Investigation of strategies for the parallel implementation of ISAT in LES/FDF/ISAT computations

L. Lu [1a], Z. Ren [a], S. R. Lantz [b], V. Raman [c], S. B. Pope [a], H. Pitsch [c]
[a] Sibley School of Mechanical and Aerospace Engineering, Cornell University, NY
[b] Cornell Theory Center, Cornell University, NY
[c] Center for Turbulence Research, Stanford University, CA

Abstract

The LES/FDF approach for turbulent combustion offers the benefits of both large eddy simulation (LES) to treat the turbulent flow, and the PDF approach to treat turbulence-chemistry interactions (in terms of the filtered density function, FDF). The approach is implemented as a particle mesh method and computationally the most expensive aspect is determining the change in particle composition over a time step due to reaction. This cost can be significantly reduced by using *in situ* adaptive tabulation (ISAT). In this work we investigate the computational performance of several strategies for the parallel implementation of ISAT in LES/FDF calculations. The capability of performing LES/FDF/ISAT computations of turbulent flames is developed by incorporating the ISAT algorithm in the Stanford structured large eddy simulation (LES) and composition "filtered density function" (FDF) code. The LES/FDF/ISAT simulation of a spatially developing mixing layer is used as the test case to study the performance and load balancing of different ISAT strategies for idealized turbulent flames of both hydrogen and methane. Detailed 9-species and 35-species mechanisms are employed for the hydrogen flame and the methane flame, respectively. The results show that when it is almost always possible to retrieve from the ISAT table, then using purely local processing (without any message passing) is optimal. But when a significant number of direct integrations of the chemical kinetic equation is required, then parallel strategies, such as the uniform random distribution (URAN) strategy, are advantageous. Finally, a simple model is developed to explain the observed computational performance of the different parallel strategies in different simulations.

1. Introduction

For the modeling of turbulent combustion, large eddy simulation (LES) has the advantage of accounting for the large scale unsteady motions, which account for the bulk of the spatial transport. Within the past decade, large eddy simulation of turbulent reacting flows has been the subject of widespread investigation. A significant recent development of the subgrid scale (SGS) modeling is based on the "filtered density function" (FDF) [1] which provides a complete description of the subfilter-scale composition fluctuations. Therefore the effect of chemical reactions appears in closed form and the FDF offers the ability to treat finite-rate chemistry and the turbulence-chemistry interactions. The composition FDF is considered by many investigators [2–8] and this combined methodology LES/FDF (with very simple chemistry) has been applied to the Sandia D flame by Sheikhi *et al.* [9], and to the Sydney bluff-body flame by Raman *et al.* [10].

The computational cost of directly using realistic combustion chemistry in the LES/FDF simulation of turbulent combustion is expensive (and sometimes prohibitive) because of the strong nonlinearities of the reaction source term and the wide range of time scales in the chemical kinetics. The *in situ* adaptive tabulation (ISAT) algorithm [11] greatly facilities the incorporation of realistic combustion chemistry (e.g., of order 10 and 20 species for hydrogen and methane, respectively) and reduces the CPU time dramatically, which makes the LES/FDF simulation of turbulent combustion with realistic chemistry feasible. Different ISAT parallel strategies are implemented using the *x2f_mpi* software developed by Pope & Lantz [12].

The parallel LES/FDF/ISAT capability for the

[1]Corresponding author: ll267@cornell.edu

numerical simulations of turbulent reactive flows with realistic combustion chemistry is developed ([13]) by incorporating ISAT into the Stanford structured LES [14, 15] code with a Lagrangian Monte Carlo implementation of the composition FDF method [16]. Also a preliminary study of the parallel LES/FDF/ISAT capability has been carried out [13].

