



ELSEVIER

Available online at www.sciencedirect.com

SciVerse ScienceDirect

Proceedings of the Combustion Institute 34 (2013) 205–215

**Proceedings
of the
Combustion
Institute**

www.elsevier.com/locate/proci

Large-scale parallel simulations of turbulent combustion using combined dimension reduction and tabulation of chemistry

Varun Hiremath^{a,*}, Steven R. Lantz^b, Haifeng Wang^a,
Stephen B. Pope^a

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

^b Center for Advanced Computing, Cornell University, Ithaca, NY 14853, USA

Available online 25 September 2012

Abstract

Simulations of turbulent reacting flows with chemistry represented using detailed kinetic model involving a large number of species and reactions are computationally expensive. Here we present a combined dimension reduction and tabulation strategy for implementing chemistry in large scale parallel *Large-Eddy Simulation* (LES)/*Probability Density Function* (PDF) computations of turbulent reacting flows. In this approach, the dimension reduction is performed using the *Rate Controlled Constrained-Equilibrium* (RCCE) method, and tabulation of the reduced space is performed using the *In Situ Adaptive Tabulation* (ISAT) algorithm. In addition, we use *x2f_mpi* — a Fortran library for parallel vector-valued function evaluation (used with ISAT in this context) — to efficiently redistribute the chemistry workload among the participating cores in parallel LES/PDF computations to reduce the overall wall clock time of the simulation. We test three parallel strategies for redistributing the chemistry workload, namely (a) PLP, purely local processing; (b) URAN, the uniform random distribution of chemistry computations among all cores following an early stage of PLP; and (c) P-URAN, a Partitioned URAN strategy that redistributes the workload within partitions or subsets of the cores. To demonstrate the efficiency of this combined approach, we perform parallel LES/PDF computations (on 1024 cores) of the Sandia Flame D with chemistry represented using a 38-species C_1 – C_4 skeletal mechanism. We show that relative to using ISAT alone with the 38-species full representation, the combined ISAT/RCCE approach with 10 represented species (i) predicts time-averaged mean and standard deviation statistics with a normalized root-mean-square difference of less than 3% (30 K) in temperature, less than 2% (0.02 kg/m^3) in density, less than 2.5% in mass fraction of major species, and less than 8% in mass fraction of minor species of interest; and (ii) reduces the simulation wall clock time by over 40% with the P-URAN strategy.

© 2012 The Combustion Institute. Published by Elsevier Inc. All rights reserved.

Keywords: Large-Eddy Simulation; Probability Density Function method; *In situ* adaptive tabulation; Rate controlled constrained-equilibrium; Parallel function evaluation

* Corresponding author.

E-mail address: vh63@cornell.edu (V. Hiremath).

1. Introduction

Detailed chemical mechanisms of hydrocarbon fuels may involve hundreds or thousands of species and thousands of reactions [1,2]. Incorporating directly such detailed chemistry in the combustion flow calculations is computationally prohibitive, even using distributed parallel computing. The current efforts aimed at reducing the computational cost of representing chemistry can be placed under three main categories: (1) *mechanism reduction*, the generation of reaction mechanisms involving fewer species and reactions [3–5]; (2) *dimension reduction*, the representation of chemistry using a reduced number of variables [6–9]; and (3) *tabulation*, the use of storage-retrieval methods to reduce significantly the cost of expensive evaluations of the reaction mappings involving ODE integrations [10–13]. Combined methodologies have also been developed to use reduced reaction mechanisms or dimension reduction methods in conjunction with tabulation [15–17].

Since most of the modern day reactive flow simulations are performed in parallel on multiple cores using distributed computing, in addition to the aforementioned techniques, strategies are needed to efficiently distribute the chemistry workload among the participating cores to reduce the overall wall clock time of the simulation [18–21]. We recently demonstrated scalable parallel strategies implemented using *x2f_mpi* for the efficient redistribution of chemistry workload in large scale parallel Large-Eddy Simulation (LES)/Probability Density Function (PDF) computations [22].

In this paper we further extend our LES/PDF solver with the capability of representing chemistry using our combined dimension reduction and tabulation approach [15]. In this approach, the dimension reduction is performed using the Rate Controlled Constrained-Equilibrium (RCCE) [6,23,24] method followed by tabulation using the *In Situ* Adaptive Tabulation (ISAT) [10,11] algorithm. In [15], we tested our combined dimension reduction and tabulation approach using the partially-stirred reactor for methane and ethylene chemistry, and the main conclusions drawn from this work are that the ISAT/RCCE approach

- yields the same level of accuracy as other reduced (based on the Quasi-Steady State Assumption, QSSA) or skeletal mechanisms with relatively fewer represented species;
- yields significant speedup relative to using ISAT alone with the detailed mechanism.

Here, for the first time, the ISAT/RCCE approach is being demonstrated in the context of full-scale LES/PDF simulations of turbulent

reacting flows using realistic chemistry. In this study, the accuracy and efficiency of this combined approach is demonstrated by performing full-scale LES/PDF simulation of the Sandia Flame D [25].

The outline of the remainder of the paper is as follows: in Section 2 we describe our combined LES/PDF solver; in Section 3 we describe the combined dimension reduction and tabulation strategy; in Section 4 we briefly describe the parallel strategies implemented using *x2f_mpi* for redistributing the chemistry workload in large scale LES/PDF computations; in Section 5 we present computational details for simulating the Sandia Flame D; in Section 6 we present simulation results; and finally in Section 7 we state our conclusions.

2. Combined LES/PDF solver

In this study we use a combined LES/PDF solver developed at Cornell as described in more detail in [22,26]. Below we mention some of the key aspects of this solver.

2.1. LES solver

The LES solver is based on a Stanford LES code [27,28]. The solver uses structured non-uniform grids; supports cylindrical coordinate system; is second-order accurate in space and time; and is parallelized (using MPI) by domain decomposition in two dimensions.

2.2. PDF solver

The PDF solver, *HPDF* [26], has second-order accuracy in space and time; supports Cartesian and cylindrical coordinate systems; is parallelized (using MPI) by domain decomposition in two dimensions; and has a general interface to facilitate coupling with existing LES (or RANS) solvers.

