PDF Simulations of Turbulent Combustion Incorporating Detailed Chemistry

VIVEK SAXENA* and STEPHEN B. POPE
Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853 USA

Recently, the technique of in situ adaptive tabulation (ISAT) was proposed to accomplish efficiently the hitherto formidable task of including detailed chemistry in particle methods for turbulent-reactive flows. This paper addresses the performance issues (accuracy, efficiency, and storage requirements) involved in the implementation of ISAT in a particle-mesh code to solve for the joint probability density function (JPDF) of velocity-frequency-composition.

Detailed error-control analysis for a skeletal methane mechanism is performed using a pairwise mixing stirred reactor (PMSR) test. The results of the PMSR tests are used to specify the error tolerance for the JPDF simulation of a piloted non-premixed jet flame. We demonstrate that JPDF simulations of realistic turbulent-reactive flows with detailed chemistry are now possible within reasonable CPU time and storage requirements.

INTRODUCTION

It is widely held that the probability density function (PDF) approach is the most suited method for turbulent-reactive flows since complex reactions can be treated without modeling assumptions. In the computational implementation of PDF methods [1], the fluid within the solution domain is represented by a large number of computational particles. The composition of each particle evolves according to a set of ordinary differential equations. Typically, a simulation based on particle methods involves 10^6 particles evolving over 10^3 time steps. Consequently, the number of integrations for the stiff ordinary differential equations, representing the composition evolution, may be of the order of a billion. For example, on an SGI Indigo workstation used for the present study, with the 16-species mechanism for methane combustion used in the tests reported below, the average CPU time required to numerically integrate these equations over a single time step is 0.05 s. A straightforward numerical integration is referred to as *direct integration* or DI. Over the whole simulation period (i.e., about a billion such integrations), DI will require 1.5 years of CPU time. This is generally deemed to be prohibitively expensive.

A number of reduction, storage, and retrieval techniques have been employed in the past to deal with the problem of excessive computer time requirement. One approach is to reduce radically the number of degrees of freedom in the description of chemistry (see e.g., [2–4] and [5]). The number of degrees of freedom is reduced typically to four or less, at the cost of introducing assumptions of uncertain generality and accuracy. Another technique, used in PDF methods, involves tabulation (e.g., [6, 7]). In the preprocessing stage of this approach, DI is used to construct a table of values to cover the realizable region of the composition space. Table look-up, with a multi-linear interpolation, is used to determine composition evolution during the PDF calculations. However, the direct application of tabulation to detailed kinetics is not feasible. For example, with $D = 10$ degrees of freedom and with a mesh of 10 nodes in each direction, the resulting table will require $D10^D = 10^{11}$ words of storage and one multi-linear interpolation would require at least $D2^D \approx$ 10,000 operations.

Recently, a new algorithm called in situ adaptive tabulation (ISAT) was proposed by Pope [8] for the efficient computation of detailed chemistry. ISAT has been shown to reduce the computer time needed to solve the reaction equations by up to three orders of magnitude, with reasonable storage requirement, and with good control of interpolation errors. On a separate front, the Monte Carlo method has been successfully applied to solve the velocity-frequency-composition joint PDF equations [9]

---

*Corresponding author. E-mail: saxen@mae.cornell.edu

COMBUSTION AND FLAME 117:340–350 (1999)

0010-2180/99/$–see front matter
PII S0010-2180(98)00081-9

and has been implemented for various flows, e.g., [10–12] and [13]. Here, we present the results of implementing the ISAT scheme with the above JPDF approach.

PDF calculations of piloted turbulent nonpremixed flames have been attempted in the past by Masri and Pope [10], Masri, Subramaniam, and Pope [14], and Norris and Pope [15] to demonstrate the application of velocity-frequency-composition joint PDF method. The conserved scalar approach [16] was used in the first work, while an assumption of self-similar chemistry was made for the second study. The third study used the intrinsic low-dimensional manifold (ILDM) technique for reducing chemistry. Chen et al. [17] have used a hybrid scheme that employs composition PDF method for scalars with a second moment closure method for modeling turbulent velocity statistics. Two reduced mechanisms involving four and five scalars were used for modeling reaction and a modified coalescence-dispersion model was used for modeling mixing. The comparisons with experimental data were encouraging and a need to implement detailed chemistry was clearly recognized to predict the finite-rate kinetic effects in all of these studies. However, as explained above, the computer time and the storage requirements of such an endeavor have been a deterrent.

COMPOSITION EVOLUTION IN PDF METHODS

In the computational implementation of the PDF methods [1], the fluid within the solution domain is represented by a large number of computational particles. At any point and time in a reactive gaseous flow, the thermochemical state of the mixture can be characterized by the mass fractions Y_i ($i = 1, 2, \ldots, n_s$) of the n_s species, the enthalpy h, and the pressure P. We consider the broad class of flows in which P differs by a very small fraction from a fixed reference pressure P_0, so that, given P_0, the state is determined by

$$\begin{aligned}\hat{\phi} &= \{\hat{\phi}_1, \hat{\phi}_2, \ldots, \hat{\phi}_{n_s+1}\} \\ &= \{Y_1, Y_2, \ldots, Y_{n_s}, h\}. \qquad (1)\end{aligned}$$

The components of $\hat{\phi}$ are not linearly independent because the mass fractions sum to unity. There may also be other dependencies related to element and enthalpy conservation. If there are n_1 linear dependencies, then there are

$$D = n_s + 1 - n_1 \qquad (2)$$

degrees of freedom in the thermochemistry. We define the *composition*

$$\phi = \{\phi_1, \phi_2, \ldots, \phi_D\}, \qquad (3)$$

to be a linearly independent subset of $\hat{\phi}$. Given P_0 and a knowledge of the linear dependencies, the thermochemical state of the fluid is completely determined by ϕ.