The LES/FDF code employs domain decomposition, and hence the particles are on different processors. The computational goal is to complete the reaction fractional step in the shortest wall clock time. In this work, the parallel LES/FDF/ISAT simulation of a spatially developing mixing layer is used as the test case to study the performance of different ISAT strategies for idealized turbulent flames of both hydrogen and methane. For the hydrogen flame, we use the detailed mechanism of Li [17], which has 9 species. For the methane flame, the GRI3.0 mechanism [18] without nitrogen chemistry is employed. This is generated from the standard GRI3.0 mechanism by stripping out all the N containing species and reactions except for N_2. In the next sections, the ISAT implementation in the parallel LES/FDF/ISAT capability is briefly reviewed. Then the performance of different ISAT strategies in the mixing layer simulations is discussed and investigated.

2. In situ adaptive tabulation (ISAT) implementation

In the parallel LES/FDF simulations of turbulent combustion, if domain decomposition is employed in the parallel computation, each processor is assigned the computational work of a subdomain. Moreover, if a splitting scheme is used to separate the combustion chemistry from other processes such as convective and diffusive species transport, the reaction fractional step corresponds to the reaction of a homogeneous, adiabatic, isobaric system consisting of n_s chemical species. At time t, the composition is represented by the $n_\phi = n_s + 1$ variables $\boldsymbol{\phi}(t) \equiv \{\phi_1, \phi_2, \ldots, \phi_{n_\phi}\}$, which can be taken to be the species specific moles and enthalpy. Modeling reaction kinetics in this system leads to a set of stiff ordinary differential equations (ODEs)

$$\frac{d\boldsymbol{\phi}(t)}{dt} = \mathbf{S}(\boldsymbol{\phi}(t)), \qquad (1)$$

where $\mathbf{S}$ is the rate of change of composition given by the detailed chemical mechanism (and the conservation of enthalpy). The task in the reaction fractional step is to determine the solution to Eq. 1, the composition $\boldsymbol{\phi}(t + \triangle t)$ after reaction has occurred for the time-step interval Δt starting from the initial composition $\boldsymbol{\phi}(t)$ at time t. Direct integration (DI) of Eq. 1 using an ODE integrator is computationally expensive. Consequently, when the reaction fractional step is performed a huge number of times with many different initial compositions such as in the LES/FDF simulations, significant computational resources are required.

The alternative to DI for computing the solution $\boldsymbol{\phi}(t + \triangle t)$ is ISAT, which is a storage-retrieval methodology (with error control) introduced by Pope (1997). An ISAT table stores pairs of values of $\boldsymbol{\phi}$ before and after reaction $\{\boldsymbol{\phi}(0), \boldsymbol{\phi}(\Delta t)\}$, so that given $\boldsymbol{\phi}(t)$ the corresponding value $\boldsymbol{\phi}(t + \Delta t)$ can be retrieved from the table. (Since the governing ordinary differential equations (ODEs) are autonomous, the value of t is immaterial, and without loss of generality we can set $t = 0$.) As needed, pairs of values $\{\boldsymbol{\phi}(0), \boldsymbol{\phi}(\Delta t)\}$ are added to the table by integrating the governing ODEs.

Different ISAT parallel strategies have been developed and these strategies are implemented using the *x2f_mpi* software. In the parallel LES/FDF simulation of a combustion problem, each processor has its own ISAT table. For each processor, in the reaction fractional step, the initial compositions of all the particles in the subdomain are passed to the *x2f_mpi* software. In *x2f_mpi*, different message-passing strategies are implemented, which determine how the particles compositions are processed by the ISAT tables on the different processors. In this study, we focus on two strategies: purely local processing (LOCAL) and uniform random distribution (URAN). For the purely local processing strategy, the particles on each processor are locally processed by using the local ISAT table and there is no message passing. For the uniform random distribution strategy, all the particles on each processor are randomly distributed uniformly among all the processors by the *x2f_mpi* software (using MPI) and processed there using the local ISAT table. Then the compositions after reaction are passed back to the original processor by *x2f_mpi*. Hence the URAN strategy requires a large amount of message passing.