In the PDF approach, the thermochemical composition of the fluid within the solution domain is represented by a large number of particles. The HPDF solver has three main components

1. *transport*: to account for the change in *position* of a particle due to advection in the physical space (including a random-walk component to represent the effects of subgrid-scale turbulent advection and molecular diffusion);
2. *mixing*: to account for the change in *composition* of a particle due to mixing with neighboring particles (which models the effects of molecular mixing); and
3. *reaction*: to account for the change in *composition* of a particle due to chemical reaction.

These components are implemented in *fractional steps* using splitting schemes [29]. In this study, to simulate the Sandia Flame D, we use the first-order TMR splitting scheme (which is found to perform as well as the second-order splitting scheme for jet flames [26]). The TMR splitting scheme denotes taking fractional steps of *transport*, \mathbb{T} ; *mixing*, \mathbb{M} ; and *reaction*, \mathbb{R} in this order on each time-step. The Kloeden and Platen (KP) [30] stochastic differential equation (SDE) scheme is used to integrate the *transport* equations; and the *mixing* is represented using the modified Curl [31] mixing model. The *reaction* fractional step is implemented using the combined dimension reduction and tabulation approach which we will discuss in more detail in the later sections.

2.3. Domain decomposition

The LES computations are performed on a structured non-uniform grid in the cylindrical coordinate system. We denote the grid used for LES computations by $N_x \times N_r \times N_\theta$ (in the axial, radial and azimuthal directions). In performing parallel LES/PDF computations (using the combined LES/HPDF solver) on N_c cores, the computational domain is decomposed into N_c sub-domains and each core performs the computations of one sub-domain. The domain decomposition is done in the first two principal directions x and r , and is denoted by $D_x \times D_r$, where $D_x D_r = N_c$. For instance, in this study we perform LES/PDF simulations of the Sandia Flame D using a non-uniform LES grid of size $N_x = 192$, $N_r = 192$, $N_\theta = 96$ on $N_c = 1024$ cores using a domain decomposition with $D_x = 64$ and $D_r = 16$.

3. Combined dimension reduction and tabulation

In this section we briefly describe the combined dimension reduction and tabulation approach used for representing chemistry using ISAT/RCCE. More detailed description can be found in [15].

3.1. Particle representation

We consider a reacting gas-phase mixture consisting of n_s chemical species, composed of n_e elements. We consider an isobaric system with a fixed specified pressure p , and so the thermochemical state of the mixture (at a given position and time) is completely characterized by the mixture enthalpy h , and the n_s -vector \mathbf{z} of specific moles of the species.

In the *reaction fractional step*, a particle's chemical composition \mathbf{z} evolves (at constant enthalpy h) in time according to the following set of n_s coupled ordinary differential equations (ODEs)

$$\frac{d\mathbf{z}(t)}{dt} = \mathbf{S}(\mathbf{z}(t)), \quad (1)$$

where \mathbf{S} is the n_s -vector of chemical production rates determined by the chemical mechanism used to represent the chemistry.

Given a reaction fractional time step Δt , the *reaction mapping*, $\mathbf{z}(\Delta t)$, is defined to be the solution to Eq. (1) after time Δt starting from the initial composition $\mathbf{z}(0)$. The reaction mapping obtained by directly integrating the set of ODEs given by Eq. (1) is referred to as a *direct evaluation* (DE). We use DDASAC [32] for performing ODE integration.

Owing to the large cost of direct evaluation of reaction mappings involving large numbers of species, we use a combined dimension reduction (using RCCE) and tabulation (using ISAT) strategy for representing chemistry. This combined methodology can be applied to chemical systems involving a large number of species (100–1000) by first applying the dimension reduction to reduce the dimensionality of the system to say around 20 (depending on the level of accuracy needed) and then using ISAT to tabulate the reaction mappings in the reduced dimension.

3.2. Dimension reduction

In this section we briefly describe the procedure followed for dimension reduction in our implementation of the RCCE method; a more detailed description can be found in [14,15].

In our implementation of RCCE, to perform the dimension reduction we specify a set of n_{rs} *represented* (constrained) species selected from the full set of n_s species. Consequently, we have $n_{us} = n_s - n_{rs}$ *unrepresented* species.

The selection of good represented species is crucial for the accuracy of the RCCE dimension reduction method. We have devised an automated Greedy Algorithm with Local Improvement (GALI) [14,15] to select good represented species based on a specified measure of dimension reduction error. The greedy algorithm selects represented species in stages one-by-one which minimizes the specified measure of dimension reduction error [14].

The *reduced representation* of the species composition is denoted by $\mathbf{r} \equiv \{\mathbf{z}^r, \mathbf{z}^{u,e}\}$, where \mathbf{z}^r is an n_{rs} -vector of specific moles of the represented species; and $\mathbf{z}^{u,e}$ is an n_e -vector of specific moles of the elements in the unrepresented species (for atom conservation). Thus, \mathbf{r} is a vector of length $n_r = n_{rs} + n_e$, and the dimension of the system is reduced from n_s to n_r . At any time t , the *reduced representation*, $\mathbf{r}(t)$, is related to the *full representation*, $\mathbf{z}(t)$, by

$$\mathbf{r}(t) = \mathbf{B}^T \mathbf{z}(t), \quad (2)$$

where \mathbf{B} is a constant $n_s \times n_r$ matrix which can be determined for a specified set of represented species. (In general, the reduced representation in RCCE can be a linear or non-linear function of the full representation [33].)

In the HPDF solver, with dimension reduction, the particles carry only the reduced representation $\{\mathbf{r}, h\}$. Given the reduced representation \mathbf{r} , the temperature T and density ρ are approximated using the method described in the Appendix of [15]. Later in the Results (Section 6.1) we show that this approximation method yields values of temperature and density that match closely with those obtained with the full representation \mathbf{z} .

