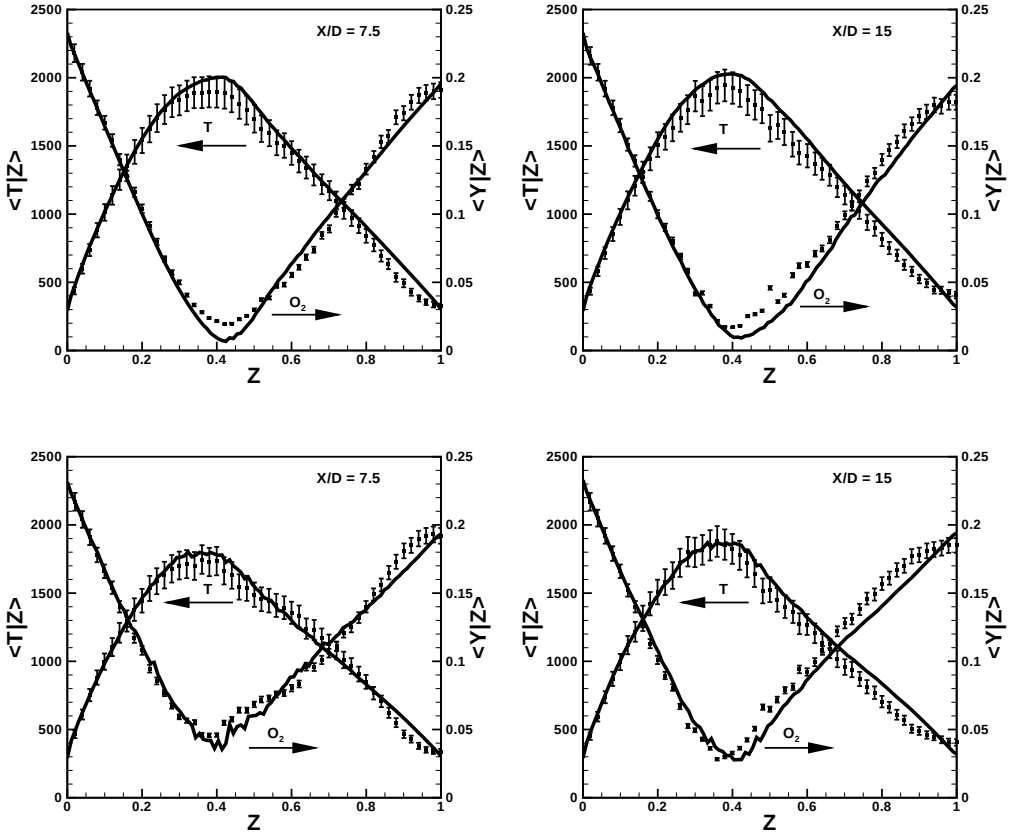

FIGURE 6. Conditional mean temperature and O_2 mass fraction at two different axial locations for flame D (top row) and flame E (bottom row)

rate algorithm has been implemented in the cylindrical coordinate reference frame. A non-premixed bluff-body stabilized flame was used to demonstrate the accuracy and the feasibility of the new scheme. Using Eulerian solutions of the mixture-fraction, accuracy was established. For the first time, a full LES-FDF simulation with a 16-species detailed chemistry mechanism was used to simulate the flame and good agreement was obtained between simulation results and experimental data. This consistent algorithm was then applied to the simulation of two Sandia flame configurations. By using a dynamically determined time-scale ratio, all tunable constants in the simulation are removed. Simulations of flame D and E show that this new approach has excellent predictive capability and can capture local extinction in these flames.

Acknowledgments

The authors would like to thank Prof. Rodney Fox for many useful discussions. The authors would also like to acknowledge the NASA Supercomputing Center for computing time on their COLUMBIA system.

REFERENCES

Akselvoll, K. & Moin, P. 1996 Large eddy simulation of turbulent confined coannular jets. *J. Fluid Mech.* **315**, 387–411.

Barlow, R. S. & Frank, J. H. 1998 Effects of turbulence on specific mass fractions in methane/air jet flames. *Twenty-Seventh Symposium (International) on Combustion* **27**, 1087–1095.