The composition ϕ of each particle evolves as:

$$\frac{d\phi(t)}{dt} = M(t) + S(\phi[t]) \qquad (4)$$

where S is the rate of change due to chemical reactions, and M is the rate of change due to molecular transport and is referred to as the *mixing* term. Given the chemical kinetics of the system, the components of S are known in terms of ϕ. It is valuable to treat the composition ϕ as a vector; or, equivalently, as a point in the D-dimensional *composition space*.

The composition evolution equation is usually solved by a simple splitting method. For example, the solution is advanced from time t_0 for a small time step Δt by the following procedure:

1. From the initial condition $\phi(t_0)$, the mixing equation

$$\frac{d\phi(t)}{dt} = M(t) \qquad (5)$$

   is integrated for a time Δt, and the solution is denoted by ϕ^0.
2. From the initial condition ϕ^0, the reaction equation

$$\frac{d\phi(t)}{dt} = S(\phi[t]) \qquad (6)$$

   is integrated for a time Δt, to obtain $\tilde{\phi}(t_0 + \Delta t)$—the approximation to $\phi(t_0 + \Delta t)$.

A simple analysis shows that, in the limit as Δt tends to zero, the (local) error on each time-step $|\tilde{\phi}(t_0 + \Delta t) - \phi(t_0 + \Delta t)|$ is $O(\Delta t^2)$ so that the (global) error for a fixed time interval is $O(\Delta t)$, i.e., the method is first-order accurate. In practice, the time-step is selected to be small compared to the mixing timescale (e.g., $\Delta t = \frac{1}{10} \tau_{\text{mix}}$), but this may be several orders of magnitude larger than the smallest chemical time-scale. Although the simple analysis is not valid for such a choice of Δt, numerical tests nevertheless confirm that the method is indeed first-order accurate (Yang and Pope [18]).

ISAT ESSENTIALS

The essential ideas and ingredients of the ISAT approach [8] are as follows:

1. An operator splitting is employed so that reaction problem is reduced to determining the mapping $R(\phi)$, which is the solution to the reaction equation, Eq. (6) above, after a time Δt from the initial condition ϕ.
2. A table is built up in situ, as the reactive flow calculation is performed, so that only the accessed region of the composition space is tabulated. A crucial observation is that the accessed region is a small subset of the realizable region.
3. A table entry (or record) consists of: a composition ϕ^0; the mapping $R(\phi^0)$; the gradient of the mapping $A(\phi^0)$; and the specification of an ellipsoid of accuracy (EOA). The gradient of the mapping is used to obtain the linear approximation

$$R(\phi^q) \approx R(\phi^0) + A(\phi^0)(\phi^q - \phi^0) \tag{7}$$

   to the mapping for a query composition ϕ^q. The EOA is an ellipsoidal region, centered at ϕ^0, within which the linear approximation Eq. (7) is known to be accurate.
4. An ellipsoid of accuracy (EOA) is initialized and grown to ensure that (with high probability) the error involved in Eq. (7) is within a specified tolerance.
5. The records are stored in a binary tree which, given a query composition ϕ^q, can be traversed to obtain a table entry ϕ^0 which in some sense is close to ϕ^q.
6. Given a query ϕ^q, the tree is traversed to the table entry ϕ^0, and if the linear approximation Eq. (7) based on the record ϕ^0 is sufficiently accurate, then this value of $R(\phi^q)$ is returned. Otherwise a new table entry is generated, based on the accurate numerical integration of Eq. (6).
7. As the calculation proceeds, with increasing probability, the query composition ϕ^q lies within the EOA of a table entry ϕ^0, so that the mapping is efficiently retrieved.

ACCURACY

It is not straightforward to conduct the error analysis on the JPDF computations because of the inherent statistical variations between different simulations. Consequently, we rely upon the pairwise mixing stirred reactor (PMSR) test case, which corresponds to a zero-dimensional PDF calculation, similar to the one described by Pope [8] for this purpose.

Test Case: PMSR

The partially stirred reactor (PaSR) model was developed by Correa and Braaten [19] and Chen [20] to study the effects of turbulence-chemistry interactions and to provide a test bed for chemical and mixing schemes for use in particle-based methods. The PMSR (see e.g., [8] and [18]) is designed to overcome the undesirable property of the PaSR that, in steady state, the accessed composition space is a one-dimensional manifold. The PMSR yields a large accessed region which provides a more stringent test for combustion chemistry. At any time t, the PMSR consists of an even number N of particles, the ith particle having composition $\phi^{(i)}(t)$. With Δt being the specified time step, at the discrete times $k\Delta t$ (k integer) events occur corresponding to outflow, inflow, and pairing, which can cause $\phi^{(i)}(t)$ to change discontinuously. Between these discrete times, the composition evolves by a *mixing* fractional step and a reaction fractional step.

The particles are arranged in pairs: particles 1 and 2, 3, and 4, . . . $N - 1$ and N are partners. The mixing fractional step consists of pairs (p and q, say) evolving by

TABLE 1

Parameters used in the PMSR tests[a]

Number of particles	N	100
Time step	Δt	0.1 ms
No. of sub-steps	N_{sub}	5
Substep	Δt_{sub}	0.02 ms
Residence time	τ_{res}	10 ms
Mixing timescale	τ_{mix}	1 ms
Pairing timescale	τ_{pair}	1 ms

[a] The sub-stepping allows the reaction and mixing equation to be solved for a fractional step of 20 μs.

$$\frac{d\phi^{(p)}}{dt} = -(\phi^{(p)} - \phi^{(q)})/\tau_{mix}, \tag{8}$$

where τ_{mix} is the specified mixing timescale.

In the reaction fraction step, each particle evolves by the reaction equation (Eq. (6)).

