Scaling and Efficiency of PRISM in Adaptive Simulations of Turbulent Premixed Flames

John B. Bell[a], Nancy J. Brown[a], Marcus S. Day[a],
Michael Frenklach[a,b], Joseph F. Grcar[c], Richard M. Propp[a] and Shaheen R. Tonse[a]

[a] Lawrence Berkeley National Laboratory, Berkeley, CA
[b] Mechanical Engineering Department, UC Berkeley, Berkeley, CA
[c] Sandia National Laboratories, Livermore, CA

Abstract

The dominant computational cost in modeling turbulent combustion phenomena numerically with high fidelity chemical mechanisms is the time required to solve the ordinary differential equations associated with chemical kinetics. One approach to reducing that computational cost is to develop an inexpensive surrogate model that accurately represents evolution of chemical kinetics. One such approach, PRISM, develops a polynomial representation of the chemistry evolution in a local region of chemical composition space. This representation is then stored for later use. As the computation proceeds, the chemistry evolution for other points within the same region are computed by evaluating these polynomials instead of calling an ordinary differential equation solver. If initial data for advancing the chemistry is encountered that is not in any region for which a polynomial is defined, the methodology dynamically samples that region and constructs a new representation for that region. The utility of this approach is determined by the size of the regions over which the representation provides a good approximation to the kinetics and the number of these regions that are necessary to model the subset of composition space that is active during a simulation. In this paper, we assess the PRISM methodology in the context of a turbulent premixed flame in two dimensions. We consider a range of turbulent intensities ranging from weak turbulence that has little effect on the flame to strong turbulence that tears pockets of burning fluid from the main flame. For each case, we explore a range of sizes for the local regions and determine the scaling behavior as a function of region size and turbulent intensity.

Introduction

Reactive computational fluid dynamics (CFD) calculations with large grid sizes and large chemical reaction sets are not practical on today's computers. In order to reduce the severity of a brute-force approach to chemical kinetics problems, various approaches are used, often in combination with one another. Some of the methods used are: simple strategies for reduction in the number of individual chemical reactions, and associated with that, reduction in number of chemical species (N_s) [1, 2], steady-state and partial equilibrium approximations [3, 4], principal component analysis [5], and parametrization of intrinsic low-dimensional manifolds (ILDM) [6]. In the latter approach, introduced by Maas and Pope [6], chemical compositions are viewed as points in a multi-dimensional space and the kinetics problem is addressed from a dynamical systems perspective. In this context, they observe that rapid equilibration of reactions having fast time scales causing arbitrary mixtures to quickly evolve onto a low-dimensional manifold (attractor).

Another class of methods to reduce the costs associated with complex chemistry is to develop surrogates or models for the evolution of the chemical kinetics equations. In these types of approaches the evolution of the chemical kinetics for a time interval is viewed as a mapping from one chemical composition (and temperature) to another. The methods then try to construct an approximation to that mapping that is inexpensive to evaluate. Approaches within this class include Solution Mapping [7,8], piecewise implementation of the above (PRISM) [9], fifth- to eighth-order polynomial parameterizations to the solution of the ordinary differential equations (ODEs) [10], and *in situ* adaptive tabulation (ISAT) [11, 12].

The construction of a suitable surrogate model typically requires the solution of the kinetics equations multiple times. But once a surrogate model is developed, its utilization does not require additional ODE integration and is decoupled from the specific chemical knowledge of the problem. The model accuracy can be controlled by the surrogate construction. Because of the strongly nonlinear character of kinetics equations, surrogate models are typically only valid in some localized neighborhood. Thus, the key questions in using this type of approach to reduce the cost of reacting flow computations are: 1) How many local surrogates are required to represent the chemical compositions and temperatures needed in a given problem? and 2) Is there sufficient reuse of the surrogates to offset the cost of their construction?

