

## The performance of *in situ* adaptive tabulation in computations of turbulent flames

B. J. D. LIU and S. B. POPE\*

Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14853, USA

(Received 30 August 2004; in final form 8 February 2005)

This paper presents a detailed characterization of the local and global errors associated with the *in situ* adaptive tabulation (ISAT) algorithm, which is used in conjunction with a transported PDF method. Calculations of a non-premixed turbulent methane/air piloted jet flame (Sandia flame D) using a skeletal chemical mechanism were performed using ISAT coupled with the computational fluid dynamics (CFD) code FLUENT. The three strategies implemented in ISAT for the growing of the ellipsoids of accuracy (EOAs) are discussed, and the cumulative distribution function (CDF) of the local error is presented for each of the three growing strategies. Computations are also performed to characterize the global error in the ISAT/PDF calculation. The computations used to characterize the global error were performed in parallel to achieve substantial savings in computational time.

In general the local error is well controlled, but there is a small probability of relatively large errors. Results from the investigation suggest that large retrieve errors are due to the region of accuracy (ROA) being non-convex, where the ROA is the connected region for which the error does not exceed the error tolerance,  $\epsilon_{\text{tol}}$ . The global error in ISAT is found to be small compared to statistical error for  $\epsilon_{\text{tol}} \leq 10^{-4}$ , and is found to vary linearly with  $\epsilon_{\text{tol}}$ .

**Keywords:** Turbulent combustion; PDF methods; ISAT algorithm

### 1. Introduction

A major hurdle to accurate calculations of complex turbulent reactive flows is that typical detailed combustion mechanisms involve tens or hundreds of species, hundreds or thousands of reactions, and a wide range of timescales. The system of ordinary differential equations governing chemical reactions is thus large and extremely stiff, making the task of solving these equations computationally expensive. The *in situ* adaptive tabulation (ISAT) algorithm, introduced by Pope (1997) [13], has been shown to speed up these chemistry calculations by up to a factor of 1000 [13], and can be used in the context of PDF methods for turbulent combustion, as well as in other approaches (e.g. finite-difference methods for laminar flames [15]).

Previously, Xu and Pope (2000) [18] and Tang *et al.* (2000) [17] have used ISAT in conjunction with PDF methods in the computation of turbulent non-premixed piloted jet flames. ISAT has also been used with artificial neural networks (ANN) and LES by Kapoor *et al.*

---

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

(2001) [8], and has also been successfully combined with dimension reduction strategies (e.g. quasi-steady-state assumption, rate-controlled constrained equilibrium, etc.) [16].

While the performance gains from using ISAT are well demonstrated, a comprehensive study of the accuracy of the error associated with the storage and retrieval algorithm in ISAT has not previously been performed. Also, apart from an investigation by Singer and Pope [15] for a laminar premixed flame test case, there has been no study of the global error in ISAT to date. This paper addresses these issues by providing a characterization of the local and global errors associated with ISAT using, as a test case, a turbulent flame calculation by the composition PDF method with a 16-species skeletal mechanism.

The experimental reacting flow chosen for this study is a piloted jet methane/air flame (flame D) measured by Barlow and Frank [1]. The next section provides a description of this test flame, and is followed by a discussion of the local (retrieve) error in ISAT, and by the presentation of computed cumulative distribution functions (CDFs) of local error. A characterization of the global error in ISAT is then given, followed by some observations and conclusions.

## 2. Description of the test flame

This section briefly summarizes the experimental conditions [1] of the piloted jet test flame used for the calculations in this paper. The fuel jet with radius  $R_j = 3.6$  mm is accompanied by an annular pilot with radius  $R_p = 9.2$  mm. The jet fuel is 25% CH<sub>4</sub> and 75% dry air by volume, and the pilot burns a lean premixture of C<sub>2</sub>H<sub>2</sub>, H<sub>2</sub>, air, CO<sub>2</sub>, and N<sub>2</sub> with the same nominal enthalpy and equilibrium composition as methane/air at an equivalence ratio of 0.77. The bulk velocity of the fuel jet is 49.6 m s<sup>-1</sup>, and the jet Reynolds number is 22,400. The jet and pilot are surrounded by a slow co-flow of air, with a free-stream velocity of 0.9 m s<sup>-1</sup>. Detailed information about the flame can be obtained from Barlow and Frank (2003) [2]. PDF calculations of this flame have been performed by Xu and Pope (2000) [18], Tang *et al.* (2000) [17], and Lindstedt and Louloudi (2000) [10], among others.

## 3. Computational specifications

All computations were performed using the FLUENT CFD code, which solves the Reynolds averaged Navier–Stokes (RANS) equations for the mean conservation of mass, momentum and energy, using the  $k - \varepsilon$  model for turbulence modelling. The model constants used in the  $k - \varepsilon$  model are given in table 1. A transport equation for the composition PDF is coupled and solved using a Lagrangian particle-based Monte Carlo method. The modified Curl (MC) mixing model [7] is used with  $C_\phi = 2.0$ , and the chemical kinetics are treated using ISAT with a skeletal mechanism for methane consisting of 16 species and 41 reactions (the same as that used in [14]). The wall functions used in this calculation are based on the proposal of Launder and Spalding [9]. No radiation effects are considered, and the effect of gravity is neglected.

The solution domain, shown in figure 1, is an axisymmetric, 2-D domain. The origin is placed at the centre of the jet exit plane. The axial direction is denoted by  $x$ , while  $y$  and  $r$  are used interchangeably to denote the radial direction. The solution domain extends 5 jet

Table 1. Model constants used in the  $k - \varepsilon$  model.

$C_\mu$	$C_{\varepsilon 1}$	$C_{\varepsilon 2}$	$\sigma_k$	$\sigma_\varepsilon$	$\sigma_\Phi$
0.09	1.52	1.92	1.0	1.3	0.7




Figure 1. Sketch of the solution domain.

diameters upstream and 100 jet diameters downstream ( $x = 720$  mm) in the axial direction; and 50 jet diameters from the centerline in the radial direction ( $y = 360$  mm). The wall thickness of the pilot flame is treated as infinitesimally thin. The grid is comprised of 2352 cells with non-uniform grid spacing and some cells made non-orthogonal. More detailed information about the grid is given in table 2.

Boundary conditions are given in table 3. The compositional inflow boundary condition in the pilot is determined by matching the measurements at  $x/d = 1$  with calculations of laminar unstrained premixed  $\text{CH}_4/\text{air}$  flames and then extrapolating to the conditions at the burner exit plane, based on the estimated convective time up to  $x/d = 1$ . The pilot burnt gas velocity is determined from the cold mass flow rate, the density at the estimated exit condition, and the flow area of the pilot annulus. The pilot composition is thus taken as that of an unstrained  $\text{CH}_4/\text{air}$  premixed  $\phi = 0.77$  flame at the point in the flame profile where the temperature is 1880 K, following the process outlined above [2]. In the computation, the inlet boundary velocity profiles for the jet and coflow are specified. The pilot inlet velocity profile is assumed to be flat, except for thin boundary layers. Since this inlet velocity profile is obtained directly from experimental data, it eliminates inaccurate boundary conditions as a source of error in our computations. In table 3, the turbulence length scale,  $L$ , is such that the dissipation  $\varepsilon$  is

Table 2. Detailed information of the grid used in the calculation.

	x			y		
	From (mm)	To (mm)	Cells	From (mm)	To (mm)	Cells
Fuel jet	-36.0	0	8	0	3.6	6
Pilot	-36.0	0	8	3.6	9.1	9
Coflow	-36.0	720	49	9.1	360	33

Table 3. Boundary conditions for the calculations.