With τ_{res} being the specified residence time, outflow, and inflow consists of selecting $\frac{1}{2} N\Delta t/\tau_{res}$ pairs at random and replacing their compositions with inflow compositions, which are drawn from a specified distribution. With τ_{pair} being the specified pairing timescale, $\frac{1}{2} N\Delta t/\tau_{pair}$ pairs of particles (other than the inflowing particles) are randomly selected for pairing. Then these particles and the inflowing particles are randomly shuffled so that (most likely) they change partners.

The present PMSR is slightly different form that of Pope [8] in that a specified number of substeps can be taken for reaction and mixing, while pairing and inflow-outflow occur at the specified time-step. This allows us to choose very small substeps Δt_{sub} without affecting the pairing and inflow-outflow in the reactor.

The PMSR test calculations involve the same thermo-chemistry and approximately the same mixing and reaction time-steps (time-substeps here) as those chosen for the JPDF simulation described in the next section. A 16 species-41 reaction skeletal mechanism, described elsewhere [18], was selected for modelling chemistry. Because of element conservation, only 13 of the 16 species are independent, and the 14th composition variable is taken to be enthalpy (the pressure is constant). The values of the parameters used in the test are given in Table 1.

There are three inflowing streams: air (79% N_2, 21% O_2) at 300K; methane at 300K; and a pilot stream consisting of an equilibrium, stoichiometric fuel/air mixture at 2600K. The mass flow rates of these streams are in the ratio 0.85:0.05:0.1. Initially ($t = 0$), all particles are set to the pilot-stream composition and the pressure is atmospheric for all streams.

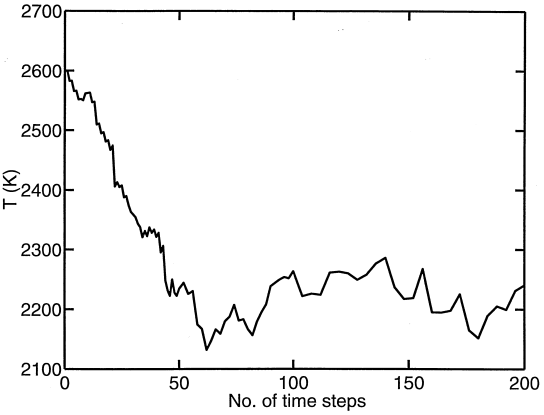

Fig. 1. Mean temperature evolution for the PMSR test.

Figure 1 shows the evolution of mean temperature over 200 time steps (1000 sub time steps) which appears to become statistically stationary around 2200 K. The significant fluctuations are due entirely to the stochastic nature of the PMSR test case: they do not reflect error of any kind. For accuracy testing, the PMSR simulations with different ϵ_{tol} values were performed for 100 time steps and particle properties were compared with those obtained using DI.

Error Control

The local error is the difference between the reaction mapping $R^{DI}(\phi)$ obtained by DI and that using the ISAT technique $R^{ISAT}(\phi)$. This error is controlled satisfactorily [8] by specifying an error tolerance ϵ_{tol}. The objective here is to quantify the global errors, a measure of which is defined below, incurred in computing thermochemistry by user specified ϵ_{tol} for controlling local interpolation errors. Detailed testing for five values of ϵ_{tol} spanning over three decades was performed as follows and the results are presented in Figs. 2, 3, 4a, b.

A measure of average global error (ϵ_G) for species, that lumps the absolute errors incurred for major and minor species together is given by

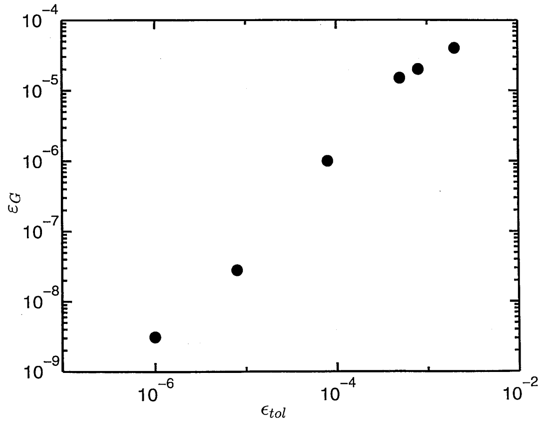

Fig. 2. Global error ϵ_G (Eq. 9) against the specified error tolerance ϵ_{tol} for the PMSR tests.

$$\epsilon_G \equiv \frac{1}{KN} \sum_{k=1}^{K} \sum_{n=1}^{N} |(\phi^{(n)}[k\Delta t] - \phi_{DI}^{(n)}[k\Delta t])|, \tag{9}$$

where $\phi^{(n)}(k\Delta t)$ is the composition (mole fractions) of the nth particle on the kth time step, and $\phi_{DI}^{(n)}(k\Delta t)$ is the corresponding value obtained when the entire calculation is performed using DI. Recall that $N = 100$ is the number of particles and $K = 100$ is the number of time steps. Figure 2 shows that average global errors are well controlled.

We now investigate the global errors for individual species. Figure 3 plots the average error (percentage) in mole fractions for the 16 species, relative to their mean mole fractions, vs the corresponding average mean mole fractions for the 16 species. The points towards the left side of the abscissa are for the minor species. For a given ϵ_{tol}, the minor species incur larger relative global errors than do the major species. For major species (i.e., those with mole fractions in excess of 0.01) a relative error of less than 2% is achieved with $\epsilon_{tol} = 8 \times 10^{-4}$, whereas the same accuracy for all species requires $\epsilon_{tol} = 8 \times 10^{-5}$.