3. Flow configuration and numerical specification

The LES/FDF/ISAT simulation of ideal two-dimensional, spatially developing, reacting, plane mixing layers is considered. In the simulations, the density is taken to be constant and therefore there is no feedback from the FDF to the LES computation. The primary objective of the simulation is to investigate the performance of different ISAT strategies. No attempt is made to study numerical accuracy, to investigate the effect of models and model parameters, or to make comparisons with experimental data.

In the simulations, the reacting planar mixing layer is composed of cold diluted hydrogen in one stream (fuel stream) and hot lean combustion product in the other stream (oxidant stream). For the hydrogen flame, the composition for the fuel stream is taken to be (in volume percentage) N_2 (75%) and H_2 (25%); the temperature is 305K. The composition for the oxidant stream is taken to be (in volume percentage) N_2 (75%), O_2 (15%) and H_2O (10%); the temperature is 1045K. For the methane flame, the composition for the fuel stream is taken to be (in volume percentage) N_2 (52.9%), O_2 (14.1%) and

CH_4 (33%); the temperature is 320K. The composition for the oxidant stream is taken to be (in volume percentage) N_2 (73.1%), O_2 (12%) and H_2O (14.9%); the temperature is 1500K.

In the simulations, the flow variables are normalized with respect to selected reference quantities, denoted by the subscript r. The reference velocity U_r is chosen to be the fuel stream velocity 66 m/s, the reference length L_r is taken to be 3.3mm, and the non-dimensional time is given by $t^* = U_r t/L_r$. The kinematic viscosity ν_r in the simulation is taken to be a constant throughout, which is based on the fuel stream properties. The Reynolds number $Re = U_r L_r/\nu_r$ is 15200. The two-dimensional computational domain in non-dimensional units in the simulation is the rectangular region, where x^* from -10 to 50 and y^* from 0 to 15. The splitter plate lies in the region, where x^* from -10 to 0 and y^* from 7.2 to 7.8. The present two-dimensional calculations are performed on a nonuniform mesh with 128 grid nodes in the streamwise direction (x^*) and 64 grid nodes in the transverse direction (y^*). The mesh in the small region near the splitter plate is refined. The simulation uses parallel computation with 8 processors, and the computational domain is decomposed into 8 blocks in the x^* direction with each block having the same number of grid cells.

In the simulation, the inlet velocity profiles at x^*=-10 are taken to be laminar plug profiles and turbulence develops in the downstream layer due to the instability of the shear layer. The ratio of oxidant and fuel stream velocities is fixed at 0.4, with the non-dimensional fuel stream velocity being unity, and oxidant stream velocity being 0.4. These values of the free stream velocities are selected so that the flame exists throughout most of the flow. The wall boundary conditions at the splitter plate are Neumann conditions for all scalars and pressure, and no-slip Dirichlet conditions for velocity. All the open and outflow boundaries use the convective boundary condition.

In FDF, the Monte Carlo particles are distributed at t=0 throughout the computational domain. The nominal particle number density of 30 particles per cell is used in the simulations. Particle cloning and clustering are utilized to minimize the fluctuations in the number density. New particles are introduced at the inlet at a rate proportional to the local mass flow rate, and the composition of incoming particles is set according to the composition of the fluid at the point of entry.

4. ISAT performance and discussion

In parallel LES/FDF/ISAT simulations of turbulent combustion, the reaction fractional step is usually the most time-consuming. Therefore the computational goal is to complete the reaction fractional step in the shortest wall clock time. The objective of this work is to investigate and compare the performance of different ISAT strategies (which are used to treat combustion chemistry in the reaction fractional step) under different parameter regimes.

Table 1: Parameters in different simulations: the ISAT error tolerance ϵ, the maximum assigned storage, the fraction of direct integration P_{DI}, the nonuniformity factor λ, the average CPU time per retrieve t_R, the average communication time per particle composition t_C, and the average CPU time per direct integration t_{DI}.