Given the reduced representation at the beginning of the reaction fractional step $\mathbf{r}(0)$, and the reaction fractional time step Δt , the *reduced reaction mapping* $\mathbf{r}(\Delta t)$ (at constant enthalpy) is computed using the following steps:

1. *species reconstruction*: given $\mathbf{r}(0)$, we compute the constrained-equilibrium composition at constant enthalpy, $\mathbf{z}^{CE}(\mathbf{r}(0))$, using CEQ [34];
2. *reaction mapping*: starting from $\mathbf{z}^{CE}(\mathbf{r}(0))$, we solve the full system of n_s ODEs Eq. (1) to obtain the reaction mapping, $\mathbf{z}(\Delta t)$;
3. *reduction*: we obtain the *reduced reaction mapping* as, $\mathbf{r}(\Delta t) = \mathbf{B}^T \mathbf{z}(\Delta t)$.

The above steps of course make the computation of the reaction mapping even more expensive than directly solving the full set of ODEs Eq. (1), due to the additional species reconstruction and reduction steps. However, when ISAT is used in conjunction with dimension reduction, the computational cost is reduced significantly as explained in the next section.

A more efficient way of obtaining the reduced reaction mapping, $\mathbf{r}(\Delta t)$, is to directly solve a reduced system of n_r ODEs for the constraints, \mathbf{r} , or for the constraint-potentials (as is done in the classical RCCE approach [35,36]). We are currently working on implementing this method which should give a further improvement in performance. Nevertheless, even with our current implementation of RCCE, we achieve significant reduction in computational cost relative to the detailed chemistry calculation as shown in this work.

3.3. Tabulation

We use *in situ* adaptive tabulation (ISAT) [10] for tabulating the reaction mappings. The ISAT algorithm has been successfully applied in many combustion chemistry calculations involving up to $n_s \leq 50$ species [11,15]. However, with chemistry involving more than 50 species, the direct use of ISAT may not be very efficient, due to the large table size and search times [15].

Hence, for chemistry involving more than say $n_s \geq 30$ species, we use the RCCE dimension reduction method to represent the chemistry using a reduced representation involving fewer n_r variables. Note, for very large mechanisms involving thousands of species, the direct use of RCCE/GALI may still result in $n_r \gg 30$ to achieve an acceptable level of accuracy. In such cases it will be more efficient to use ISAT/RCCE with a skeletal mechanism (based on the detailed mechanism) involving hundreds of species.

We use ISAT to tabulate the *reduced reaction mapping* in the reduced dimension n_r , which reduces significantly the overall computational cost because

1. the exact reduced reaction mapping is computed (using the steps listed in the previous section) only for a small fraction of particles (typically less than 1%); and for the majority of the particles a linear approximation to the reduced reaction mapping is obtained using the tabulated data;
2. since the tabulation is performed in a reduced dimension, n_r , the ISAT table size is reduced, which in turn reduces the table search and retrieve times; and
3. since the particle compositions are also stored in a reduced dimension, fetching particles from the memory is faster.

Consequently, the combined dimension reduction and tabulation approach using ISAT/RCCE is found to give an additional speedup by a factor of $\mathcal{O}(10)$ relative to using ISAT alone with the full representation for tests performed using the 111-species C_1 – C_4 USC Mech II detailed mechanism [15]. A more detailed description of our combined dimension reduction and tabulation approach is provided in [15].

4. Parallel strategies for implementing chemistry

In performing parallel LES/PDF computations on multiple cores using our LES/HPDF solver with chemistry represented using the combined dimension reduction and tabulation approach, each core has its own ISAT table for tabulating the chemistry. On the reaction fractional step, the reaction mappings for all the particles in the computational domain need to be evaluated. However, the chemical reactivity is in general not uniform across the entire domain. For example, in simulation of jet flames, the sub-domains in the flame front are chemically more reactive than sub-domains in the outer coflow/air. Thus, a direct call to ISAT on each core at the reaction fractional step can create *load imbalance* among the cores, leading to increase in the overall simulation wall clock time. Hence, at

the reaction fractional step, we use parallel strategies implemented using *x2f_mpi* to redistribute the particles among the participating cores for reaction mapping evaluation, thereby achieving a better load balance and reducing the overall simulation wall clock time.

In [22], we presented three parallel strategies, denoted by PLP, URAN and P-URAN, for redistributing the chemistry workload. We give a brief description of these strategies below.

1. Purely Local Processing (PLP):

In this strategy, the reaction mapping of all the particles on a core is evaluated using the local ISAT table without any message passing or load redistribution. This in some sense is the same as invoking ISAT directly from HPDF on each core without using the *x2f_mpi* interface. This strategy thus leads to significant load imbalance.

2. Uniform Random (URAN):

This strategy is the extreme opposite of the PLP strategy and aims at achieving statistically ideal load balancing by evenly distributing the chemistry workload among all the participating cores. The strategy involves one initial step of PLP to initialize the local ISAT tables. In the subsequent steps, on each core, first an attempt is made to retrieve the reaction mapping of particles from the local ISAT table (also referred to as a “quick try”). Following this, there is a uniform random distribution of all the unresolved particles (for which “quick try” failed) to all the cores. This strategy thus ensures that the workload is evenly balanced over all the cores, however, it also results in a large amount of all-to-all message passing.

3. Partitioned Uniform Random (P-URAN):

This strategy aims at achieving a balance between communication cost and load imbalance by using the PLP and URAN strategies over smaller partitions of cores. The P-URAN strategy works in two stages: in stage 1, for a specified duration of time τ (hours) the PLP strategy is used to resolve particles; then in stage 2, for the remainder of the time, the participating cores are partitioned into smaller groups of κ cores, and within each partition the URAN strategy is used to uniformly distribute the chemistry workload among the cores in that partition. We use the notation P-URAN $[\tau, \kappa]$ to describe the P-URAN strategy.

In [22], based on LES/PDF simulations of Sandia Flame D using ISAT alone we showed that among the aforementioned three strategies, the P-URAN strategy yields the lowest wall clock time. We also showed that the P-URAN strategy shows

good scaling up to 9000 cores. In this work we use these strategies in conjunction with combined dimension reduction and tabulation to compare their relative performance. Here we focus more on the gains achieved using the combined dimension reduction and tabulation approach and show that the simulation wall clock time can be further reduced using our combined ISAT/RCCE approach without losing too much accuracy.

5. LES/PDF simulation of Sandia Flame D

To test the chemistry implementation we perform LES/PDF simulations of the Sandia Flame D.