Stream	Condition	Value
Fuel jet	Bulk velocity (m s <sup>-1</sup> )	49.6
	Temperature (K)	294
	Turbulence intensity (%)	4
	Turbulence lengthscale, <i>L</i> (m)	0.0002
	<i>Y</i> <sub>CO<sub>2</sub></sub>	0.00046
	<i>Y</i> <sub>O<sub>2</sub></sub>	0.1962
	<i>Y</i> <sub>CH<sub>4</sub></sub>	0.15605
	<i>Y</i> <sub>N<sub>2</sub></sub>	0.6473
Pilot	Velocity (m s <sup>-1</sup> )	11.4
	Temperature (K)	1880
	Turbulence intensity (%)	1
	Turbulence lengthscale, <i>L</i> (m)	0.0001
	<i>Y</i> <sub>H<sub>2</sub>O</sub>	0.0942
	<i>Y</i> <sub>CO<sub>2</sub></sub>	0.1098
	<i>Y</i> <sub>O<sub>2</sub></sub>	0.054
	<i>Y</i> <sub>CO</sub>	0.00407
	<i>Y</i> <sub>H</sub>	2.48 × 10 <sup>-5</sup>
	<i>Y</i> <sub>H<sub>2</sub></sub>	0.000129
	<i>Y</i> <sub>N<sub>2</sub></sub>	0.7378
Coflow	Velocity (m s <sup>-1</sup> )	0.9
	Temperature (K)	291
	Turbulence intensity (%)	1
	Turbulence lengthscale, <i>L</i> (m)	0.001
	<i>Y</i> <sub>H<sub>2</sub>O</sub>	0.00581
	<i>Y</i> <sub>O<sub>2</sub></sub>	0.23113
	<i>Y</i> <sub>CO<sub>2</sub></sub>	0.00055
<i>Y</i> <sub>N<sub>2</sub></sub>	0.7625	
Jet wall	Adiabatic wall (zero heat flux)	
Pilot wall	Adiabatic wall	
Outer boundary	Symmetry	

determined as  $\varepsilon = ck^{3/2}/L$ , where  $k$  is the turbulent kinetic energy and  $c$  is equal to  $0.09^{3/4}$ .  $Y_i$  denotes the mass fraction of species  $i$ .

The numerical conditions selected are as follows: there are nominally 40 particles per cell ( $N_{pc}$ ), and approximately 100,000 particles in total. The ordinary differential equation (ODE) solver DDASAC [4] is used to perform direct integrations where necessary. The absolute and relative error tolerances in DDASAC are set to  $10^{-6}$  and  $10^{-9}$ , respectively, for all variables. The Courant number used for local time-stepping is set to 0.5 for convection, diffusion, and mixing.

Discretization is achieved via an implicit, segregated finite-volume scheme which solves the governing equations sequentially. The values of momentum, turbulent kinetic energy, and turbulent dissipation rate at cell faces are calculated using a multidimensional linear reconstruction approach [3], yielding second-order accuracy at cell faces through a Taylor series expansion of the cell-centered solution about the cell centroid. This scheme is also known as the second-order upwind scheme.

The Semi-Implicit Method for Pressure-Linked Equations (SIMPLE) family of algorithms [12] is used for pressure-velocity coupling (i.e. introducing pressure into the continuity equation). The Pressure Staggering Option (PRESTO!) pressure interpolation scheme is used to obtain the value of pressure at the cell faces from the cell values via a discrete continuity balance for a control volume about the face. This procedure is similar in spirit to the staggered-grid schemes used with structured meshes [12].

The iterative process operates as follows: a FLUENT finite-volume iteration is performed, followed by a particle iteration in which the particle properties are advanced for one pseudo-time step [6]. The base case corresponds to a computation in which iterative convergence is obtained. This fact is ascertained by observing that quantities at monitoring locations become statistically stationary. Further details about the finite-volume calculations are given by Liu [11].

#### 4. Brief overview of ISAT

The *In Situ* Adaptive Tabulation (ISAT) algorithm introduced by Pope (1997) [13] is a storage and retrieval methodology. This section outlines essential ideas about the ISAT algorithm.

Consider an isobaric reactive gaseous flow (a good assumption for most low-Mach number flames) with  $n_s$  species. For a given pressure  $P$ , the thermochemical composition vector  $\phi$  is given by

$$\phi = \{Y_1, Y_2, \dots, Y_{n_s}, h\}, \quad (1)$$

where  $Y_k$  is the mass fraction of species  $k$ , and  $h$  is the specific enthalpy:  $\phi$  is a vector of length  $n_\phi \equiv n_s + 1$ .

In the present computational implementation, the fluid within the solution domain is represented by a large number of particles, whose compositions evolve due to reaction and mixing. These two processes are treated in separate fractional steps. In the reaction fractional step, the composition evolves due to chemical reactions at fixed pressure and enthalpy according to the ordinary differential equation

$$\frac{d\phi(t)}{dt} = \mathbf{S}(\phi[t]). \quad (2)$$

Here  $S_k$  is the net creation rate of species  $k$  (for  $k = 1, 2, \dots, n_s$ ), and  $S_{n_\phi} \equiv 0$ , corresponding to constant enthalpy. From the initial condition  $\phi^0$  at time  $t_0$ , (2) is integrated for a time  $\Delta t$  to obtain  $\phi(t_0 + \Delta t)$ . For fixed  $\Delta t$ , the *reaction mapping*

$$\mathbf{R}(\phi^0) \equiv \phi(t_0 + \Delta t), \quad (3)$$

is uniquely determined by  $\phi^0$ . ISAT uses the ODE solver DDASAC to integrate (2) and stores the reaction mappings in a binary tree, with each termination node (or leaf) representing a record consisting of (among other information) the tabulation point  $\phi^0$ , its associated reaction mapping  $\mathbf{R}(\phi^0)$ , and the *mapping gradient*  $\mathbf{A}(\phi^0)$ , where

$$A_{ij}(\phi^0) \equiv \frac{\partial R_i(\phi^0)}{\partial \phi_j^0}. \quad (4)$$

The *region of accuracy* (ROA) is defined to be the connected region containing  $\phi^0$  consisting of points  $\phi$  for which the local error  $\varepsilon$  does not exceed the specified tolerance  $\varepsilon_{\text{tol}}$ . At each leaf, ISAT approximates the ROA as a hyperellipsoid, known as the *ellipsoid of accuracy* (EOA) centered at  $\phi^0$ . Given a query composition  $\phi^q$ , ISAT traverses the tree until a leaf representing some  $\phi^0$  is reached, and this value of  $\phi^0$  is close to  $\phi^q$ . Now if  $\phi^q$  is determined to be within the EOA for that leaf, ISAT performs a *retrieve*, and returns a linear approximation to  $\mathbf{R}(\phi^q)$ . This linear approximation is denoted  $\mathbf{R}^\ell(\phi^q)$ , which is defined as

$$\mathbf{R}^\ell(\phi^q) \equiv \mathbf{R}(\phi^0) + \mathbf{A}(\phi^0)(\phi^q - \phi^0). \quad (5)$$

The *local error*  $\varepsilon$  is then simply defined as the scaled difference between the exact mapping and the linearized estimate

$$\varepsilon \equiv |\mathbf{B}(\mathbf{R}[\phi^q] - \mathbf{R}^\ell[\phi^q])|, \quad (6)$$

where  $\mathbf{B}$  is a scaling matrix. For a more detailed description of the ISAT algorithm, the reader is directed to Pope (1997) [13].

## 5. Local error

### 5.1 Introduction

As described in the original ISAT paper by Pope (1997) [13], given a query point  $\phi^q$  falling within an existing ellipsoid of accuracy (EOA), the ISAT algorithm uses a piecewise linear approximation within the EOA to calculate a linear approximation of the reaction mapping. The (scaled) difference between the exact mapping,  $\mathbf{R}(\phi^q)$ , and this linearized estimate,  $\mathbf{R}^\ell(\phi^q)$ , is known as the retrieve, or *local error* [13]. Effective control of the local error ensures that the ISAT part of the turbulent reactive flow calculation is numerically accurate (for sufficiently small values of the ISAT error tolerance,  $\varepsilon_{\text{tol}}$ ).

As described in Pope (1997) [13], if a query point  $\phi^q$  does not fall within an existing EOA, but has error  $\varepsilon$  less than  $\varepsilon_{\text{tol}}$ , then the existing EOA is grown. The computational expense incurred during the growth of an EOA is offset by the increased probability of retrieving from that (larger) EOA during future evaluations of (2). For a computer with unlimited memory, the retrieve rate would be close to 100% once the (very large) ISAT table is built. In reality, however, computer systems have limited memory in which to store the ISAT table, and growing the EOAs is a cost-effective method to increase the percentage of retrieves for a given ISAT table size.