The answer to these questions is related to the geometry of the solution in chemical phase space. More precisely, we can describe the thermo-chemical state of the fluid with N_s species and reaction temperature at any instant of time as a point in an N_s+1 dimensional space, denoted hereafter by $\mathbf{C}$. For large chemical mechanisms with many species, constructing local surrogates for the entire space is impractical. The key issue is what portion of $\mathbf{C}$ is accessed over the course of a simulation. We will refer to this subset of $\mathbf{C}$ as the active part

of composition space. As noted by Pope [11], from an analytical perspective we can easily develop *a priori* bounds on the dimension of the active part of $\mathbf{C}$ that must be represented by surrogates. If we assume that the solution is smooth, then the projection of the solution as a function of space and time into $\mathbf{C}$ defines a manifold whose dimension is bounded by the physical dimensions (including time) of the problem. Thus, for the two-dimensional time-dependent flows considered here, in theory we need only cover a three-dimensional subset of composition space.

Although the dimension bounds from the preceding argument offer some hope that surrogate representations of chemical kinetics will be viable for large mechanisms, in practice the issues are more complex. As noted above, when we build a local surrogate it is valid over a limited region. For example, in PRISM the surrogate is defined on regions that are hypercubes. The size of these hypercubes is determined by accuracy requirements. This size defines a specified granularity for covering the active portion of composition space which can effect the observed, effective dimensionality. Several extremes are possible. If the active surface is sufficiently convoluted it may appear to have a higher dimension than the theoretical bounds. On the other hand, the active set may be sufficiently confined and localized that it can be covered by a small number of local surrogates. The issue, then, is whether PRISM is operating in a granularity domain at which the active part of $\mathbf{C}$ has low dimensionality.

In this work, we numerically investigate the dimensionality of the active surface for the case of a premixed hydrogen flame propagating through an unburnt turbulent mixture in two dimensions. For this study we will use the PRISM methodology described below as a surrogate for integration of the chemical kinetic ODE's. We examine how the number of hypercubes scales with the size of the cubes to characterize the geometry of the active subset at the length scales required for accuracy of the PRISM surrogate. The present analysis also addresses the reuse of generated hypercubes, which provides a measure of the potential utility of the PRISM approach.

Description of Methodology

The computations use the low Mach number equations for reacting flow. In the low Mach number setting, the equations describing momentum transport and conservation of species and enthalpy are given by

$$\rho \frac{DU}{Dt} = -\nabla \pi + \nabla \cdot \tau, \tag{1}$$

$$\frac{\partial \rho Y_m}{\partial t} + \nabla \cdot U \rho Y_m = \nabla \cdot \rho \mathcal{D}_m \nabla Y_m - \dot{\omega}_m, \tag{2}$$

$$\frac{\partial \rho h}{\partial t} + \nabla \cdot U \rho h = \nabla \cdot \lambda \nabla T + \sum_m \nabla \cdot \rho h_m(T) \mathcal{D}_m \nabla Y_m \tag{3}$$

where ρ is the density, U is the velocity, Y_m is the mass fraction of species m, h is the enthalpy of the gas mixture, and T is the temperature. Here, $\dot{\omega}_m$, the net destruction rate for ρY_m due to chemical reactions, is specified via a collection of fundamental reactions using a CHEMKIN-III [13] interface. The stress tensor is given by

$$\tau = \mu \left(\frac{\partial U_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i} - \frac{2}{3}\delta_{ij} \nabla \cdot U \right) \tag{4}$$

where $\mu(Y_m, T)$ is the viscosity, $\mathcal{D}_m$ are the species mixture-averaged diffusion coefficients [14], λ is the thermal conductivity and $h_m(T)$ is the enthalpy of species m. These evolution equations are supplemented by an equation of state:

$$p_0 = \rho R_{mix} T = \rho \mathcal{R} T \sum_m \frac{Y_m}{W_m} \tag{5}$$

where W_m is the molecular weight of species m, and by a relationship between enthalpy, species and temperature:

$$h = \sum_m Y_m h_m(T). \tag{6}$$

The numerical methodology is a variation of work discussed in detail by Day and Bell [15]. Following Day and Bell, we employ a symmetric operator-splitting scheme to decouple numerically the chemistry and CFD processes. The CFD component is treated using a fractional step scheme based on an approximate projection algorithm presented by Almgren *et al.* [16] and extended to low Mach number flows by Pember *et al.* [17]. For the chemistry, Day and Bell used VODE [18] to evolve the chemistry pointwise in time; in this work, we use PRISM to evolve the chemistry.