5.1. Sandia Flame D

The Sandia Flame D is a piloted CH_4 /Air jet flame operating at a jet Reynolds number, $Re = 22,400$. All the details about this flame and the burner geometry can be found at [25]. Here we mention only some of the important aspects of this flame.

The jet fluid consists of 25% CH_4 and 75% air by volume. The jet flows in at 49.6 m/s velocity at 294 K temperature and 0.993 atm pressure. The jet diameter is $D = 7.2$ mm. The pilot is a lean (equivalence ratio, $\Phi = 0.77$) mixture of C_2H_2 , H_2 , air, CO_2 , and N_2 with the same nominal enthalpy and equilibrium composition as that of CH_4 /Air at this equivalence ratio. The pilot velocity is 11.4 m/s. The coflow is air flowing in at 0.9 m/s at 291 K and 0.993 atm.

5.2. Computational details

We perform LES/PDF simulation of the Sandia Flame D using the coupled LES/HPDF solver. The simulation is performed in a cylindrical coordinate system. A computational domain of $80D \times 30D \times 2\pi$ is used in the axial, radial and azimuthal directions, respectively. A non-uniform structured grid of size $192 \times 192 \times 96$ (in the axial, radial and azimuthal directions, respectively) is used for the LES solver. In the HPDF solver, the number of particles per LES cell (N_{pc}) used is $N_{pc} = 40$. With a total of $192 \times 192 \times 96 \approx 3.5 \times 10^6$ LES cells, an overall 140×10^6 particles are used in the computational domain. The simulations are performed on 1024 cores using a domain decomposition of 64×16 in the axial and radial directions, respectively. All the simulations are performed on the Texas Advanced Computing Center (TACC) Ranger cluster.

The chemistry is represented using the combined dimension reduction (using RCCE) and tabulation (using ISAT) approach with a C_1-C_4 skeletal mechanism [37] involving $n_s = 38$ species composed of $n_e = 5$ elements. This mechanism is developed especially for ethylene combustion, but

is also applicable to methane flames. In the future, we want to apply the methodology developed here to study ethylene combustion.

The RCCE dimension reduction is performed by specifying $n_{rs} = 10$ represented species (which is found to be a good number of represented species to achieve less than 2% dimension reduction error based on our previous tests with chemical mechanisms involving around 30 species [14,15]), and so the reduced representation has a dimension $n_r = n_{rs} + n_e = 15$. This dimension reduction from $n_s = 38$ to $n_r = 15$ results in a 60% reduction in the storage needed for particle composition and an 84% reduction in the storage per ISAT table entry. In this preliminary study, we specify the represented species manually (to include the major species of interest for which statistics had already been collected in some of our previous LES/PDF simulations and for which experimental data is available). However, in future studies with bigger mechanisms we will use GALI [15] to select the represented species. In this work, we use the following 10 species as the represented species: CH_4 , O_2 , CO_2 , H_2O , CO , H_2 , OH , H , O and HO_2 .

A fixed ISAT error tolerance, $\epsilon_{tol} = 10^{-4}$, is used in this study. At this error tolerance, the ISAT *tabulation error* relative to direct evaluation (as defined in [15]) is found to be less than 3%. In addition, we specify a maximum ISAT table size of 1000 MB per core. In simulations with the 38-species full representation, some ISAT tables become completely filled. However in simulations with the combined ISAT/RCCE approach with 10 represented species, none of the ISAT tables have a size of more than 200 MB.

6. Results

In this section we compare the computational time and statistics of thermochemical quantities obtained using the combined dimension reduction and tabulation approach with 10 represented species relative to using tabulation alone with the 38-species C_1 – C_4 skeletal mechanism.

In order to make the comparisons, we perform separate LES/PDF simulations of the Sandia Flame D on 1024 cores with chemistry represented using the following two methods

1. ISAT: tabulation alone (no dimension reduction) with the 38-species full representation; and
2. ISAT/RCCE: combined dimension reduction and tabulation with a reduced representation involving 10 represented species (specified in the previous section) and 5 elements.

In each case, we perform LES/PDF simulation to reach a statistically stationary state. We then collect statistics for thermochemical quantities

like temperature, density, and species mass fractions time-averaged over 10,000 time steps. In addition, in each case we perform simulations for 2000 time steps using the three parallel strategies PLP, URAN and P-URAN to compare their relative performance. These simulations start from the statistically stationary state with empty ISAT tables.

6.1. Comparison of statistics

In this section we compare the radial profiles of mean and standard deviation statistics of thermochemical quantities obtained from the PDF particle data from the LES/PDF simulation using ISAT alone and using the combined ISAT/RCCE approach.

The radial statistics are azimuthally-averaged at each time step, and are also time-averaged over 10,000 time steps after reaching the statistically stationary state. For a quantity ξ , we denote the density-weighted mean statistics by $\langle \xi \rangle$, and the standard deviation by $\langle \xi'' \rangle$ which is defined as

$$\langle \xi'' \rangle \equiv \sqrt{\langle \xi^2 \rangle - \langle \xi \rangle^2}.$$

In Figs. 1 and 2, we show respectively the radial profiles of mean and standard deviation of temperature T , density ρ , and mass fraction of species CH_4 , O_2 , CO_2 , H_2O , CO , H_2 , OH at axial locations $x/D = 15, 30, 45, 60$ obtained from (i) an LES/PDF simulation using ISAT alone with the 38-species full representation; (ii) an LES/PDF simulation using ISAT/RCCE with 10 represented species; and (iii) experimentally measured statistics [25] (for reference).

We notice that the statistics obtained with ISAT/RCCE using 10 represented species match very closely with the statistics obtained using ISAT alone with the 38-species full representation. To quantify the difference between the statistics obtained using ISAT/RCCE and ISAT alone, for each quantity ξ (mean or standard deviation), we compute the normalized root-mean-square difference (RMSD) denoted by $\epsilon(\xi)$ as follows

$$\epsilon(\xi) = \frac{[\xi^r - \xi^f]_{rms}}{\xi_{ref}}, \quad (3)$$

where ξ^r and ξ^f denote respectively the quantities obtained using the reduced representation with ISAT/RCCE and the full representation with ISAT alone; and the operator $[\]_{rms}$ denotes the RMSD computed over all the radial locations at all the considered axial locations $x/D = 15, 30, 45, 60$. Here ξ_{ref} is a reference value used for normalization, which is taken to be 1000 K for temperature and 1 kg/m^3 for density. For the species mass fractions, we take ξ_{ref} to be the maximum value of the mean statistics obtained for that species, $\max(\langle \xi \rangle)$.