It is hypothesized that the growth of the EOAs is one of the major causes of rare large local errors (reported below). Accordingly, we investigate three different strategies for growing of EOAs. They are (in ascending order of conservativeness, with mode 3 the most conservative): Mode 1 (ellipsoidal growing); Mode 2 (ellipsoidal with Chew modification); and Mode 3 (conical growing). This section describes each of these growing strategies, then gives a discussion of the possible causes of local error in ISAT. Subsequently, a characterization of the local error for each of these modes is given to assess their relative merits. (It is possible to suppress the growing of the EOAs by using Mode 0, which does not allow the EOAs to grow at all.)

### 5.2 Growing strategies

**5.2.1 Grow mode 1: ellipsoidal growing.** The ellipsoidal growing strategy used in mode 1 is the growing strategy proposed by Pope in the original ISAT paper [13]. The growing strategy is illustrated in figure 2. Given an existing EOA centered at a point  $\phi^0$ , the new EOA is the unique hyper-ellipsoid of minimum volume, centered at  $\phi^0$ , which encloses both the original EOA and the point  $\phi^q$ .

The growing process can be understood more clearly by considering the linear transformation of the original EOA to a unit hypersphere. In this transformed space, the new EOA is the hyper-ellipsoid of minimum volume that encloses the unit hypersphere and the grow point. This is shown schematically in figure 3.




Figure 2. Growth of an ellipsoid of accuracy using grow mode 1.

One property of mode 1 is that, given a knowledge of the EOA before growth and the grow point  $\phi^q$ , the EOA after growth (for mode 1) is the ellipsoid of smallest volume which is sure to bound all previous grow points.

**5.2.2 Grow mode 2: ellipsoidal with Chew modification.** Grow mode 2, also called the *Chew modification*, is essentially a slightly more conservative variant of grow mode 1. This growing strategy was proposed by Chew [5]. For each EOA, the hyper-ellipsoid is first transformed into a unit hypersphere, as shown in figure 3. A hypercube is now inscribed inside the unit hypersphere, with its vertices on the surface of the hypersphere. When the EOA is grown, the new EOA (in transformed space) is the unique hyper-ellipsoid of minimum volume enclosing both the hypercube and the grow point. Modes 1 and 2 are illustrated in figure 4 for a two-dimensional case.

It is important to appreciate that while the EOA before growth is known, the sequence of previous grow points is considered to be unknown. Thus there is a sequence of previous grow points which yield the EOA before growth, and which, with mode 1, are bounded by the EOA after growth. As we can see from figure 4, mode 2 is a more conservative growing strategy than mode 1 as the grown EOA using mode 2 is slightly smaller than the grown EOA using mode

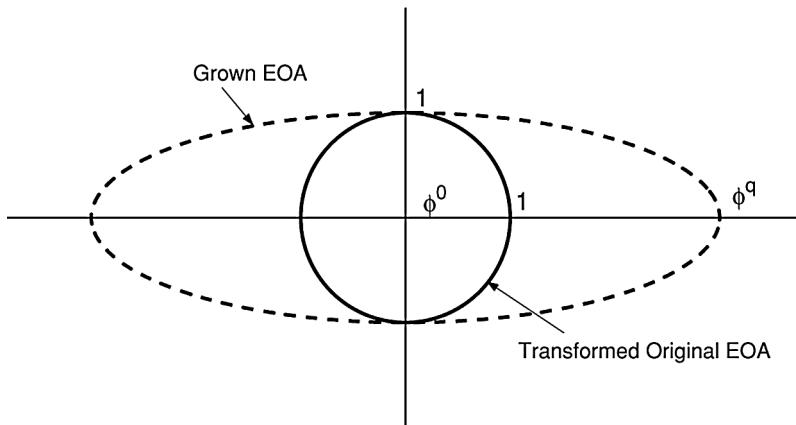


Figure 3. Growth of an ellipsoid of accuracy using grow mode 1 in transformed space.




Figure 4. Growth of an ellipsoid of accuracy using grow modes 1 and 2 in two-dimensional transformed space. Solid lines: mode 1; dotted lines: hypercube used in mode 2; dashed lines: grown EOA using mode 2.

1, and therefore the ellipsoid after growth under mode 2 may not include some previous grow points (at which  $\varepsilon < \varepsilon_{tol}$ ). However, under mode 2, the EOA after growth has the property that it is the ellipsoid of smallest volume that bounds the current grow point  $\phi^q$ , and that could bound all previous grow points. The crucial point here is that all the previous grow points are unknown, and hence an EOA grown in this way is one of smallest volume that has a possibility of bounding these old grow points. If the previous grow points are known, then it would be possible to bound all these old grow points with an ellipsoid of volume smaller than that generated by mode 1, but larger than that generated by mode 2.

**5.2.3 Grow mode 3: conical growing.** In grow mode 3, the EOAs are grown using a cone based on both the current EOA and the grow point. Figure 5 shows a sketch of the EOA transformed to the unit hypersphere, with the grow point  $\phi^q$  located on the axis at distance  $r$  away from the origin. The error at the grow point is  $\varepsilon_q$ , and (to leading order) the error increases as the square of the distance from the origin. The limit point  $\phi^p$  shown in the sketch is the location on the axis where the error is estimated to be  $\varepsilon_{tol}$ , i.e.  $\phi^p = \phi^q (\varepsilon_{tol} / \varepsilon_q)^{1/2}$ .




Figure 5. Sketch showing conical growing strategy in transformed space (grow mode 3). Solid line, original EOA; dotted line, cone used in mode 3; dashed line, grown EOA using mode 3.



A cone is created (in transformed space) with its vertex at the limit point  $\phi^p$ , as shown in figure 5. The new EOA is then the hyper-ellipsoid of maximum volume that can be inscribed within the cone, which may not include the grow point. It is possible that the new EOA has a smaller volume than the original EOA. In this case the original EOA is retained, and so no growing is performed.

This growing strategy ensures that, for the case in which the ROA is convex, the EOA after growth lies entirely within the ROA. The conical growing strategy is the most conservative of all the grow modes, and the only one with this property for convex ROAs.

The three growing strategies described above present increasingly conservative estimates for the grown EOA, with the aim of reducing local error. Before the local error associated with each of these grow modes is characterized, we consider the possible reasons for inaccuracy due to growing of the EOAs.

### 5.3 Possible reasons for inaccuracy due to growing

The region of accuracy (ROA) is defined to be the connected region containing the tabulation point,  $\phi^0$ , consisting of points  $\phi$  for which the local error  $\varepsilon$  does not exceed the specified tolerance  $\varepsilon_{\text{tol}}$ . There are three possible reasons as to why inaccuracies arise due to growing of the EOAs.

Firstly, it is possible that, over the ROA, the tabulated function is significantly non-linear, and hence the Taylor series analysis may be inaccurate, giving rise to the possibility of non-monotonic behavior of  $\varepsilon$ . If this is the case, then, given a query point  $\phi^q$ , there might exist sections on the line segment between  $\phi^0$  and  $\phi^q$  where  $\varepsilon$  is greater than  $\varepsilon_{\text{tol}}$ , even though the grow point itself has  $\varepsilon$  less than  $\varepsilon_{\text{tol}}$ . Nevertheless, in such a situation, the EOA is grown to encompass these inaccurate regions, introducing the possibility that subsequent retrieves will be inaccurate. This problem, if it exists, vanishes as  $\varepsilon_{\text{tol}}$  becomes sufficiently small, because the ROA shrinks and the accuracy of the Taylor series increases within the ROA.

Secondly, the ROA may not be convex. In this case, all growing strategies can cause the EOA after growing to include inaccurate regions. To elaborate on the geometry of the ROA, it is readily shown [13] that if the piecewise-constant approximation  $\mathbf{R}(\phi^q) \approx \mathbf{R}(\phi^0)$  is used, then the ROA (for small  $\varepsilon_{\text{tol}}$ ) is a hyper-ellipsoid, whose principal axes are given by the eigendecomposition of  $\varepsilon_{\text{tol}}(\mathbf{A}^T \mathbf{A})^{-1/2}$ . For the case of the piecewise linear approximation, the analysis is more involved, and reveals that the ROA can have a hyperbolic (i.e. non-convex) nature.