As in Day and Bell, we focus computational effort where it is need by using a block-structured adaptive mesh refinement approach. In block-structured refinement, the problem is solved on a hierarchy of nested grids with sucessively finer spacing in both time and space. In this approach, fine grids are recursively embedded in coarse grids until the solution is adequately resolved. An error estimation procedure based on user-specified criteria determines where additional refinement is needed; grids are dynamically added or deleted as necessary.

Piecewise Reusable Implementation of the Solution Mapping (PRISM) [9] makes numerical chemistry calculations more efficient by constructing a surrogate for the kinetics evolution and storing it in a data structure, to be retrieved later if the time evolution of a similar composition is required. A solution–mapping technique is used, in which the time–integration of the chemical rate equations is parameterized by a set of algebraic polynomial response surfaces in $\mathbf{C}$. The solution–mapping is done piecewise after dividing the space into cubes of dimension $\mathrm{N_s}$+1 (hypercubes) with a distinct parametrization for each hypercube. We partition chemical composition space into non-overlapping hypercubes, each adjacent

to another. Calculations are not performed, nor is storage allocated, for a hypercube until the reaction trajectory enters it for the first time. The location of the initial point, $\mathbf{r}(t_0) \equiv [\mathrm{C}_1,\mathrm{C}_2,\ldots,\mathrm{C}_{N_s},T]$, determines the first hypercube for which polynomials are to be constructed. All calculations and hypercube positions use the natural logarithm of concentration because the concentrations behave in a scale-independent manner over many orders of magnitude and the underlying physical equations typically conform better to a quadratic model under this transformation.

To parameterize the response of the time-integration, the ODE solver must be called repeatedly at selected points within the hypercube. The points are largely located on corners of the hypercube, with an additional point at the hypercube center, and one more outside each face, known as a star point. Their placement is determined by use of factorial designs [19]. Each point corresponds to a set of input concentrations, a temperature, and a time-step length. These input concentrations and temperature are propagated by the ODE solver for the length of time specified, returning a set of response concentrations and temperature. A quadratic regression is then fitted to each response, resulting in $\mathrm{N_s}+1$ polynomials for the hypercube. By evaluating these $\mathrm{N_s}+1$ polynomials we can approximate the time evolution of *any* input point located in the hypercube. Evolving a concentration through time in this manner entails successive evaluations, with the response from one time-step used as input for the next. Perhaps for a few time-steps the polynomials of this hypercube will successively map into $\mathbf{r}(t_1)$, $\mathbf{r}(t_2)$, $\mathbf{r}(t_3)$,..., all the while remaining in the same hypercube, but at some time $t = t_n$, the solution $\mathbf{r}(t_{n+1})$ will fall outside the hypercube. When this happens we must determine within which hypercube the point $\mathbf{r}(t_{n+1})$ lies and construct polynomials for it, and so on. The polynomial coefficients as well as the hypercube position and extent are stored in a data structure.

Description of Case Study

The above methodology is applied to the interaction of a two-dimensional premixed hydrogen flame with a patch of turbulence superimposed in the cold region below the flame front (see Fig. 1). Inflow occurs at the bottom of the domain, while outflow occurs at the top of the domain. The inflow velocity is set to the laminar flame speed so that in the absence of turbulence, the flame is stationary. The other two sides of the domain are periodic. We initialize the turbulent velocity field and the premixed flame at time $t = 0$, and then let the system evolve. The upward velocity due to the premixed flame causes the turbulent region to perturb the flame front.