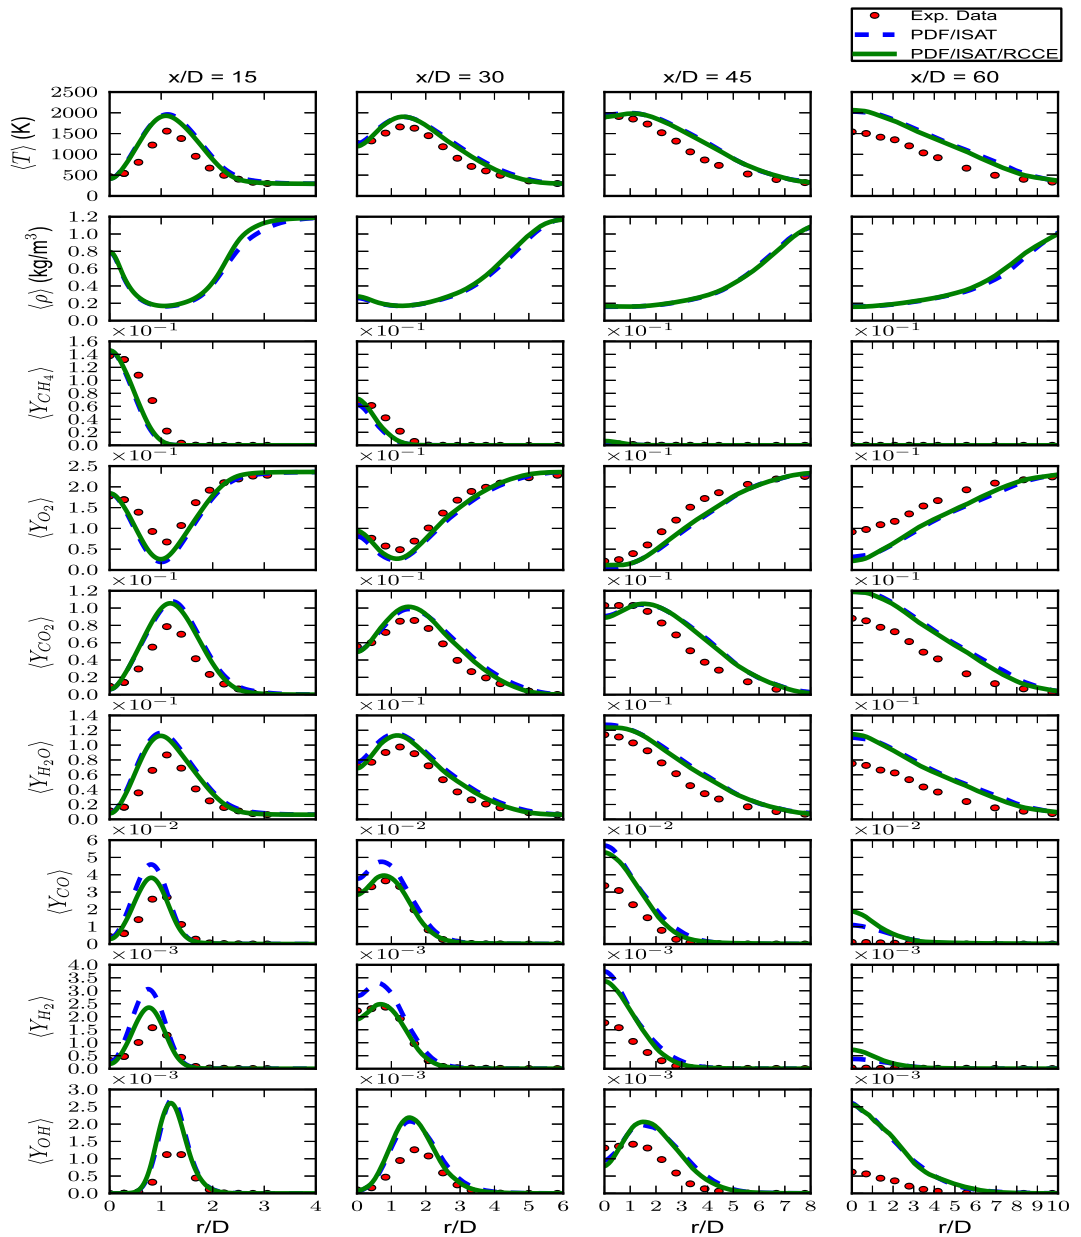


Fig. 1. Radial profiles of time-averaged mean temperature T , density ρ , and mass fraction of species CH_4 , O_2 , CO_2 , H_2O , CO , H_2 , OH at axial locations $x/D = 15, 30, 45$ and 60 obtained from (i) experimental data; (ii) an LES/PDF simulation using ISAT alone with the 38-species full representation; and (iii) an LES/PDF simulation using ISAT/RCCE with 10 represented species.

The reference value and the normalized RMSD computed using Eq. (3) for all the quantities of interest is summarized in Table 1. We notice that the normalized RMSD in the mean and standard deviation statistics is less than 3% (i.e. 30 K) for temperature; less than 2% (i.e. 0.02 kg/m^3) for density; less than 2.5% for species mass fractions of major species CH_4 , O_2 , CO_2 ,

H_2O ; and less than 8% for species mass fractions of minor species CO , H_2 , OH . In summary, these results show that the combined ISAT/RCCE approach shows good error control and the predicted statistics are well within acceptable level of accuracy (relative to using ISAT alone with the full representation) for most engineering applications. These results also show that the