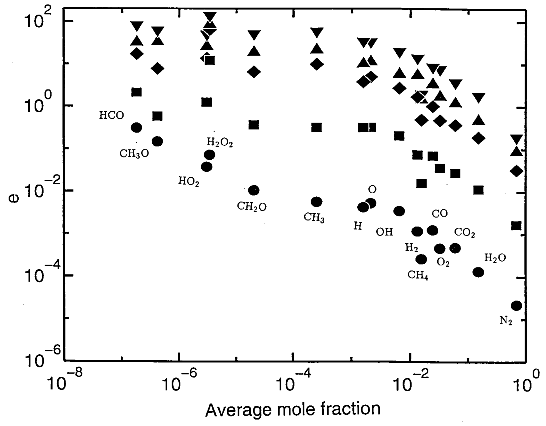

Fig. 3. Average relative error as a percentage (e) incurred for individual species for different ϵ_{tol} values. ϵ_{tol} values are circles, 8.0e-6; squares, 8.0e-5; diamonds, 8.0e-4; triangles, 2.0e-3; inverted triangles, 8.0e-3.

Another way of characterizing accuracy control is by determining the probability of obtaining the compositions to an accuracy within a certain percent of their average values for a typical query. Figure 4a quantifies the probability of ISAT results being within a certain percentage 0.5%–2.0% of DI results at a particle step for a major species CO_2. Similar results for a minor species, OH, are shown in Fig. 4b. We find that the relative errors for individual species (including the minor species) will almost always stay below 2% for $\epsilon_{tol} < 8 \times 10^{-5}$.

The conclusion, therefore, is that one can adequately control the global accuracy of results for species of interest by suitably choosing ϵ_{tol}. It is also observed that the major species exhibit better relative error control. The quantitative results for error control from above are used to select error tolerances for the JPDF simulations for a test flame below. Having fixed the level of desired accuracy, we shall now examine the performance of ISAT in PDF simulations.

ISAT PERFORMANCE IN JPDF CALCULATIONS

Test Case: Piloted Jet Diffusion Flame

Flame L of Masri et al. [21] is chosen as the test case for these simulations. It is a piloted jet diffusion flame of methane-air, for which sufficient data for velocity and thermochemistry is available. The flame is reported to be far from extinction. The simple Langevin model for turbulence was used in the calculations. Our preliminary results [22] revealed that the choice of the mixing model was crucial. The IEM (interaction through exchange with the mean) model failed to perform mixing adequately and the

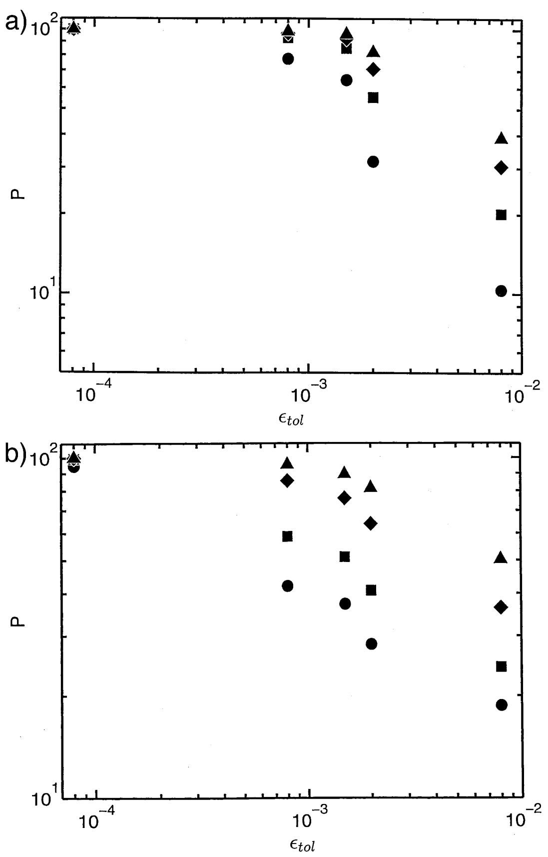

Fig. 4. a: Percentage probability (P) that the relative error incurred for CO2 is less than e percent for different ϵ_{tol} for a typical ISAT mapping. To achieve less than 2% error with certainty, an error tolerance of less than 8×10^{-4} must be chosen. e: circles, 0.5; squares, 1.0; diamonds, 1.5; triangles, 2.0. b: Percentage probability (P) that the relative errors for OH are less than e percent for a typical ISAT mapping. To achieve less than 2% error with certainty, an error tolerance of less than 8×10^{-5} must be chosen. e: circles, 1.0; squares, 2.0; diamonds, 5.0; triangles, 10.0.

flame was extinguished. This is consistent with the IEM results of [23] and [14].

Chen et al. [17] used a modified coalescence-dispersion mixing model for flame-L and found that to avoid the blow-out of the flame, they had to "artificially ignite" the flame by inserting particles with burnt stoichiometric mixture into the fuel jet just downstream of the jet pipe exit. Artificial ignition was also resorted to by Jones et al. [24] for the piloted flame of Godoy [25]. The EMST (Euclidian minimum spanning tree) based model [26], used here, seemed to better incorporate the physics of mixing in our calculations. With the EMST model, no artificial ignition was found to be necessary. The comparison between the two mixing models and detailed comparisons with experimental data will be presented in a follow-up paper. Here, only the performance issues involved in the PDF simulations with detailed chemistry and with EMST mixing model are addressed.

TABLE 2

Parameters used in the JPDF tests[a]

Domain size	X/R, Y/R	60.0, 15.0
Grid discretization	m × n	31 × 31
Max. pSRarticles per cell	N_{pc}	100
Time step	Δt	0.02 ms
Fuel jet bulk speed	U_j	41 m/s
Pilot speed	U_p	24 m/s
Co-flow air speed	U_c	15 m/s
ISAT error tolerance	ϵ_{tol}	0.0008, 0.002, 0.008

[a] R is the fuel jet radius and m and n are the number of grid nodes in axial and radial directions, respectively.