Cao, R. & Pope, S. B. 2003 Numerical integration of stochastic differential equations: Weak second-order mid-point scheme for application in the composition PDF method. *J. Comp. Phys.* **185**, 194–212.

Colucci, P. J., Jaberi, F. A. & Givi, P. 1998 Filtered density function for large eddy simulation of turbulent reacting flows. *Phys. Fluids* **10**, 499–515.

Correa, S. M. & Pope, S. B. 1992 Comparison of a Monte Carlo PDF finite-volume mean flow model with bluff-body data. In *Twenty-Fourth Symposium (International) on Combustion*, 279–285. Combustion Institute, Pittsburgh.

Curl, R. L. 1963 Dispersed phase mixing I. Theory and effects in simple reactors. *AIChE J.* **9**, 175–181.

Dally, B. B., Masri, A. R., Barlow, R. S. & Fietchner, G. J. 1998 Instantaneous and mean compositional structure of bluff-body stabilized nonpremixed flames. *Combust. Flame* **114**, 119–148.

Fox, R. O. 2003 *Computational Models for Turbulent Reacting Flows*. Cambridge University Press.

Gao, F. & O'Brien, E. E. 1993 A large-eddy simulation scheme for turbulent reacting flows. *Phys. Fluids* **5**, 1282.

Jaberi, F. A., Colucci, P. J., James, S., Givi, P. & Pope, S. B. 1999 Filtered mass density function for large-eddy simulation of turbulent reacting flows. *J. Fluid Mech.* **401**, 85–121.

Janicka, J., Kolbe, W. & Kollmann, W. 1970 Closure of the transport equation for the PDF of turbulent scalar fields. *J. Non-Equilibrium Thermody.* **4**, 47–66.

Moin, P., Squires, K., Cabot, W. & Lee, S. 1991 A dynamic subgrid-scale model for compressible turbulence and scalar transport. *Phys. Fluids A* **3**, 2746–2757.

Muradoglu, M., Jenny, P., Pope, S. B. & Caughey, D. A. 1999 A consistent hybrid finite-volume/particle method for the PDF equations of turbulent reactive flows. *J. Comp. Phys.* **154**, 342–371.

Pepiot, P. & Pitsch, H. 2004 A 16-species reduced chemistry mechanism for combustion of methane/hydrogen fuel. Unpublished.

Peters, N. 2000 *Turbulent Combustion*. Cambridge University Press.

Pierce, C. D. 2001 Progress-variable approach for large-eddy simulation of turbulence combustion. *Ph.D. Thesis*, Stanford University.

Pitsch, H. 2003 Improved pollutant predictions in large-eddy simulations of turbulent non-premixed combustion by considering scalar dissipation rate fluctuations. *Proc. Combust. Inst.* **29**, 1971–1978.

Pitsch, H. & Steiner, H. 2000 Large-eddy simulation of a turbulent piloted methane/air diffusion flame (Sandia flame D). *Phys. Fluids* **12**, 2541–2554.

Pope, S. B. 1981 A Monte Carlo method for the PDF equations of turbulent reactive flow. *Combust. Sci. Tech.* **25**, 159–174.

Pope, S. B. 1994 Lagrangian PDF methods for turbulent flows. *Ann. Rev. Fluid Mech.* **26**, 23–63.

Pope, S. B. 2000 *Turbulent Flows*. Cambridge University Press.

Raman, V. & Pitsch, H. 2005 Large-eddy simulation of bluff-body stabilized non-premixed flame using a recursive-refinement procedure. *Combust. Flame* **142**, 329–347.

Raman, V., Pitsch, H. & Fox, R. O. 2005 A consistent hybrid LES-FDF scheme for the simulation of turbulent reactive flows. *Combust. Flame* **143** (1-2), 56–78.

Subramaniam, S. & Pope, S. B. 1999 Comparison of mixing model performance for non-premixed turbulent reactive flow. *Combust. Flame* **117**, 732.

Villermaux, J. & Falk, L. 1994 A generalized mixing model for initial contacting of reactive fluids. *Chem. Eng. Sci.* **49**, 5127–5140.

Zhang, Y. Z. & Haworth, D. C. 2004 A general mass consistency algorithm for hybrid particle/finite-volume pdf methods. *J. Comp. Phys.* **194**, 156–193.