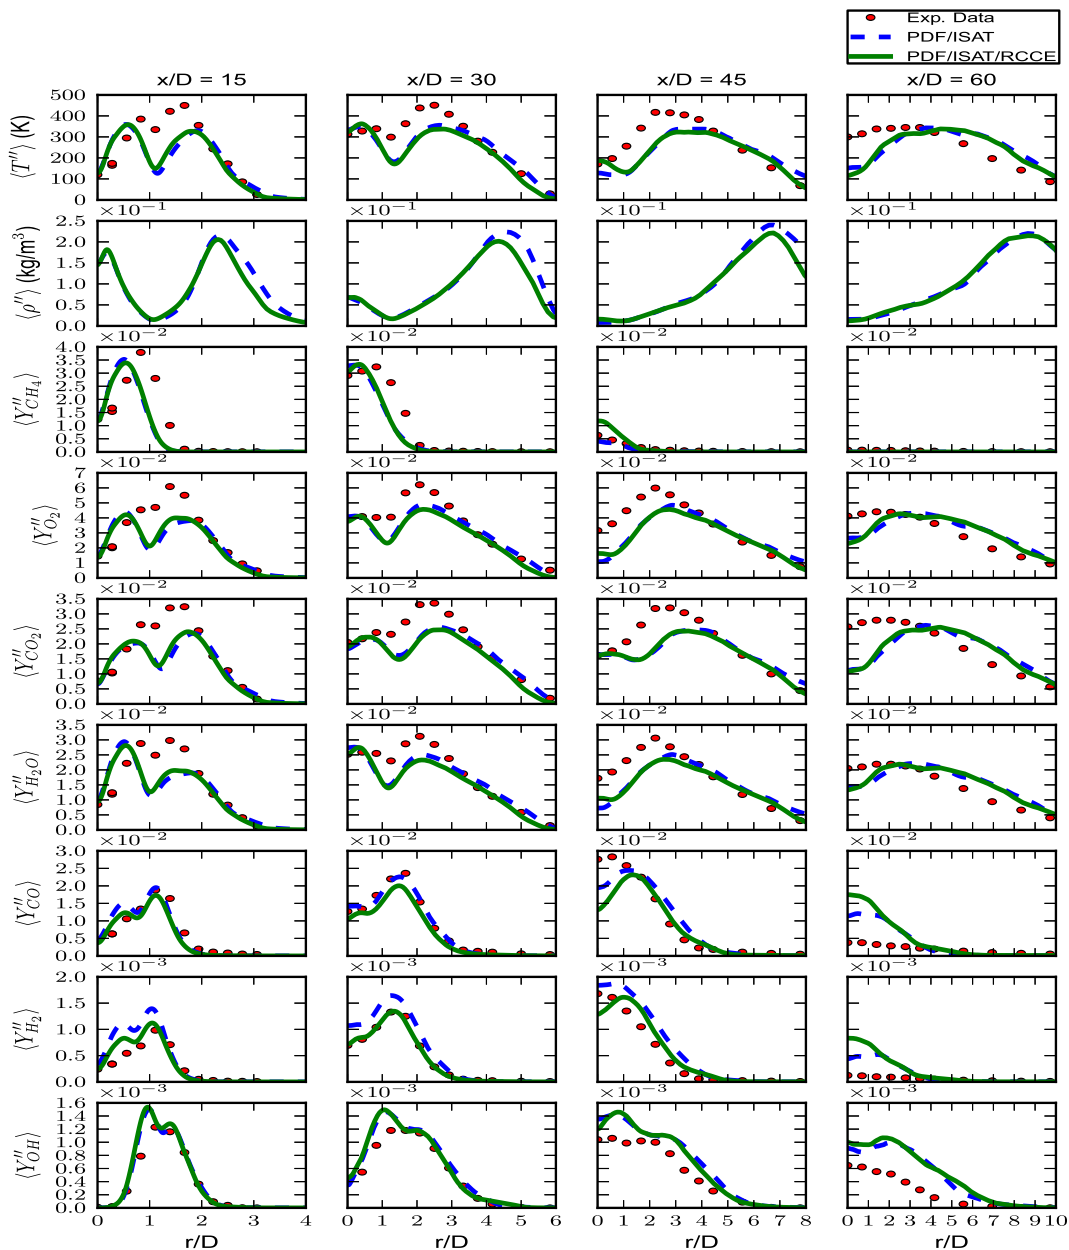


Fig. 2. Radial profiles of time-averaged standard deviation of temperature T , density ρ , and mass fraction of species CH_4 , O_2 , CO_2 , H_2O , CO , H_2 , OH at axial locations $x/D = 15, 30, 45$ and 60 obtained from (i) experimental data; (ii) an LES/PDF simulation using ISAT alone with the 38-species full representation; and (iii) an LES/PDF simulation using ISAT/RCCE with 10 represented species.

density and temperature approximation method used with the reduced representation in ISAT/RCCE [15] yields values that match closely with those obtained with the full representation. A more careful selection of represented species using GALI [15] should help further reduce the differences between the reduced and full descriptions.

The experimentally measured statistics are qualitatively well captured by the LES/PDF simulation, yet quantitatively we notice that some of the statistics differ by around 20%. The discrepancies between the LES/PDF simulation statistics and experimentally measured statistics can be attributed to (i) numerical and statistical errors

Table 1

Normalized root-mean-square difference (RMSD) (see Eq. (3)) in mean and standard deviation statistics obtained using the reduced representation with ISAT/RCCE relative to full representation with ISAT alone. The quantities listed are temperature T , density ρ , and mass fraction of species CH_4 , O_2 , CO_2 , H_2O , CO , H_2 , OH .

| Quantity | Reference value | RMSD in $\langle \xi \rangle$ | RMSD in $\langle \xi'' \rangle$ |
|-----------------------------|----------------------|-------------------------------------|---------------------------------------|
| ξ | ξ_{ref} | $\epsilon(\langle \xi \rangle)$, % | $\epsilon(\langle \xi'' \rangle)$, % |
| T (K) | 1000.0 | 2.92 | 1.95 |
| ρ (kg/m ³) | 1.0 | 1.72 | 1.07 |
| Y_{CH_4} | 1.5×10^{-1} | 1.47 | 1.17 |
| Y_{O_2} | 2.4×10^{-1} | 2.10 | 0.93 |
| Y_{CO_2} | 1.2×10^{-1} | 1.50 | 0.81 |
| Y_{H_2O} | 1.3×10^{-1} | 2.24 | 0.92 |
| Y_{CO} | 5.7×10^{-2} | 6.36 | 4.10 |
| Y_{H_2} | 3.7×10^{-3} | 7.87 | 5.29 |
| Y_{OH} | 2.7×10^{-3} | 2.18 | 1.92 |

in the simulation; (ii) experimental measurement errors; and (iii) errors in the chemical kinetic models. However, study of these errors is not the primary focus of this work. Similar level of agreement between the simulated and experimentally measured statistics is found in some of the previous studies of Sandia Flame D [38–41].