	H_2-1	H_2-2	CH_4-1
ϵ	1×10^{-5}	5×10^{-6}	1×10^{-5}
storage (Mbyte)	300	20	50
$P_{DI}(\%)$	0	0.31	0.5
λ	1	4.6	4.2
t_R (μs)	8	8	25
t_C (μs)	4	4	16
t_{DI} (μs)	1600	1600	24300

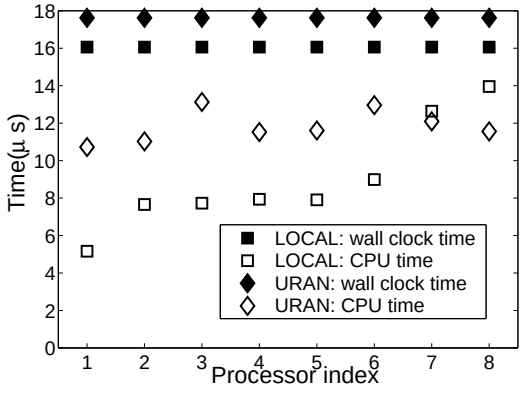

Fig. 1: The CPU time and wall clock time in reaction fractional step (in microseconds per particle step) for each processor for case H_2-1.

In order to study the ISAT performance in the statistically stationary period, the full LES/FDF/ISAT simulations are carried out, restarting from the statistically stationary flow and composition fields. In this section, the ISAT performances from three different LES/FDF/ISAT simulations (23 hours run with constant time step $dt^* = 0.06$) are presented and analyzed. The corresponding parameters (e.g., error tolerance ϵ, the maximum assigned storage) for ISAT are listed in Table 1. Figures 1, 2 and 3 show the measured CPU time and wall clock time (in microseconds per particle step) in the reaction fractional steps for the three cases.

For the case H_2-1, a large ISAT table storage and a large ISAT error tolerance are specified and hence cheap retrieves are dominant and there are no direct integrations in the simulation (see Table 1). As may be seen from Fig. 1, if the purely local processing

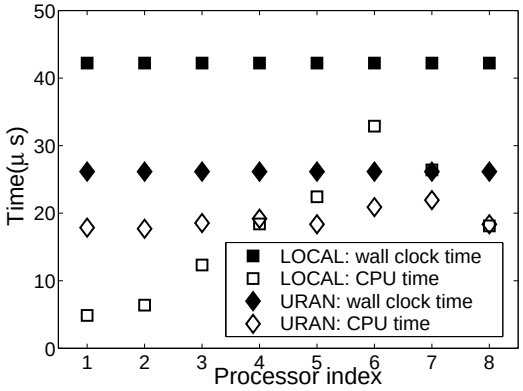

Fig. 2: The CPU time and wall clock time in reaction fractional step (in microseconds per particle step) for each processor for case H_2-2.

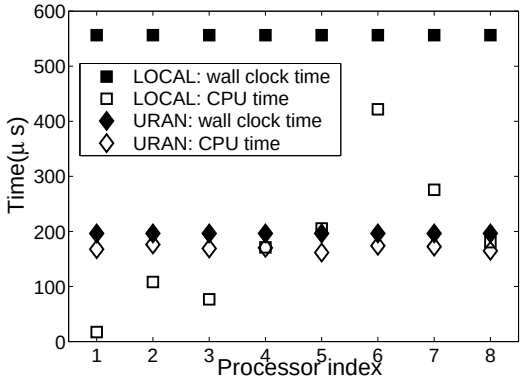

Fig. 3: The CPU time and wall clock time in reaction fractional step (in microseconds per particle step) for each processor for case CH_4-1.