PDF2DV [9] calculates the properties of statistically two-dimensional turbulent reactive flows using the joint velocity-frequency-composition PDF model. It is a mesh based particle code that solves for the evolution of the joint velocity-frequency-composition PDF using a Monte-Carlo technique. The model constants used are the same as in Masri and Pope [10]. Table 2 lists the parameters of these simulations. The skeletal mechanism, described in the PMSR tests, was used again. The goal in this paper is to quantify the performance of the ISAT implementation in PDF2DV for a reasonably fine discretization of the computational domain with a reasonable number of particles.

First, we present results to confirm that the physical features of flame-L are approximately modeled after statistical stationarity of the numerical results is achieved. ISAT performance is addressed subsequently.

Simulation Results

The simulation is allowed to run for long enough so that the flow achieves a statistically stationary state. In Fig. 5a–c, the evolution of the two components of mean velocity, mean mole-fractions of two major and two minor species are plotted at fixed nodes. These time-series for the plotted physical and composition variables confirm that statistical stationarity has

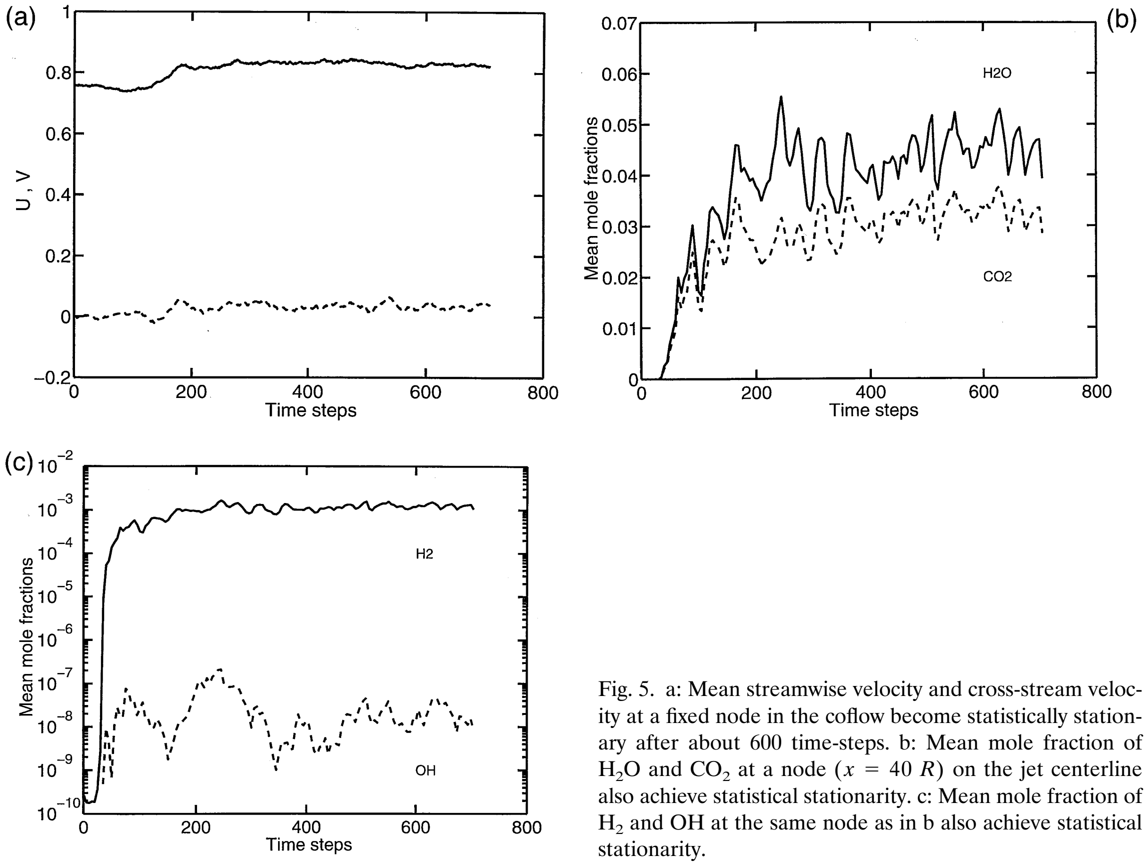

Fig. 5. a: Mean streamwise velocity and cross-stream velocity at a fixed node in the coflow become statistically stationary after about 600 time-steps. b: Mean mole fraction of H_2O and CO_2 at a node ($x = 40\ R$) on the jet centerline also achieve statistical stationarity. c: Mean mole fraction of H_2 and OH at the same node as in b also achieve statistical stationarity.

been achieved at these nodes after about 700 time-steps. Similar time-series, not included here, at other nodes were found to be consistent with these results. During this simulation, the reaction equations were solved over 50 million times, a number we shall refer to as the total number of calls or queries in the ISAT algorithm.

The radial profiles of the Favre-averaged quantities at $x/R = 40$ are plotted in Fig. 6a–d for the ISAT error tolerance ϵ_{tol} of 0.002. Comparisons with experimental data is also made. The results for mean profiles with the other two error tolerances are similar and are not included here. Profiles of Favre-mean and Favre-fluctuating velocities and Favre-averaged mass-fractions of selected major and minor species exhibit consistency with data. The numerical results are not rigorously tested for convergence with respect to the grid-discretization or with respect to the number of particles. The convergence study is currently underway. For the evaluation of ISAT performance, however, the present simulations are satisfactory since the main features of the flame are reproduced well. We now proceed to evaluate the computational efficiency of these ISAT-JPDF simulations.