Thirdly, even if the ROA is convex (e.g. ellipsoidal), the ellipsoidal growing strategies (modes 1 and 2) can lead to the inclusion of inaccurate regions.

### 5.4 Computational determination of local error

The aim of the computations described in this section is to characterize the local error  $\varepsilon$  for a large number of retrieves. This is done by estimating the cumulative distribution function (CDF) of the local error

$$F(x) \equiv \text{Prob}\{\varepsilon < x\}, \quad (7)$$

where  $x$  is the sample-space variable corresponding to  $\varepsilon$ . The CDF  $F(x)$  provides a complete characterization of the local error.

The protocol to determine  $F(x)$  is as follows: the PDF calculation described in section 3 is run for 600 iterations starting from a statistically stationary solution and an empty ISAT table. Over these initial 600 iterations, the ISAT table is built and filled, i.e. the maximum allocated

Table 4. ISAT error statistics for different growing strategies,  $\varepsilon_{\text{tol}} = 1 \times 10^{-4}$ .

Grow mode	1	2	3	No growing
$\varepsilon_{\text{ref}} \times 10^6$	342	303	67.8	—
$\varepsilon_{0.99}/\varepsilon_{\text{ref}}$	6.93	6.77	5.94	—
$\varepsilon_{0.999}/\varepsilon_{\text{ref}}$	19.77	18.65	15.04	—
Prob $\{\varepsilon > \varepsilon_{\text{tol}}\}$ (percentage)	20.02	19.89	7.14	0.0
$\varepsilon_{0.9}/\varepsilon_{\text{tol}}$	3.42	3.42	0.678	—
$\varepsilon_{0.99}/\varepsilon_{\text{tol}}$	23.7	20.5	4.03	0.016
$\varepsilon_{0.999}/\varepsilon_{\text{tol}}$	67.6	56.5	10.2	0.027

space (400 megabytes) for the ISAT table has been occupied, and hence no more adding to the table can take place. Five iterations are subsequently performed, during which, for each query point  $\phi^q$ , the ISAT method is used to determine the mapping,  $\mathbf{R}^\ell(\phi^q)$ , and in addition the exact result  $\mathbf{R}(\phi^q)$  is obtained by direct integration so that the local error  $\varepsilon$  can be measured directly. From these calculations,  $F(x)$  is constructed based on the samples of  $\varepsilon$ . This process is repeated for each grow mode for two values of  $\varepsilon_{\text{tol}}$ :  $1 \times 10^{-4}$  and  $(1/64) \times 10^{-4}$ .

In these calculations, the absolute and relative error tolerances in DDASAC are set to  $10^{-6}$  and  $10^{-9}$ , respectively, for all variables, and one binary tree is used in ISAT. As described previously, there are (nominally) 40 particles per cell, with about 100,000 particles in total, and hence 5 iterations yields half a million samples of the local error,  $\varepsilon$ . All computations reported in this section are performed in serial.

The quantity  $\varepsilon_\alpha$  is defined as

$$F(\varepsilon_\alpha) = \text{Prob}\{\varepsilon < \varepsilon_\alpha\} = \alpha, \quad \text{for } 0 < \alpha < 1, \quad (8)$$

i.e. with probability  $\alpha$ , the error is less than  $\varepsilon_\alpha$ . From  $F(x)$ , various statistics can be extracted, such as  $\text{Prob}\{\varepsilon > \varepsilon_{\text{tol}}\}$ , as well as the error statistics just defined. These statistics are normalized using both  $\varepsilon_{\text{tol}}$  and  $\varepsilon_{\text{ref}} \equiv \varepsilon_{0.9}$ . Normalization by  $\varepsilon_{\text{ref}}$  allows a comparison of the shapes of the CDFs obtained with different values of  $\varepsilon_{\text{tol}}$  and different growing strategies. Local error statistics normalized by both  $\varepsilon_{\text{tol}}$  and  $\varepsilon_{\text{ref}}$  are shown in section 5.5.

### 5.5 Measurement of local error

Tables 4 and 5 show the tabulated values of ISAT error statistics for the different grow modes with  $\varepsilon_{\text{tol}} = 1 \times 10^{-4}$  and  $\varepsilon_{\text{tol}} = (1/64) \times 10^{-4}$ . Also included are statistics corresponding to the case where no growing of the EOAs is performed. Various statistics are shown normalized by  $\varepsilon_{\text{tol}}$  and  $\varepsilon_{\text{ref}}$  (defined in section 5.4). Note that in tables 4 and 5, the data generated from

Table 5. ISAT error statistics as a variation of grow mode,  $\varepsilon_{\text{tol}} = (1/64) \times 10^{-4}$ .

Grow mode	1	2	3	No growing
$\varepsilon_{\text{ref}} \times 10^9$	9090	7000	872	—
$\varepsilon_{0.99}/\varepsilon_{\text{ref}}$	6.63	6.99	9.94	—
$\varepsilon_{0.999}/\varepsilon_{\text{ref}}$	17.82	18.29	52.52	—
Prob $\{\varepsilon > \varepsilon_{\text{tol}}\}$ (percentage)	23.36	21.64	6.76	0.0
$\varepsilon_{0.9}/\varepsilon_{\text{tol}}$	5.82	4.48	0.56	—
$\varepsilon_{0.99}/\varepsilon_{\text{tol}}$	38.6	31.3	5.55	—
$\varepsilon_{0.999}/\varepsilon_{\text{tol}}$	103.7	81.92	29.3	$2.6 \times 10^{-3}$

the calculations do not give a value of  $\varepsilon_{\text{ref}}$  for the case where growing is disallowed. This is because when no growing is allowed, there is a negligible number of samples with errors greater than  $\varepsilon_{\text{tol}}$ , and the CDF of  $\varepsilon$  does not have a tail from which the value of  $\varepsilon_{\text{ref}}$  could be determined. Hence, the normalized data for this case are unavailable.

The statistics of  $\varepsilon$  normalized by  $\varepsilon_{\text{tol}}$  show that grow modes 1 and 2 give similar results, with mode 2 being slightly more conservative than mode 1 as expected. Mode 3 gives substantially better error control (with respect to error tolerance). As expected, the absence of growing provides the best error control of all.

The plots of  $1 - F(x)$  against  $\varepsilon/\varepsilon_{\text{ref}}$  for different grow modes for  $\varepsilon_{\text{tol}} = 1 \times 10^{-4}$  and  $\varepsilon_{\text{tol}} = (1/64) \times 10^{-4}$  are shown in figure 6. The quantity  $1 - F(x) = \text{Prob}\{\varepsilon \geq x\}$  is shown to focus on the larger errors. Figure 7 shows the effect of  $\varepsilon_{\text{tol}}$  on the shape of  $1 - F(x)$ : each grow mode is plotted separately for the two values of  $\varepsilon_{\text{tol}}$ . For figures 6 and 7, note that the plots appear to show that  $1 - F(x)$  does not equal 1 for small values of error. This is due to the fact that the plot of the CDF shows samples which have a value of  $\varepsilon$  in the range between  $\varepsilon_{\text{tol}}/10^3$  and  $\varepsilon_{\text{tol}} \times 10^3$ , but about half of the samples have values of  $\varepsilon$  less than  $\varepsilon_{\text{tol}}/10^3$ .

From figure 6, we see that long tails exist for the plot of  $1 - F(x)$  associated with all three grow modes. This clearly shows that none of the grow modes exhibits excellent error control associated with a sharp cutoff in the shape of  $1 - F(x)$  at the specified error tolerance. Also, examination of figure 7 shows that  $\varepsilon_{\text{tol}}$  does not have a significant effect on the behaviour of the tail of  $1 - F(x)$  if grow modes 1 or 2 are used; the difference between the two curves is possibly within statistical variability. However, for mode 3, the finer error tolerance




Figure 6. CDF of local error for different grow modes against  $\varepsilon/\varepsilon_{\text{ref}}$ . Left plot,  $\varepsilon_{\text{tol}} = 1 \times 10^{-4}$ ; right plot,  $\varepsilon_{\text{tol}} = (1/64) \times 10^{-4}$ . Solid lines, mode 1; dotted lines, mode 2; dashed lines, mode 3.