The chemistry mechanism contains 9 species and 27 reactions and is derived by removing carbon from the GRI-Mech 1.2 reaction model [20]. The initial velocities and concentrations are computed using the PREMIX code from the CHEMKIN collection [21] and interpolated

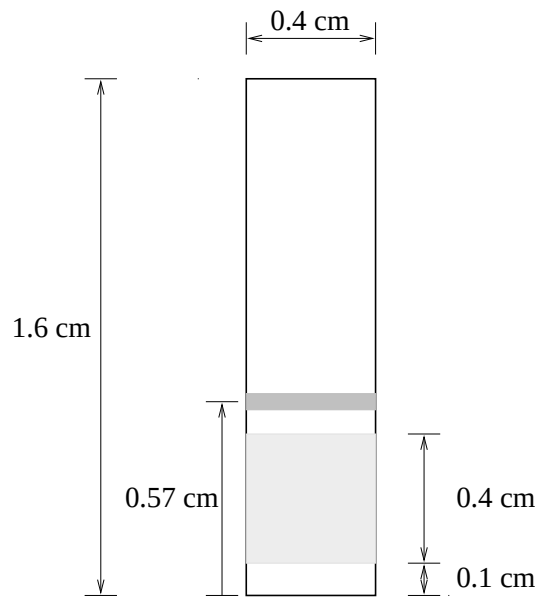

Fig. 1: Schematic of the sample problem. The flame front is initially 0.57 cm from the bottom of the domain. There is a 0.4 cm × 0.4 cm patch of turbulence below the flame.

onto the hierarchy of grids. The inflow conditions are $X_{\mathrm{H}_2} = 0.1909$, $X_{\mathrm{O}_2} = 0.0910$, $X_{\mathrm{N}_2} = 0.7181$, $T = 298K$, and velocity of 79.675 cm/s.

Below the flame front, we superimpose a turbulent velocity field onto the velocity that already exists due to the premixed flame. The isotropic turbulence is prescribed by the energy spectrum:

$$E(k) = \frac{(\frac{k}{k_i})^4}{[1 + (\frac{k}{k_i})^2]^{\frac{17}{6}}} exp[-\frac{9}{4}(\frac{k}{k_d})^{\frac{4}{3}}] \tag{7}$$

where k is the wavenumber, $k_d = 1/(2\Delta x)$, k_i is the peak frequency, and Δx is the mesh spacing. In this work, we use $\Delta x = 0.0015625$ cm and a peak frequency of 1500 cm^{-1}.

We compute the solution with three different initial turbulent intensities: 1, 10, and 100 cm/s. These intensities represent a broad range of wrinkled flame phenomena. (We note that these intensities decay in time.) For each of these intensities, we use three different cube sizes — full-size, half-size, and quarter-size. A full-size hypercube is defined to have side lengths consisting of 0.1 of $\log_{10}$(concentration) for the 3 major species H_2, O_2 and H_2O, 0.25 for the other species, and 20 K for temperature. In half-size hypercubes, the concentration widths are half of the size of the full-size cube, while the temperature length is unchanged. The quarter-size hypercubes are similarly defined.

The CFD base grid for these simulations was 64 × 256 with uniform mesh spacing. In addition, we used two levels of refinement of a factor of two each; thus the finest grid has

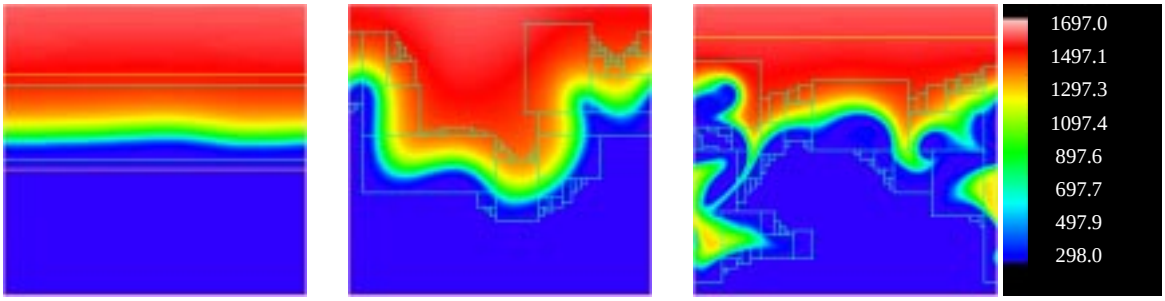

Fig. 2: Temperature near the flame front for the intensity = 1, 10, 100 cm/s cases.