ISAT Efficiency

By directly integrating the reaction equations half a million times, we estimated the average CPU time for DI chemistry to be 0.05 s per integration. All computations were performed on an SGI Indigo 2 workstation. The efficiency of ISAT in the JPDF simulation described above is investigated by observing the speed-ups achieved over DI and the ISAT-chemistry CPU time as a fraction of the overall CPU time. Table 3 compares the performance of the three ISAT calculations with the DI procedure. Table 4 gives the break-up of the CPU time, for ISAT calculations and a simulation with flamelet

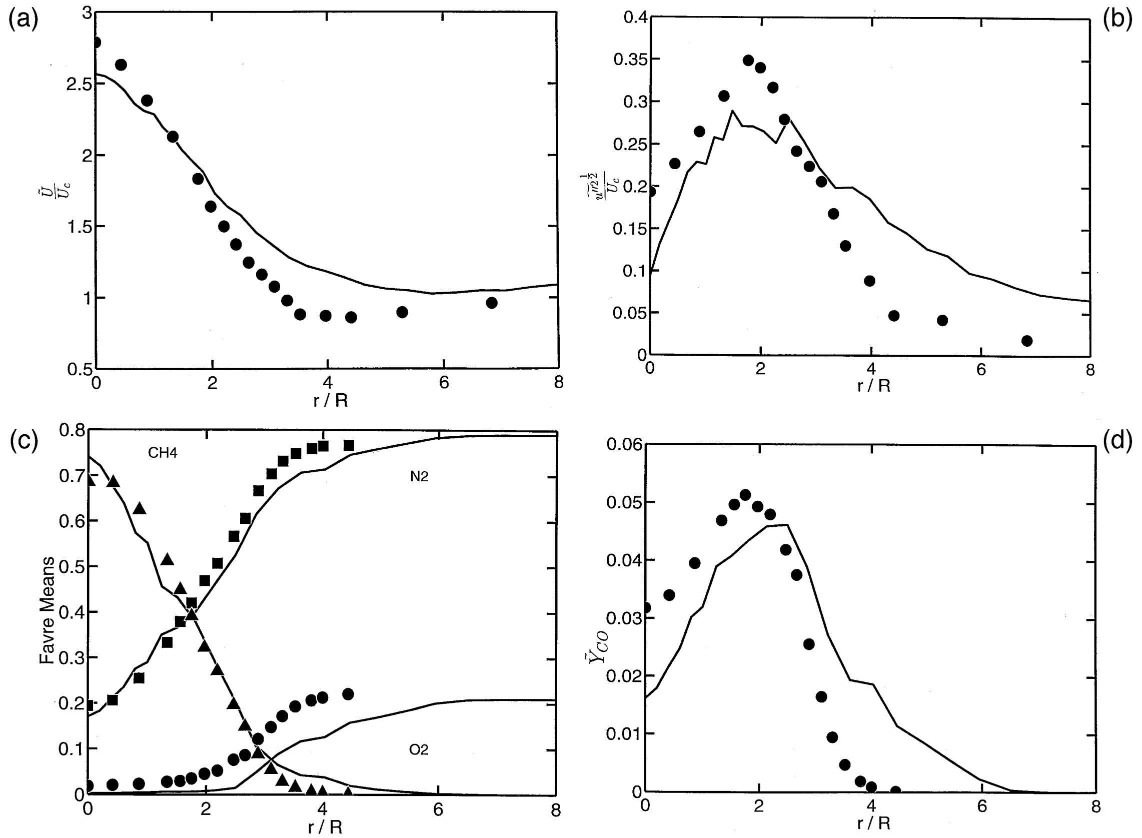

Fig. 6. Radial profiles at $x/R = 40$ of computed (a) Favre-mean streamwise velocity, (b) Favre r.m.s. velocity, (c) Favre-mean mass fractions of three major species (O_2, N_2 and CH_4), and (d) Favre-mean mass fraction of CO are the continuous lines. Data points are from (Masri, Dibble, and Bilger, 1988).

chemistry (with the same grid nodes, particle numbers, and time-step), in mixing, reaction and other calculations. Total simulation time for the most accurate ISAT calculation ($\epsilon_{\text{tol}} = 0.0008$) was 24.85 CPU hours and the storage needed was 50 Megabytes. It is observed that an asymptotic speed-up over DI of about 60 can be achieved. The performance of ISAT is presented graphically in Figures 7a–d.

We report the following further observations

TABLE 3

ISAT performance in the JPDF simulation over the entire simulation and over the last ten time-steps

	$\epsilon_{\text{tol}} = 0.008$	$\epsilon_{\text{tol}} = 0.002$	$\epsilon_{\text{tol}} = 0.0008$	DI
Over 700 time steps				
Total CPU time	14.5 h	18.5 h	24.85 h	774 h
Fraction of time on chemistry	0.3	0.46	0.6	0.98
CPU time per JPDF time step	75 s	95 s	129 s	3980 s
Speed up relative to DI	53	42	31	1
Number of records	790	7130	21975	
Over the last ten time steps				
CPU time per JPDF time step	67 s	75 s	107 s	3980 s
Speed up relative to DI	59	53	37	1
Number of records	2	45	300	

TABLE 4

The break-up of total CPU time in mixing, reaction and other calculations[a]

CPU time (h)	$\epsilon_{\text{tol}} =$ 0.0008	$\epsilon_{\text{tol}} =$ 0.002	$\epsilon_{\text{tol}} =$ 0.008	Flamelet
Mixing	3.01 (12)	3.01 (16)	3.01 (21)	0.26 (25)
Chemistry	14.87 (60)	8.51 (46)	4.5 (31)	0.07 (7)
Rest	6.98 (28)	6.98 (38)	6.98 (46)	0.7 (68)

[a] Flamelet chemistry CPU times are given for comparison. The quantities in parenthesis represent the percentages of the total CPU times.

regarding ISAT efficiency for the three chosen parameters for error control.