strategy (LOCAL) is used, even for this simple geometry and chemistry case, the load balancing is not good. The CPU time spent by processor 8 is almost 3 times the CPU time spent by processor 1. Accordingly, processor 1 has a significant amount of idle time. This is due to the following fact: in the active flame region there are strong chemical reactions and so the reaction fractional step is usually computationally expensive. In the computational domain of processor 1 (adjacent to the inlet), there is no flame at all, whereas for processor 8 (adjacent to the outlet), there is significant flame activity. (For a more complicated flame, such as the methane flame in this study, much worse load balancing is observed when the purely local processing strategy is employed.) In contrast to the purely local processing strategy, the uniform random distribution (URAN) strategy achieves much better load balancing. However, the URAN strategy gives a larger wall clock time compared to the LOCAL strategy. For this particular case, as far as the wall clock time is concerned (as it should), the LOCAL strategy has advantage over the URAN strategy even though the URAN mode achieves better load balancing.

Hence load balancing is not the issue: wall clock time is. The optimal ISAT strategy, which minimizes the wall clock time for the combustion chemistry, may not give the best load balancing.

Figures 2 and 3 compare the parallel ISAT performance for both the hydrogen flame H_2-2 and the methane flame CH_4-1. For both these cases, a relatively small ISAT table storage and a small error tolerance ϵ are specified, hence as may be seen from Table 1, there are significant fractions of direct integrations. The fraction of direction integration is 0.31% for case H_2-2 and 0.5% for case CH_4-1. As may be seen from Figs. 2 and 3, if purely local processing is employed, there is a significant load imbalance for both of these cases (this is also implied by the value of the nonuniformity factor λ as discussed later), and the URAN strategy has better load balancing. More important, for the case H_2-2, the URAN strategy is about 1.6 times faster than the LOCAL strategy; and for the case CH_4-1, the URAN strategy is about 2.9 times faster than the LOCAL strategy. Hence for these two cases, the URAN strategy has a clear advantage over the LOCAL strategy.

The relative performance of the above two parallel ISAT strategies (LOCAL and URAN) depends on the following five parameters: the fraction of reaction fractional steps treated by direct integration P_{DI} (averaged over all particle steps on all processors); the average CPU time per retrieve t_R (averaged over all the retrieve events on all processors); the average communication time (two way message passing) per particle composition t_C; the average CPU time per direct integration t_{DI} (averaged over all the direction integration events on all processors); and the nonuniformity factor λ ($\lambda \geq 1$). The parameter λ measures the ratio of the maximum (over processors) of the fractions of reaction fractional steps treated by direct integration (averaged over the particle steps on the processor) to P_{DI} (averaged over all particle steps on all processors). Hence the product of λP_{DI} is the maximum (over processors) of the fractions of reaction fractional steps treated by direct integration (averaged over the particle steps on the processor), and the processor with the maximum of the fractions of direct integrations is generally most time-consuming. The corresponding parameters for the three cases considered are listed in Table 1. As may be observed from Table 1, $t_C/t_R \approx 0.5$ and 0.6, while $t_{DI}/t_R \approx 200$ and 970 for the hydrogen and methane flames, respectively.

Here a simple model is developed to explain the observed performance of the different ISAT strategies in the three simulations considered. Given the parameters, the wall clock time per query for pure local processing (LOCAL) strategy can be approximated as

$$t_L = \lambda P_{DI} \times t_{DI} + (1 - \lambda P_{DI}) \times t_R, \qquad (2)$$

where the first term is the contribution of the di-

Table 2: Measured (from simulations) and predicted (by the model, in parentheses) wall clock time (per query) and γ values.