Figure 7. CDF of local error against  $\varepsilon/\varepsilon_{\text{ref}}$ . Left plot, mode 1; middle plot, mode 2; right plot, mode 3. Solid lines,  $\varepsilon_{\text{tol}} = 1 \times 10^{-4}$ ; dashed lines,  $\varepsilon_{\text{tol}} = (1/64) \times 10^{-4}$ .

$\varepsilon_{\text{tol}} = (1/64) \times 10^{-4}$  clearly shows the tail of  $1 - F(x)$  shifted to the right significantly as compared with the case of  $\varepsilon_{\text{tol}} = 1 \times 10^{-4}$ , indicating that error control performance worsens when  $\varepsilon_{\text{tol}}$  is decreased.

After this study was completed, it was discovered that the computation of a Householder vector used in the algorithm to “grow” EOAs is on rare occasions ill-conditioned. This can result in EOA grows being inaccurate, and hence retrieves from such EOAs being inaccurate. Some of the observed large retrieve errors may be attributable to this ill-conditioning. The retrieve errors measured and reported here are, therefore, upper bounds on the errors that would be incurred in calculations without the ill-conditioning.

## 5.6 Conclusions on local error

Here, some conclusions can be drawn. Recall that the statistics of the local retrieve error  $\varepsilon$  are fully characterized by the CDF,  $F(x) \equiv \text{Prob}\{\varepsilon < x\}$ .

1. Large local errors are observed, albeit with small probability. For example, with ellipsoidal growing (grow mode 1) and  $\varepsilon_{\text{tol}} = 10^{-4}$ ,  $\varepsilon$  exceeds about  $20\varepsilon_{\text{ref}}$  or  $68\varepsilon_{\text{tol}}$  with 0.1% probability (see table 4).
2. The observed large local errors result from the growing process. When growing is suppressed, the largest errors are just a few percent of  $\varepsilon_{\text{tol}}$ .
3. The large local errors are not caused by non-monotonic behaviour of  $\varepsilon$ . If this were the case, error control would improve as  $\varepsilon_{\text{tol}}$  is decreased, but this is not observed.

Table 6. Centre of mixture fraction range of width 0.1 for conditional means.

Quantity	T	CO <sub>2</sub>	H <sub>2</sub> O	CO	H <sub>2</sub>	OH
Centre of range	0.37	0.37	0.41	0.47	0.53	0.33

- The large errors are associated with non-convex regions of accuracy. This follows from the fact that large errors are observed even with conical growing (grow mode 3), which guarantees error control (i.e.  $\varepsilon \leq \varepsilon_{\text{tol}}$ ) if the ROA is convex (for small  $\varepsilon_{\text{tol}}$ ).
- The CDFs of  $\varepsilon/\varepsilon_{\text{ref}}$  are generally the same for all growing strategies and values of  $\varepsilon_{\text{tol}}$ . Hence, there is no reason not to use the simplest growing strategy, i.e. ellipsoidal growing (grow mode 1).

## 6. Global error

### 6.1 Introduction and motivation

This section describes the characterization of the global error in the PDF calculations arising from the local errors incurred in the ISAT algorithm. This is achieved by considering the error in means of species mass fractions and temperature conditional upon the mixture fraction being in a specified range of width 0.1 around stoichiometric, estimated from the PDF calculations of flame D. The global error of any quantity  $X$ ,  $\varepsilon_{G,X}$  is defined as the difference between the computed value of  $X$  using ISAT and the value obtained using direct integration. If the global error arising from the ISAT algorithm is controlled, the quantities considered (conditional species mass fractions and temperature) should converge to the value given by direct integration as the error tolerance  $\varepsilon_{\text{tol}}$  tends to zero.

### 6.2 Testing protocol

The quantities examined are means of species mass fractions and temperature conditional upon the mixture fraction being in a specified range of width 0.1 in the vicinity of stoichiometric. In the experiments [1], the Bilger mixture fraction is used; in the calculations, since equal diffusivities are assumed, all definitions of mixture fraction are equivalent. The ranges for the variables (which are shown in table 6) are chosen to be centered on the peak of the mean conditioned on mixture fraction. Preliminary investigation revealed that these quantities, which are denoted by  $\langle Y_{H_2} | \xi_s \rangle$ , etc., are much more sensitive to ISAT errors than are unconditional means. The conditional means at each axial location are evaluated using PDF particles in 9 bins defined in  $x - r$  space. The ranges of the bins for each axial location are given in table 7. The lower and upper bound of the bins in the  $x$ -direction are denoted as  $x_{\text{low}}$  and  $x_{\text{up}}$ , respectively, and the upper bound of the bins in the  $r$ -direction is denoted  $r_{\text{up}}$ . The lower bound in the  $r$ -direction is zero for all axial locations.

Table 7. Ranges in  $x$  and  $r$  of bins used to evaluate conditional means.

Axial location ( $x/d$ )	1	2	3	7.5	15	30	45	60	75
$x_{\text{low}}/d$	0.75	1.7	2.7	7	14.3	28.8	43.5	58	73
$x_{\text{up}}/d$	1.25	2.3	3.3	8.0	15.7	31.2	46.5	62	77
$r_{\text{up}}/d$	2.5	2.5	2.5	3.1	4.0	7.0	10	12	15

The computations made to characterize the global error were performed in parallel. Each parallel computation node is equipped with 2 GB of RAM and two 2.4 GHz Intel Xeon processors. The nodes are connected using gigabit ethernet.

The computations are performed as follows: starting from the statistically stationary solution, 600 iterations are performed using 12 parallel processes starting from an empty ISAT table. The domain is partitioned into rectangular subdomains, each of which extends over the whole radius, and over a fraction of the axial extent of the domain, with each subdomain having its own (independent) ISAT table. For all variables, the absolute and relative integration error tolerances in DDASAC are set to  $10^{-6}$  and  $10^{-9}$ , respectively. The maximum storage of the ISAT table is set to 400 megabytes, and grow mode 2 is used. To reduce statistical variability, the number of particles per cell is increased to 240. These 600 iterations are performed to build up the ISAT tables. Subsequently, a further 200 iterations are performed. Values of conditional means are time-averaged over these 200 iterations. A series of 6 computations is performed with  $\varepsilon_{\text{tol}}$  varying (by factors of 4) between  $(1/64) \times 10^{-4}$  and  $16 \times 10^{-4}$ . For all values of  $\varepsilon_{\text{tol}}$  except  $\varepsilon_{\text{tol}} = (1/64) \times 10^{-4}$ , 5 statistically identical runs are performed and the 95% confidence intervals are calculated.

### 6.3 Results

Figures 8–13 show the conditional temperature and species mass fractions (linear scale) against  $\varepsilon_{\text{tol}}$  (log scale) at different axial locations. The symbols are from the computations. The vertical lines (most clearly seen in figure 9) correspond to the 95% confidence intervals, and