- Figure 7a shows that, for chemistry computations alone, the CPU time speed-up over DI is about 200 for the coarsest ϵ_{tol} chosen. For ϵ_{tol} = 0.0008, the speed-up factor is about 50. There is a sudden drop in these curves after about 10,000 calls which corresponds to the rapid growth of the table at this stage as seen in Fig. 8. The speed-up has somewhat flattened but not necessarily reached an asymptote for smaller ISAT error tolerances when the simulation was stopped.
- Overall speed-up of the JPDF simulations is shown in Fig. 7b. Consistent with the fraction of the total CPU time spent in ISAT (Fig. 7c), the maximum speed-up factor for the coarsest error tolerance is about 55. To confirm whether simulations with smaller ϵ_{tol} values are close to reaching asymptotic speed-ups, the performance was analysed for the last 10 time steps as shown in Table 3. We expect that, for larger number of queries, the final speed-ups will approach the asymptotic limit of about 60.
- The ISAT CPU time remains within about 30–60% of the total simulation time.
- A comparison with a JPDF simulation using the flamelet model is shown in Table 4. Note

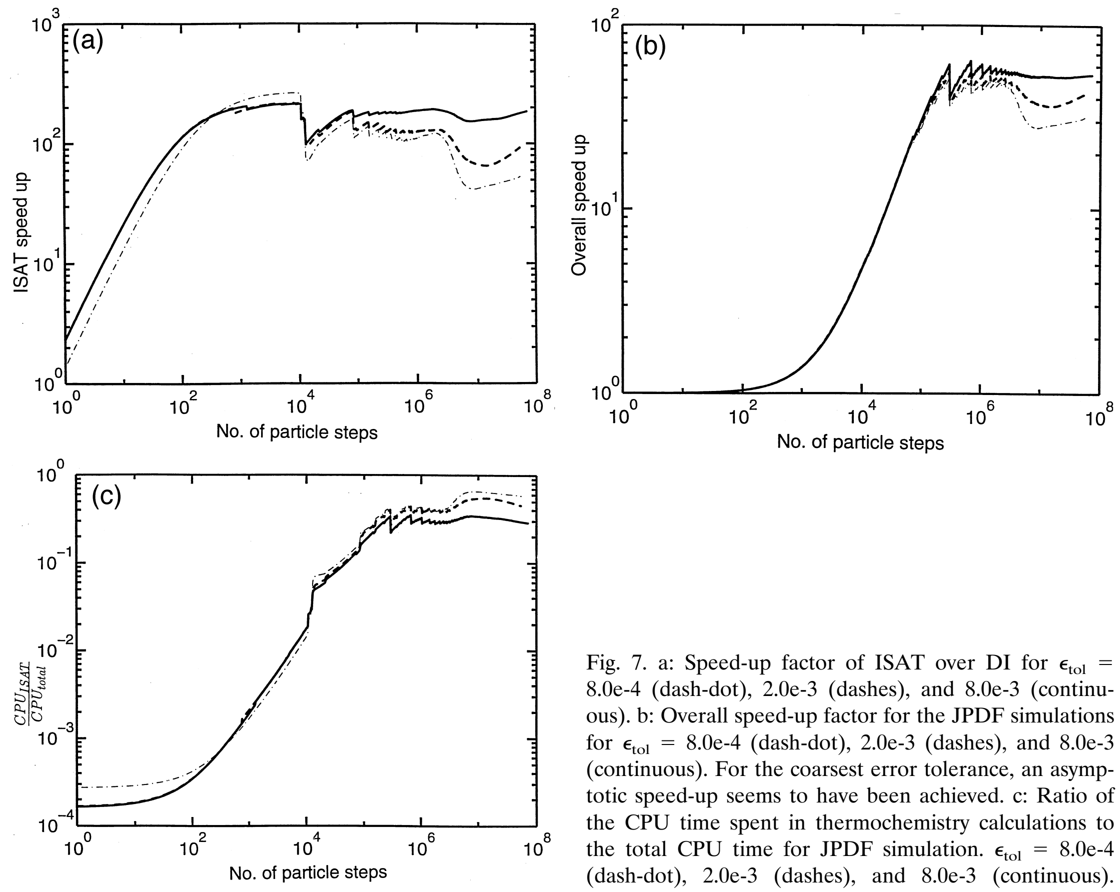

Fig. 7. a: Speed-up factor of ISAT over DI for ϵ_{tol} = 8.0e-4 (dash-dot), 2.0e-3 (dashes), and 8.0e-3 (continuous). b: Overall speed-up factor for the JPDF simulations for ϵ_{tol} = 8.0e-4 (dash-dot), 2.0e-3 (dashes), and 8.0e-3 (continuous). For the coarsest error tolerance, an asymptotic speed-up seems to have been achieved. c: Ratio of the CPU time spent in thermochemistry calculations to the total CPU time for JPDF simulation. ϵ_{tol} = 8.0e-4 (dash-dot), 2.0e-3 (dashes), and 8.0e-3 (continuous). Chemistry takes only 30–60% of the total CPU time.

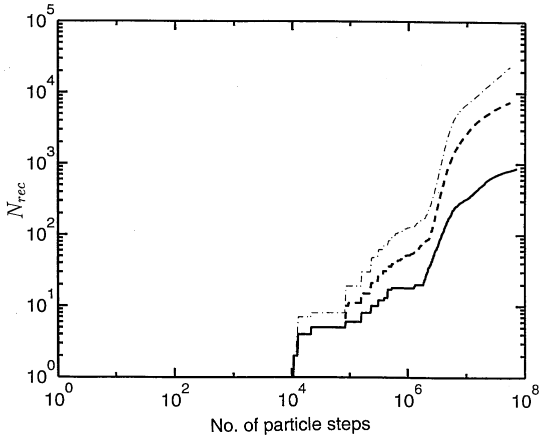

Fig. 8. Table build-up for the JPDF simulations. ϵ_{tol} = 8.0e-4 (dash-dot), 2.0e-3 (dashes), and 8.0e-3 (continuous); 25,000 records for the finest error tolerance correspond to about 50 Mb of memory.

that in the flamelet model, only one scalar (the mixture fraction) is needed.