6.2. Computational performance

In this section we compare the wall clock time required to perform LES/PDF simulation of Sandia Flame D for 2000 time steps using the combined ISAT/RCCE approach relative to using ISAT alone. In addition we compare the relative performance of the PLP, URAN and P-URAN parallel strategies. In each case, the LES/PDF simulation is started from a fixed statistically stationary state with empty ISAT tables. We measure a moderate ISAT *build time* (see [15]) of about 1 hour in these simulations, i.e., after 1 hour of simulation, most of the particles are resolved by ISAT retrieves.

In Fig. 3, the bottom three bars show the wall clock time taken to perform 2000 simulation time steps using the combined ISAT/RCCE approach with 10 represented species with the PLP, URAN and P-URAN[0.2 h, 32] parallel strategies. In each case, we also show the breakdown of time spent in LES, HPDF (outside reaction), Reaction (including $x2f_mpi$ communication), and Waiting (average idle time) as defined in [22]. We see that the P-URAN strategy yields the lowest wall clock time among the three strategies. The Waiting time (average idle time), which is indicative of the *load imbalance* is maximum for PLP, minimum for URAN and moderate for P-URAN (as also seen in our previous studies [22]).

The LES/PDF simulation of Sandia Flame D using ISAT alone with the 38-species full

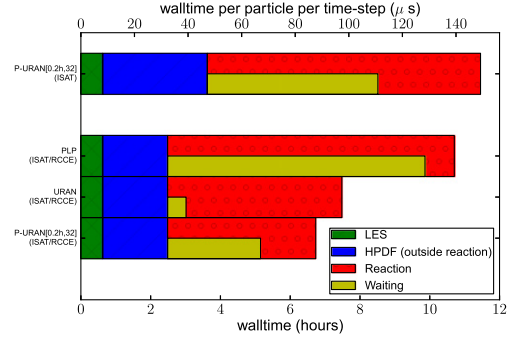


Fig. 3. For LES/PDF simulation of Sandia Flame D, wall clock time for 2000 time steps along with breakdown of time spent in LES, HPDF (outside reaction), Reaction (including $x2f_mpi$ communication) and Waiting (average idle time) using different parallel strategies. Top: using ISAT alone with the 38-species full representation with the P-URAN[0.2 h, 32] parallel strategy. Bottom three: using combined ISAT/RCCE with 10 represented species using (i) PLP; (ii) URAN; and (iii) P-URAN[0.2 h, 32] parallel strategies.

representation is performed using the PLP, URAN and P-URAN strategies [22]. Among these the P-URAN strategy again yields the lowest wall clock time. In Fig. 3, for comparison, the top bar shows the wall clock time for 2000 time steps using the 38-species full representation with the P-URAN[0.2 h, 32] strategy.

We notice that relative to the simulation using ISAT alone with the 38-species full representation, the combined ISAT/RCCE approach with 10 represented species using the P-URAN strategy

1. yields more than 40% reduction in HPDF time (outside Reaction). This is because with dimension reduction, the particles in PDF simulation carry only the reduced representation (in this case involving 15 variables). As a result, a) particles require 60% less storage, which in turn reduces the particle communication cost; and b) less time is required for collecting species (LES cell mean) statistics.
2. reduces the Reaction time by over 40% due to smaller ISAT table sizes and faster retrieve times (statistics given in Table 2 and explained below); and
3. consequently, reduces the overall wall clock time of the simulation by more than 40%.

In Table 2, we list the ISAT statistics collected from the simulations with the P-URAN strategy using a) ISAT alone with the 38-species full representation; and b) ISAT/RCCE with 10 represented species. We see that in both the cases, over 99.9% of the queries result in retrieves which shows the high efficiency of ISAT tabulation. Compared to ISAT/RCCE with 10 represented

Table 2

Cumulative ISAT statistics for the LES/PDF simulation of Sandia Flame D using (i) ISAT alone and (ii) ISAT/RCCE with the P-URAN parallel strategy.

| Method | Variables | Queries | Retrieves (%) | Adds (%) | Direct evals* (%) | Retrieve time (μ s) |
|-----------|-----------|--------------------|---------------|------------------------|------------------------|--------------------------|
| ISAT | 38 | 3×10^{11} | 99.974 | 5.572×10^{-3} | 1.665×10^{-3} | 8 |
| ISAT/RCCE | 15 | 3×10^{11} | 99.984 | 2.742×10^{-3} | 0.0 | 3 |

* Performed only if the ISAT table is completely filled.

species, the simulation with 38-species full-representation results in almost twice the number of adds, and also results in some direct evaluations because some of the tables get completely filled. The average retrieve time with ISAT/RCCE is only 3 μ s compared to 8 μ s with ISAT alone.

7. Conclusions

We have successfully extended our LES/PDF solver with the capability of performing turbulent combustion calculations with realistic combustion chemistry, wherein the chemistry in the PDF solver is represented using the combined dimension reduction (using RCCE) and tabulation (using ISAT) approach. The chemistry workload is efficiently redistributed using the P-URAN strategy implemented using the *x2f_mpi* library. We have shown that for performing LES/PDF simulation of Sandia Flame D, relative to using ISAT alone with the 38-species full representation, the ISAT/RCCE approach with 10 represented species (i) yields acceptable level of accuracy in mean and standard deviation statistics of major thermochemical quantities of interest like temperature, density and species mass fractions; and (ii) reduces the overall simulation wall clock time by more than 40% with the P-URAN strategy.