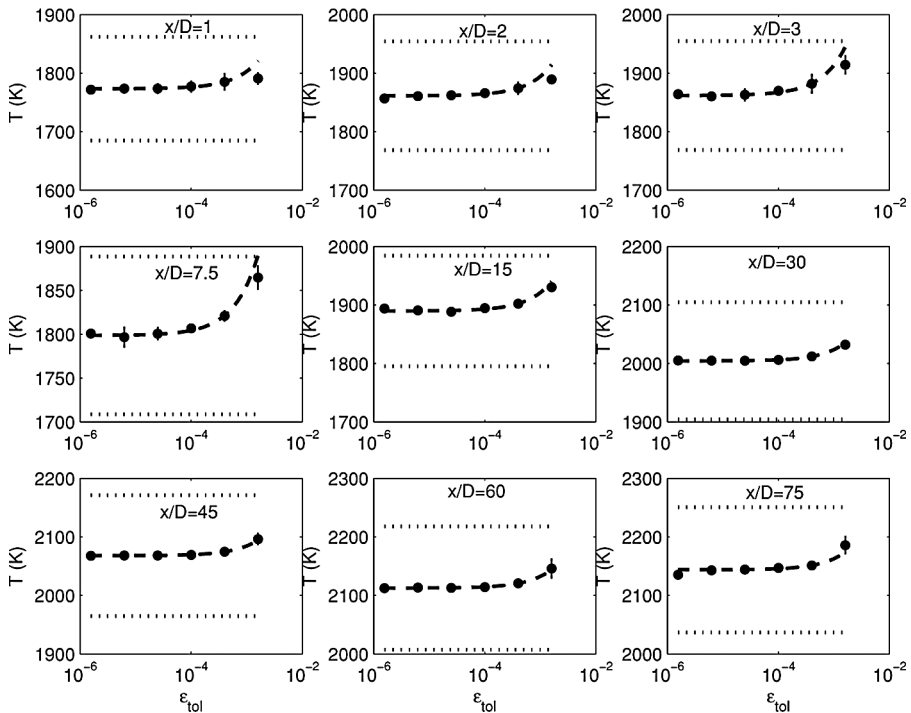


Figure 8. Plot of mean temperature conditional on stoichiometric ( $T|\xi_s$ ) against error tolerance  $\varepsilon_{\text{tol}}$  with confidence intervals on log-linear scale. Symbols, computations; dashed line, linear fit through middle four points; dotted lines, plus and minus 5% error.

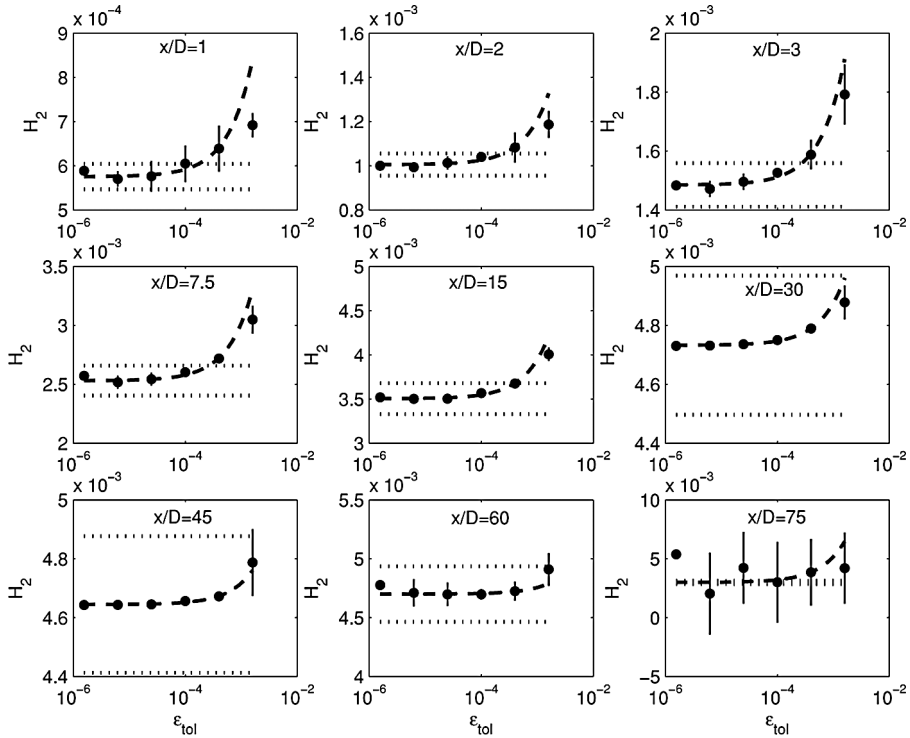


Figure 9. Plot of  $\langle Y_{H_2} | \xi_s \rangle$  against  $\epsilon_{\text{tol}}$  with confidence intervals on log-linear scale. Symbols, computations; vertical solid lines, 95% confidence intervals; dashed line, linear fit through middle four points; dotted lines, plus and minus 5% error.

these are plotted for all quantities and all values of  $\epsilon_{\text{tol}}$  except  $\epsilon_{\text{tol}} = (1/64) \times 10^{-4}$ . In most cases, the confidence interval is no larger than the symbol size in the figures, and hence cannot be seen. The dashed line is a linear fit through the middle four points at each axial location, which appears as a curve in these semi-log plots. The point with the largest value of  $\epsilon_{\text{tol}}$  is not included in the linear fit because it has very large errors; and the point with the smallest value of  $\epsilon_{\text{tol}}$  is not included in the linear fit as it is based on a single simulation and hence is prone to greater statistical variability. The dotted lines show plus and minus 5% about an estimate of the true value (obtained by extrapolating the straight line fit to the y-axis).

Before describing the results, three points need to be stressed. First, the conditional statistics are selected for examination because they are found to be much more sensitive than unconditional statistics. Second, the results (except for  $\epsilon_{\text{tol}} = (1/64) \times 10^{-4}$ ) are based on 5 independent simulations, each with 240 particles per cell (i.e. effectively 1200 particles per cell). As a consequence, the statistical errors are small compared to those of typical PDF calculations. Lastly, it is important to remember that the result with  $\epsilon_{\text{tol}} = (1/64) \times 10^{-4}$  is based on a single simulation (compared to 5 independent simulations for other values of  $\epsilon_{\text{tol}}$ ) and hence is subject to greater statistical variability.

One observation is that, for many of the quantities plotted, the confidence intervals are large for the one or two largest values of  $\epsilon_{\text{tol}}$ . This is due to the fact that for large values of  $\epsilon_{\text{tol}}$ , there exist large errors, and these are random (due to the nature of the PDF/ISAT calculations). These errors can be viewed as bias and statistical errors, both of which

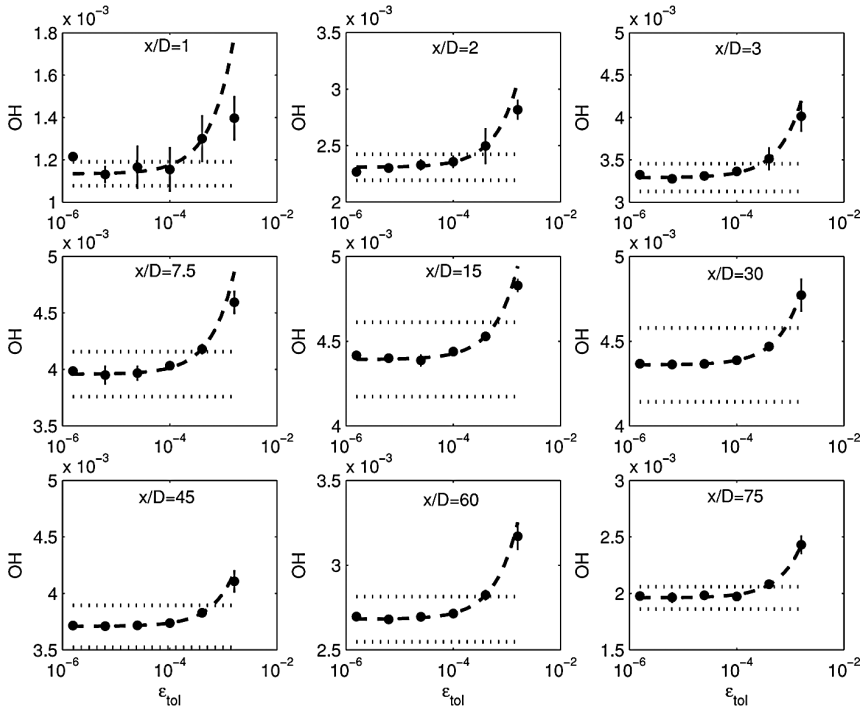


Figure 10. Plot of  $\langle Y_{OH} | \xi_s \rangle$  against  $\epsilon_{tol}$  with confidence intervals on log-linear scale. Symbols, computations; vertical solid lines, 95% confidence intervals; dashed line, linear fit through middle four points; dotted lines, plus and minus 5% error.




Figure 11. Plot of  $\langle Y_{CO} | \xi_s \rangle$  against  $\epsilon_{tol}$  with confidence intervals on log-linear scale. Symbols, computations; vertical solid lines, 95% confidence intervals; dashed line, linear fit through middle four points; dotted lines, plus and minus 5% error.