	H_2-1	H_2-2	CH_4-1
t_L (μs)	16 (8)	42 (31)	556 (535)
t_U (μs)	17 (12)	26 (17)	196 (163)
γ (measured)	0.9	1.6	2.9
γ (predicted)	0.67	1.8	3.3

rect integration time and the second term is the retrieve time. When the uniform random distribution (URAN) strategy achieves good load balancing as expected, the wall clock time per query can be estimated as

$$t_U = P_{DI} \times t_{DI} + (1 - P_{DI}) \times t_R + t_C, \quad (3)$$

where the last term is an estimation for the communication time. Hence the relative performance of these two strategies can be measured by

$$\gamma = \frac{t_L}{t_U} = \frac{\lambda P_{DI} + (1 - \lambda P_{DI}) \times t_R/t_{DI}}{P_{DI} + (1 - P_{DI}) \times t_R/t_{DI} + t_C/t_{DI}}. \quad (4)$$

If $\gamma < 1$, the LOCAL strategy has advantage over the URAN strategy, whereas for $\gamma > 1$, the URAN strategy has advantage.

As may be observed from Eq. 4, if the cheap retrieves are dominant and there are no direct integrations in the simulation, then $\gamma = \frac{t_R}{t_R + t_C} < 1$ and the URAN strategy has no advantages over the LOCAL strategy. (Even though the URAN strategy may achieve better load balancing, the wall clock time for combustion chemistry is larger compared with the LOCAL strategy.) This is the case for case H_2-1 where a large ISAT table storage and a large ISAT error tolerance are specified and the ISAT table does not become full. By contrast, the URAN strategy will have advantage over the LOCAL strategy only when there is a significant fraction of direct integrations in the simulation. This is the case for H_2-2 and CH_4-1, where a small ISAT table storage and a small error tolerance ϵ are specified.

Based on the simple model, we can estimate the wall clock time and the relative performance parameter γ for the three cases considered. Table 2 lists the measured values and the corresponding predicted values by the model. As may be seen from Table 2, the model underpredicts the wall clock time for all the three cases. This is mainly because that the model does not count the idle time, which is the execution time used for waiting. (For the LOCAL strategy, the idle time on each processor is the corresponding difference between the wall clock time and CPU time on each processor. For the URAN strategy, part of the difference between the wall clock time and CPU time is due to idle time.) Moreover the model does not count the time spent on other processes in the ISAT algorithm and assumes a uniform retrieve time for all the ISAT tables on different processors. When the contribution of the direct integration time to the total wall clock time becomes significant, the model gets more accurate, such as in case CH_4-1.

Nevertheless, the simple model gives a reasonable prediction for the performance parameter γ. As may be seen from Table 2, according to the model, we have $\gamma = 0.67$, 1.8 and 3.3 for cases H_2-1, H_2-2 and CH_4-1, respectively. These values are consistent with the measured values (0.9, 1.6 and 2.9) read from the Figs. 1, 2 and 3, i.e., the simple model well predicts the performance of different ISAT strategies.

Based on this simple model (Eq. 4), the URAN strategy has advantage over the LOCAL strategy when the fraction of direct integrations (P_{DI}) satisfies

$$P_{DI} > \frac{t_C/t_{DI}}{(1 - t_R/t_{DI})(\lambda - 1)}. \quad (5)$$

Under normal circumstances where $t_R/t_{DI} \ll 1$ and $\lambda P_{DI} \ll 1$ (this is the case for the current study), Eq. 5 can be further simplified as

$$P_{DI} > \frac{t_C/t_{DI}}{(\lambda - 1)}. \quad (6)$$

The above analysis is based on the assumption that the URAN strategy would achieve good load balancing as theoretically suggested. However, we do observe poor load balancing of URAN strategy for some pathological cases, which is probably due to the complicated structure in the ISAT table. The investigation of this issue will be part of the future work.

5. Conclusion

The LES/FDF/ISAT capability for the numerical simulations of turbulent combustion with realistic chemistry is developed by incorporating ISAT into the Stanford structured LES code with a Lagrangian Monte Carlo implementation of the composition FDF method. Different ISAT parallelization strategies are implemented using the *x2f_mpi* software.