Acknowledgments

V.H.'s work on dimension reduction methodologies is supported by Office of Energy Research, Office of Basic Energy Sciences, Chemical Sciences, Geosciences and Biosciences Division of the US Department of Energy (DOE) under Grant No. DE-FG02-90ER. The work of V.H. and S.R.L. on the ethylene mechanism is supported by Grant No. FA9550-09-1-0611 funded by the National Center for Hypersonic Combined Cycle Propulsion, sponsored by the AFOSR and NASA ARMD. S.R.L.'s initial work on this project was supported by NASA Grant No. NNX08AB36A. This research was also supported in part by the National Science Foundation through TeraGrid resources provided by the Texas Advanced Computing Center under Grant No. TG-CTS090020. S.B.P. has a financial interest in Ithaca Combustion Enterprise, LLC., which

has licensed the software ISAT-CK and CEQ used in this work.

References

- [1] C.K. Westbrook, W.J. Pitz, O. Herbinet, H.J. Curran, E.J. Silke, *Combust. Flame* 156 (2009) 181–199.
- [2] S. Sarathy, C. Westbrook, M. Mehl, W. Pitz, C. Togbe, P. Dagaut, et al., *Combust. Flame* 158 (12) (2011) 2338–2357.
- [3] T. Lu, C.K. Law, *Proc. Combust. Inst.* 30 (2005) 1333–1341.
- [4] P. Pepiot-Desjardins, H. Pitsch, *Combust. Flame* 154 (2008) 67–81.
- [5] T. Nagy, T. Turanyi, *Combust. Flame* 156 (2009) 417–428.
- [6] J.C. Keck, D. Gillespie, *Combust. Flame* 17 (1971) 237–241.
- [7] S.H. Lam, D.A. Goussis, *Int. J. Chem. Kinet.* 26 (1994) 461–486.
- [8] U. Maas, S.B. Pope, *Combust. Flame* 88 (1992) 239–264.
- [9] Z. Ren, S.B. Pope, A. Vladimirov, J.M. Guckenheimer, *J. Chem. Phys.* 124 (2006). Art. No. 11411.
- [10] S.B. Pope, *Combust. Theory Model.* 1 (1997) 41–63.
- [11] L. Lu, S.B. Pope, *J. Computat. Phys.* 228 (2) (2009) 361–386.
- [12] S.R. Tonse, N.W. Moriarity, N.J. Brown, M. Frenklach, *Israel J. Chem.* 39 (1999) 97–106.
- [13] T. Turanyi, *Comput. Chem.* 18 (1994) 45–54.
- [14] V. Hiremath, Z. Ren, S.B. Pope, *Combust. Theory Model.* 14 (5) (2010) 619–652.
- [15] V. Hiremath, Z. Ren, S.B. Pope, *Combust. Flame* 158 (11) (2011) 2113–2127.
- [16] Z. Ren, G.M. Goldin, V. Hiremath, S.B. Pope, *Combust. Theory Model.* 15 (6) (2011) 827–848.
- [17] Q. Tang, S.B. Pope, *Proc. Comb. Inst.* 29 (2002) 1411–1417.
- [18] L. Lu, Z. Ren, V. Raman, S.B. Pope, H. Pitsch, Proceedings of the CTR Summer Program, 2004, pp. 283–294.
- [19] L. Lu, Z. Ren, S.R. Lantz, V. Raman, S.B. Pope, H. Pitsch, Proceedings of the 4th Joint Meeting of US Sections of the Combustion Institute, Philadelphia, PA, 2005.
- [20] L. Lu, S.R. Lantz, Z. Ren, S.B. Pope, *J. Computat. Phys.* 228 (15) (2009) 5490–5525.
- [21] Y. Shi, W.H. Green, H.-W. Wong, O.O. Oluwole, *Combust. Flame* 158 (5) (2011) 836–847.
- [22] V. Hiremath, S.R. Lantz, H. Wang, S.B. Pope, *Combust. Flame*, in press <http://dx.doi.org/10.1016/j.combustflame.2012.04.013>.
- [23] J.C. Keck, *Prog. Energy Combust. Sci.* 16 (1990) 125–154.

- [24] W. Jones, S. Rigopoulos, *Combust. Flame* 142 (3) (2005) 223–234.
- [25] R.S. Barlow, J.H. Frank, *Proc. Combust. Inst.* 27 (1998) 1087–1095, <http://www.sandia.gov/TNF/DataArch/FlameD.html>.
- [26] H. Wang, S.B. Pope, *Proc. Combust. Inst.* 33 (2011) 1319–1330.
- [27] C.D. Pierce, *Progress-variable approach for Large-Eddy Simulation of turbulent combustion*, Ph.D. thesis, Stanford University, 2001.
- [28] C.D. Pierce, P. Moin, *J. Fluid Mech.* 504 (2004) 73–97.
- [29] H. Wang, P.P. Popov, S.B. Pope, *J. Computat. Phys.* 229 (2010) 1852–1878.
- [30] P. Kloeden, E. Platen, Springer Verlag, Berlin, 1992.
- [31] J. Janicka, W. Kolbe, W. Kollmann, *J. Non-Equilib. Thermodyn.* 4 (1970) 47–66.
- [32] M. Caracotsios, W.E. Stewart, *Comput. Chem. Eng.* 9 (1985) 359–365.
- [33] G.P. Beretta, J.C. Keck, M. Janbozorgi, H. Metghalchi, *Entropy* 14 (2) (2012) 92–130.
- [34] S.B. Pope, *Combust. Flame* 139 (2004) 222–226.
- [35] M. Janbozorgi, H. Metghalchi, *Int. J. Thermodyn.* 12 (2009) 44–50.
- [36] M. Janbozorgi, S. Ugarte, H. Metghalchi, J.C. Keck, *Combust. Flame* 156 (2009) 1871–1885.
- [37] G.E. Esposito, H.K. Chelliah, *Combust. Flame* 158 (3) (2011) 477–489.
- [38] M. Sheikhi, T. Drozda, P. Givi, F. Jaberi, S. Pope, *Proc. Combust. Inst.* 30 (2005) 549–556.
- [39] M. Ihme, H. Pitsch, *Combust. Flame* 155 (2008) 90–107.
- [40] R. Mustata, L. Valio, C. Jimnez, W. Jones, S. Bondi, *Combust. Flame* 145 (2006) 88–104.
- [41] M.B. Nik, S.L. Yilmaz, P. Givi, M.R.H. Sheikhi, S.B. Pope, *Am. Inst. Aeronaut. Astronaut* 48 (2010) 1513–1522.