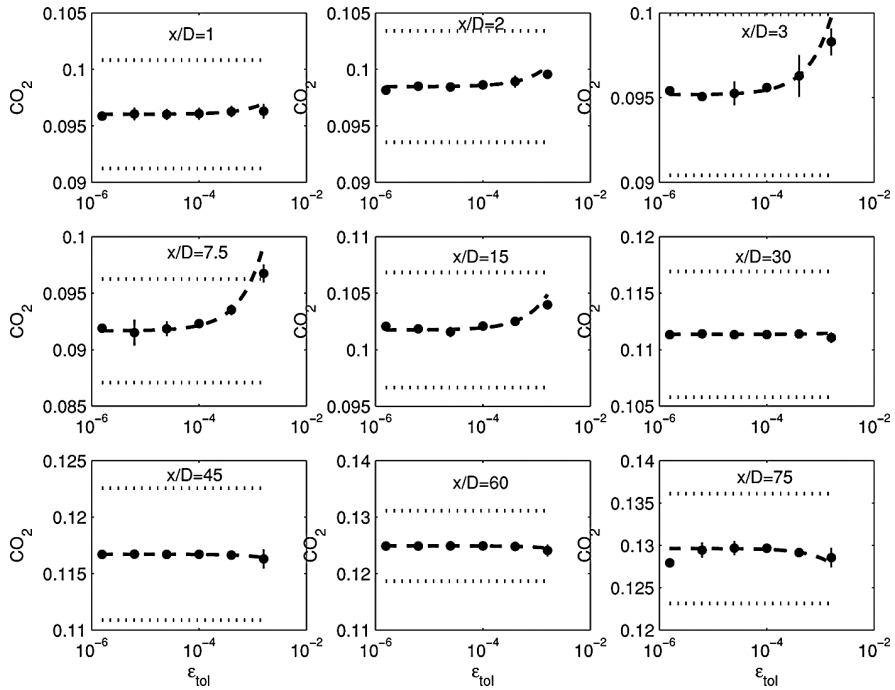


Figure 12. Plot of  $\langle Y_{CO_2} | \xi_s \rangle$  against  $\epsilon_{tol}$  with confidence intervals on log-linear scale. Symbols, computations; vertical solid lines, 95% confidence intervals; dashed line, linear fit through middle four points; dotted lines, plus and minus 5% error.

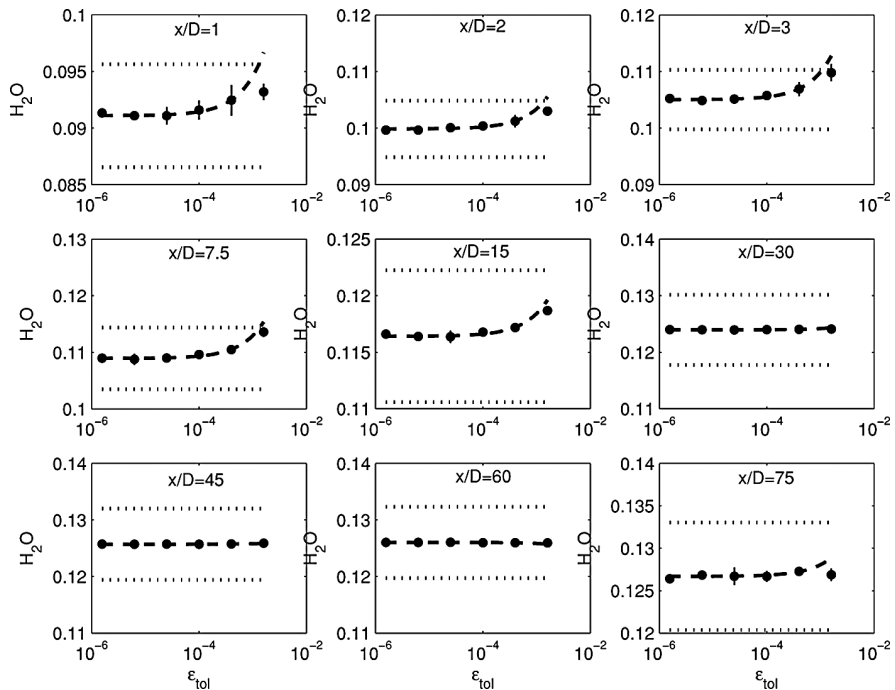


Figure 13. Plot of  $\langle Y_{H_2O} | \xi_s \rangle$  against  $\epsilon_{tol}$  with confidence intervals on log-linear scale. Symbols, computations; vertical solid lines, 95% confidence intervals; dashed line, linear fit through middle four points; dotted lines, plus and minus 5% error.

Table 8. Values of  $\varepsilon_{\text{tol}}$  for which the global error is one percent.

	Axial location ( $x/d$ )								
	1	2	3	7.5	15	30	45	60	75
$\varepsilon_T^{0.01} \times 10^4$	6.00	5.65	3.60	3.16	6.06	10.40	12.55	10.66	11.30
$\varepsilon_{\text{H}_2}^{0.01} \times 10^4$	0.35	0.50	0.56	0.52	0.79	3.32	6.44	8.49	0.14
$\varepsilon_{\text{OH}}^{0.01} \times 10^4$	0.28	0.49	0.58	0.70	1.28	1.61	1.25	0.75	0.68
$\varepsilon_{\text{CO}}^{0.01} \times 10^4$	0.39	0.63	0.74	1.04	2.10	15.78	10.96	6.25	0.35
$\varepsilon_{\text{CO}_2}^{0.01} \times 10^4$	16.07	8.76	3.32	1.93	5.20	158.15	61.16	52.61	11.70
$\varepsilon_{\text{H}_2\text{O}}^{0.01} \times 10^4$	2.62	2.83	2.17	2.72	5.82	55.86	141.5	76.10	9.34

increase with  $\varepsilon_{\text{tol}}$ , and hence it is logical that the confidence intervals are larger for large  $\varepsilon_{\text{tol}}$ .

The results for conditional temperature at  $x/D = 7.5$  are representative of most of the other results also. It may be seen from figures 8–13 that for the four smallest values of  $\varepsilon_{\text{tol}}$  ( $1/64$ ,  $1/16$ ,  $1/4$  and  $1 \times 10^{-4}$ ) all of the calculated results are in agreement (to within confidence intervals). Indeed, this observation is true for other axial locations, with the possible exception of  $x/D = 75$ . This is indicative of the errors due to ISAT being small relative to statistical errors for  $\varepsilon_{\text{tol}} = 10^{-4}$ , and of their decreasing as  $\varepsilon_{\text{tol}}$  is further decreased. On the other hand, for  $\varepsilon_{\text{tol}} = 4 \times 10^{-4}$ , and more so for  $\varepsilon_{\text{tol}} = 16 \times 10^{-4}$ , there are substantial errors due to ISAT – about 60 K at this location. Future PDF calculations of this test case using ISAT should be made with  $\varepsilon_{\text{tol}} = 10^{-4}$  to ensure an efficient calculation without sacrificing accuracy. However, in general, the appropriate values of  $\varepsilon_{\text{tol}}$  depends on many factors (e.g. the chemical mechanism, time step range, etc.).

From figures 8–13, we see that in general, the straight line fit passes through at least the points corresponding to the five smallest values of  $\varepsilon_{\text{tol}}$ . This clearly shows that the ISAT global error varies linearly with  $\varepsilon_{\text{tol}}$  (as does the local error and as is expected). Tables 8 and 9 show the values of  $\varepsilon_{\text{tol}}$  which result in a global error of 1% and 10%, respectively, for each of the five species considered, as well as for temperature at each axial location. These values are computed from the linear fit through the middle four points in each plot. We denote by  $\varepsilon_{\beta}^{\alpha}$  the value of  $\varepsilon_{\text{tol}}$  for which statistic  $\beta$  has a global error of  $\alpha$ . For example, the value of  $\varepsilon_{\text{tol}}$  for which (conditional) temperature has a 5% error is denoted as  $\varepsilon_T^{0.05}$ . Clearly, the global error due to ISAT in the PDF computation of this piloted jet flame test case is well controlled.