ISAT Storage Requirements

Up to 10,000 particle steps, the number of records in Fig. 8 is 1 which is sufficient to give the speed up observed in Fig. 7a. This is because all particles in any stream (fuel, pilot, or coflow) have same initial conditions. In the first time step, once the composition of one particle in the fuel stream is advanced, the resulting record is retrieved for the remaining particles in the fuel stream for this time step. We further observe that after about 10,000 initial calls, as the flow evolves, a rapid build up of the in situ table ensues. This growth is then seen to be in small steps before a second spurt. These features are reflected in the speed up curves where rapid table building is accompanied with slowing down of the calculations as reflected by the flattening of and dips in the speed up curves.

The storage requirement ranges from a couple of Megabytes (1000 records) for the coarsest error tolerance to about 50 Megabytes (25,000 records) for the finest one. The storage needed for the intermediate tolerance was approximately 30 Megabytes (8000 records).

The total storage has been shown [8] to scale as D^2, where D is the degrees of freedom in the thermochemistry. It is evident, therefore, that for an accurate tabulation ($\epsilon_{\text{tol}} \rightarrow 0$), with larger degrees of freedom, we need to reduce the dimensionality of the table to allow for smaller storage. Research is currently underway in this direction [27].

CONCLUSIONS

The results presented here demonstrate, for the first time, that routine PDF calculations of realistic flows with detailed chemistry are now possible. The ISAT algorithm allows us to accomplish this formidable task while keeping the errors in chemistry calculations under control.

Overall, for reasonably accurate JPDF simulations, we have demonstrated that speed-ups of up to a factor of 60, compared to the DI procedure for integrating the reaction equations for the entire simulation, are achieved. It is also shown that the storage requirement increases rapidly with increased accuracy desired in the ISAT calculations. Work is underway to reduce the dimensionality of the ISAT table to enable more accurate calculations with lower storage requirements.

The work was supported in part by AFOSR grant F49620-97-1-0126 (program manager Dr. Julian Tishkoff) and in part by the U.S. Department of Energy, Federal Energy Technology Center, Cooperative Agreement no. DE-FC21-92MC29061, through AGTSR (program manager Dr. Daniel Fant).

REFERENCES

1. Pope, S. B., *Prog. Energy Combust. Sci.* 11:119 (1985).
2. Smooke, M. D., *Reduced kinetic mechanisms and asymptotic approximations for methane-air flames*, Lecture notes in Physics 384, Springer, New York (1991).
3. Lam, S. H., and Goussis, D. A., *Proc. 22nd Intl. Symposium on Combustion*, 1988, p. 931.
4. Maas, U. A., and Pope, S. B., *Combust. Flame* 88:239 (1992).
5. Keck, J. C., and Gillespie, D., *Combust. Flame* 17:237 (1971).
6. Chen, J.-Y., Chang, W. C., and Koszykowski, M., *Combust. Sci. Tech.* 110:505 (1995).
7. Taing, S., Masri, A. R., and Pope, S. B., *Combust. Flame* 95:133 (1989).
8. Pope, S. B., *Combustion Theory Modelling* 1:1–24 (1997).
9. Pope, S. B., PDF2DV. *Fortran code*, *Unpublished*, MAE Dept., Cornell University (1994).

10. Masri, A. R., and Pope, S. B., *Combust. Flame* 81: 13–29 (1990).
11. Delarue, B., and Pope, S. B., *Phys. Fluids* 10:487–498 (1998).
12. Dreeben, T., and Pope, S. B., *J. Fluid Mech*. 357:141–166 (1997).
13. Anand, M. S., Hsu, A. T., and Pope, S. B., PDF calculations for swirl combustors. *AIAA J*. 35:1143–1150 (1996).
14. Masri, A. R., Subramaniam, S., and Pope, S. B., *26th Intl. Symposium on Combustion*, 1996, pp. 49–57.
15. Norris, A., and Pope, S. B., *Combust. Flame* 100:211–220 (1996).
16. Bilger, R. W., Turbulent reacting flows. *Topics Applied Phy*. Libby and Williams, Eds.; Vol. 44, Springer, New York (1980).
17. Chen, J. Y., Kollman, W., and Dibble, R. W., *Combust. Sci. Tech*. 64:315–346 (1989).
18. Yang, B., and Pope, S. B., *Combust. Flame* 112:16–32 (1998).
19. Correa, S. M., and Braaten, M. E., *Combust. Flame* 94:469 (1993).
20. Chen, J. Y., *WSS/CI*, *Spring Meeting*, WSS 93-071 (1993).
21. Masri, A. R., Dibble, R. W., and Bilger, R. W., *Combust. Flame* 71:245 (1988).
22. Saxena, V., and Pope, S. B., *ESS/CI Fall meeting 1997*, Hartford, CT (1997).
23. Zurbach, S., Garreton, D., and Kanniche, M., *HE-44/97/025/A*, Direction des etudes et recherches—Electricite de France (1997).
24. Jones, W. P., and Kakhi, M., *Combust. Sci. Tech*. in press (1998).
25. Godoy, S., *Turbulent Diffusion Flames*, Ph.D. thesis, University of London (1982).
26. Subramanian, S., and Pope, S. B., *Combust. Flame*, in press (1998).
27. Pope, S. B., *Cornell University Report*, FDA 96-04 (1996).

Received 23 December 1997; revised 1 May 1998; accepted 9 May 1998