In this paper, the LES/FDF/ISAT simulation of a spatially developing mixing layer is used as the test case to study the performance of two different ISAT strategies for idealized turbulent flames of both hydrogen and methane. Detailed mechanisms with 9-species and 35-species are employed for the hydrogen flame and the methane flame, respectively. The results show that when it is almost always possible to retrieve from the ISAT table, then using purely local processing ISAT (without any message passing) is optimal even though the load balancing may not be good (see case H_2-1).

But when there is a significant fraction of direct integrations in the simulation, uniform random distribution strategy (URAN) is advantageous (see case H_2-2 and CH_4-1).

The relevant parameters for the ISAT performance are identified and a simple model is proposed to explain the observed performance of different ISAT strategies. This model could be used to choose the appropriate one among different ISAT strategies in different parameter regimes. The model prediction of the performance parameter γ is consistent with the simulation results.

In the future work, the quantitative study of the performance of other parallel ISAT strategies in different circumstance will be performed.

Acknowledgements

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. This work is supported by the Air Force Office of Scientific Research under Grant No. F-49620-00-1-0171.

References

[1] POPE, S. B., *Proc. Combust. Inst.* **23**, 591-612 (1990).

[2] GAO, F. AND O'BRIEN E. E., *Phys. Fluids A* **5**, 1282-1284 (1993).

[3] COLUCCI, P.J., JABERI, F. A., GIVI, P. AND POPE, S. B., *Phys. Fluids* **10**, 499-515 (1998).

[4] JABERI, F. A., COLUCCI, P.J., JAMES, S., GIVI, P. AND POPE, S. B., *J. Fluid Mech.* **401**, 85-121 (1999).

[5] ZHOU, X. Y. AND PEREIRA, J. C. F., *Flow, Turbulence and Combustion* **64**, 279-300 (2000).

[6] JAMES, S. AND JABERI, F. A., *Combust. Flame* **123**, 465-487 (2000).

[7] GIVI, P., *AIAA Paper*, 2003-5081 (2003).

[8] RAMAN, V., PITSCH, H. AND FOX, R. O., in: *CTR Annual Research Briefs, Center For Turbulence Research*, 261-275 (2003).

[9] SHEIKHI, M. R. H., DROZDA, T. G., GIVI, P., JABERI, F.A. AND POPE, S. B., *Proc. Combust. Inst.* **30**, (in press 2004).

[10] RAMAN, V., PITSCH, H. AND FOX, R. O., *TNF7 Workshop Proceedings*, 56-57 (2004).

[11] POPE, S. B., *Combust. Theo. Modelling* **1**, 41-63 (1997).

[12] POPE, S. B. AND LANTZ, S. R., *x2f_mpi* Software for the efficient load distribution of function evaluations in a multi-processor MPI enviroment. (in preparation 2004).

[13] LU, L., REN, Z., RAMAN, V., POPE, S. B. AND PITSCH, H., In *Proceedings of the Summer Program 2004, Center For Turbulence Research*, 283-294 (2004).

[14] PIERCE, C. D. , Progress-variable approach for large-eddy simulation of turbulent combustion. PhD thesis, Stanford University (2001).

[15] PIERCE, C. D. AND MOIN, P. *J. Fluid Mech.* **504**, 73-97 (2004).

[16] RAMAN, V., FOX, R. O., HARVEY, A. D. AND WEST, D. H. *Industrial and Engineering Chemistry Research*, **42**, 2544-2557 (2003).

[17] LI, J., ZHAO, Z., KAZAKOV, A. AND DRYER, F. L. An updated comprehensive kinetic model for H_2 combustion. *Fall Technical Meeting of the Eastern States Section of the Combustion Institute, Penn State University, University Park, PA, October 26-29, 2003.*

[18] G. P. Smith, D. M. Golden, M. Frenklach, N. W. Moriarty, B. Eiteneer, M. Goldenberg, C. T. Bowman, R. K. Hanson, S. Song, W. C. Gardiner, V. V. Lissianski, and Z. Qin, *http://www.me.berkeley.edu/gri_mech/*