Table 9. Values of  $\varepsilon_{\text{tol}}$  for which the global error is 5%.

	Axial location ( $x/d$ )								
	1	2	3	7.5	15	30	45	60	75
$\varepsilon_T^{0.05} \times 10^4$	30.00	28.23	18.00	15.80	30.29	52.00	62.74	53.28	56.48
$\varepsilon_{\text{H}_2}^{0.05} \times 10^4$	1.74	2.50	2.78	2.61	3.97	16.59	32.19	42.44	0.69
$\varepsilon_{\text{OH}}^{0.05} \times 10^4$	1.39	2.44	2.92	3.50	6.39	8.04	6.23	3.75	3.38
$\varepsilon_{\text{CO}}^{0.05} \times 10^4$	1.94	3.16	3.71	5.21	10.51	78.94	54.79	31.25	1.75
$\varepsilon_{\text{CO}_2}^{0.05} \times 10^4$	80.37	43.81	16.61	9.66	25.98	790.73	305.82	263.05	58.51
$\varepsilon_{\text{H}_2\text{O}}^{0.05} \times 10^4$	13.09	14.13	10.86	13.58	29.11	279.3	707.5	380.5	46.64

## 7. Conclusions

The ISAT algorithm was used in conjunction with the CFD code FLUENT to perform calculations of a turbulent non-premixed piloted methane–air jet flame (Barlow and Frank flame D) to characterize the local and global error characteristics of ISAT. This section summarizes, in turn, conclusions drawn regarding the local and global error due to ISAT.

With regard to local error, the following conclusions can be drawn. Recall that the local error  $\varepsilon$  is completely characterized by its CDF,  $F(x) = \text{Prob}\{\varepsilon < x\}$ , and  $\varepsilon_{\text{ref}}$  denotes the 90th percentile error.

- (i) From figures 6 and 7, it is observed that large local errors are present for all growing strategies, albeit with small probability. From tables 4 and 5, it is shown that when growing is suppressed, large local errors are no longer observed, indicating that large local errors result from the growing process.
- (ii) From figure 7, the CDFs of  $\varepsilon$  are found to be largely independent of  $\varepsilon_{\text{tol}}$ , suggesting that large errors are not due to the non-monotonicity of  $\varepsilon$ .
- (iii) Large local errors are associated with non-convex ROAs. This follows from the fact that large errors are observed even with grow mode 3, which guarantees error control for convex ROAs.
- (iv) When  $\varepsilon$  is scaled by  $\varepsilon_{\text{ref}}$ , the CDFs of  $\varepsilon$  are generally the same for all grow modes, as can be seen from figure 6. Hence, the simplest strategy (i.e. ellipsoidal growing, grow mode 1) should be used.

Parallel computations were performed to obtain conditional mean quantities in order to examine the global error in ISAT. Statistical variability in the conditional mean quantities was reduced considerably by time averaging the quantities over 200 iterations and using a large number of particles per cell ( $N_{pc} = 240$ ). The following observations and conclusions regarding the global error can be made from figures 8–13.

- (i) The confidence intervals for the plotted quantities are large for the one or two largest values of  $\varepsilon_{\text{tol}}$ , due to the occurrence of large statistical errors at large  $\varepsilon_{\text{tol}}$ .
- (ii) In general, for the four smallest values of  $\varepsilon_{\text{tol}}$ , the calculated results are in agreement, indicating that ISAT errors are small relative to statistical errors for  $\varepsilon_{\text{tol}} \leq 10^{-4}$ .
- (iii) Figures 8–13 and tables 8 and 9 show that the global error  $\varepsilon_G$  varies linearly with  $\varepsilon_{\text{tol}}$ , as expected.

Future work is focused on the improvement of ISAT local error control through the determination of the geometry of the region of accuracy, and the development of an enhanced growing algorithm which limits the extent by which EOAs grow to encompass inaccurate regions. The eventual aim is to achieve excellent local error control, while still maintaining the excellent speed-up provided by ISAT. It is important to appreciate that even in the absence of effective local error control, the global error in ISAT is small for  $\varepsilon_{\text{tol}} \leq 10^{-4}$ , and found to vary linearly with  $\varepsilon_{\text{tol}}$  in this calculation.

## Acknowledgements

The authors wish to thank S. A. Lantz and L. P. Chew for assistance in various parts of this work. This work is supported by the US Department of Energy grant number DE-FG02-90ER14128. This research was conducted using the resources of the Cornell Theory Center, which receives funding from Cornell University, New York State, federal agencies, foundations, and corporate partners.

## References

- [1] Barlow, R.S. and Frank, J.H., 1998, Effects of turbulence on species mass fractions in methane/air jet flames. *Proceedings of the Combustion Institute*, **27**, 1087–1095.
- [2] Barlow, R.S. and Frank, J.H., 2003, Piloted CH<sub>4</sub>/Air flames C, D, E, and F—Release 2.0 *TNF Workshop*. Available online at: <http://www.ca.sandia.gov/TNF/DataArch/FlameD/SandiaPilotDoc20.pdf>. Accessed 1 July 2004.
- [3] Barth, T.J. and Jespersen, D., 1989, The design and application of upwind schemes on unstructured meshes. *Technical Report AIAA-89-0366, AIAA 27th Aerospace Sciences Meeting*, Reno, Nevada.
- [4] Caracotsios, M., and Stewart, W.E., 1985, Sensitivity analysis of initial-value problems with mixed ODEs and algebraic equations. *Computers & Chemical Engineering*, **9**, 359–365.
- [5] Chew, L.P., Private Communication.
- [6] FLUENT 6.1 User's Guide, FLUENT Inc. Available online at: [http://www.fluentusers.com/fluent61/doc/ori/html/ug/main\\_pre.htm](http://www.fluentusers.com/fluent61/doc/ori/html/ug/main_pre.htm). Accessed 1 July 2004.
- [7] Janicka, J., Kolbe, W. and Kollman, W., 1977, Closure of the transport equation for the probability density function of turbulent scalar fields. *Journal of Non-Equilibrium Thermodynamics*, **4**, 47–66.
- [8] Kapoor, R., Lentati, A. and Menon, S., 2001, Simulations of Methane–Air Flames using ISAT and ANN. In *37th AIAA/ASME/SAW/ASEE Joint Propulsion Conference, AIAA 2001–3847*.
- [9] Launder, B.E. and Spalding, D.B., 1974, The numerical computation of turbulent flows. *Computer Methods in Applied Mechanics and Engineering*, **3**, 269–289.
- [10] Lindstedt, R.P., Louloudi, S.A. and Vaos, E.M., 2000, Joint scalar probability density function modelling of pollutant formation in piloted turbulent jet diffusion flames with comprehensive chemistry. *Proceedings of the Combustion Institute*, **29**, 149–156.
- [11] Liu, B.J.D., 2004, Performance of *in situ* adaptive tabulation in computation of non-premixed turbulent piloted jet flames. MS thesis, Cornell University.
- [12] Patankar, S.V., 1980, *Numerical Heat Transfer and Fluid Flow* (Hemisphere: Washington, DC).
- [13] Pope, S.B., 1997, Computationally efficient implementation of combustion chemistry using *in situ* adaptive tabulation. *Combustion Theory and Modelling*, **1**, 41–63.
- [14] Saxena, V. and Pope, S.B., 1998, PDF calculations of major and minor species in a piloted turbulent jet flame. *Proceedings of the Combustion Institute*, **27**, 1081.
- [15] Singer, M.A. and Pope, S.B., 2004, Exploiting ISAT to solve the reaction-diffusion equation. *Combustion Theory and Modelling*, **8**, 361–383.
- [16] Tang, Q., 2003, Computational modelling of turbulent combustion with detailed chemistry. Ph.D. thesis, Cornell University.
- [17] Tang, Q., Xu, J., and Pope, S.B., 2000, Probability density function calculations of local extinction and NO production in piloted-jet turbulent methane/air flames. *Proceedings of the Combustion Institute*, **28**, 133–139.
- [18] Xu, J., and Pope, S.B., 2000, PDF calculations of turbulent nonpremixed flames with local extinction. *Combustion and Flame*, **123**, 281–307.