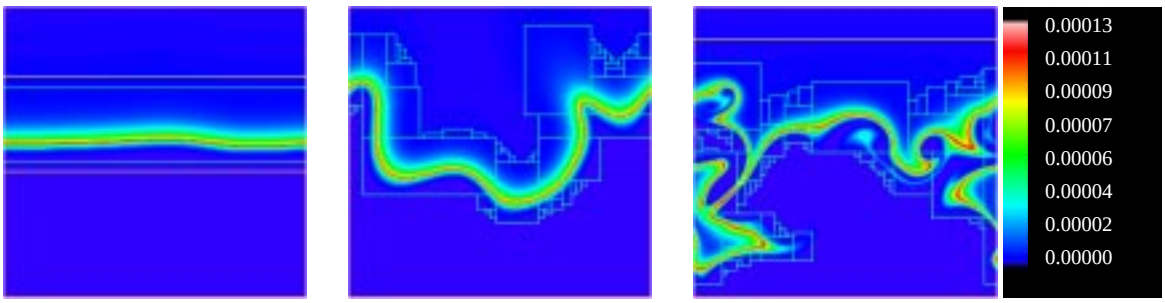

Fig. 3: $X_{\mathrm{HO_2}}$ near the flame front for the intensity = 1, 10, and 100 cm/s cases.

an equivalent resolution of 256 $\times$ 1024. Level 1 and level 2 cells were created where the temperature difference between neighboring cells was greater than 50 K. Level 1 cells were created where the magnitude of vorticity was greater than 500 s^{-1}.

Data, Analysis and Commentary

In Fig 2 we show the temperature near the flame front for each level of turbulent intensity. The difference in morphology of the flame as a function of turbulent intensity is apparent. For the lower intensity, the turbulence barely perturbs flame. At medium intensity the flame is substantially wrinkled. In the high intensity case, portions of the flame have been broken off and there are ephemeral fingers of hot gas extending into the cold region. Figure 3 shows the mole fraction of $X_{\mathrm{HO_2}}$, which turned out to provide a distinctive marker for the reaction zone. For the two lower intensities the flame is perturbed but $X_{\mathrm{HO_2}}$ is relatively uniform across the flame. For the higher intensity, there is considerable variability in $X_{\mathrm{HO_2}}$ indicating substantial perturbations of the chemistry resulting from the flame-turbulence interaction.

The changes in the complexity of the flow are reflected in the number of hypercubes required by PRISM to represent the flow. The number of cubes required by the code as a function of turbulent intensity and cube size is summarized in Table 1. We see a dramatic increase in the number of cubes between the low and medium intensity case and a modest

Table 1: Number of hypercubes for each simulation.

Initial Turbulent Intensity (cm/s)	Hypercube Size	number of hypercubes
1	full	572
1	half	1523
1	quarter	5223
10	full	4570
10	half	20935
10	quarter	125562
100	full	7241
100	half	31294
100	quarter	151633

Table 2: Effective dimensionality of region covered by hypercubes.

Intensity (cm/s)	full $\rightarrow$ half	half $\rightarrow$ quarter
1	1.412	1.778
10	2.196	2.584
100	2.112	2.277

additional increase for the high intensity case.

We can characterize the geometry of the active portion of $\mathbf{C}$, the region covered by the hypercubes, by computing its effective dimension. This dimension was computed by changing the composition-space granularity, i.e., by halving the lengths of the hypercube edges. For example, for a one-dimensional surface, halving the edges of the hypercubes will only double the number required to cover the surface. In general, for dimensionality d the number of hypercubes used should rise by 2^d. Effective dimensions for the full-size to half-size transition and the half-size to quarter-size transition for each intensity are summarized in Table 2.

Over the range of scales represented by the hypercubes used here the weak intensity case is only active in a region of dimension less than two in composition space. In this case the dimensionality suggests a "thickened" one-dimensional surface, e.g., a narrow band. This then suggests that the chemical composition stays in the neighborhood of a one-dimensional surface. In the more intense cases the dimension is between 2 and 3 suggesting that the chemical composition stays in a region near a two-dimensional sheet. In Fig 4 we present three-dimensional projections showing the location of active hypercubes for the three different intensities. The structure indicated in the figures matches the expectations from our analysis of the dimensionality.

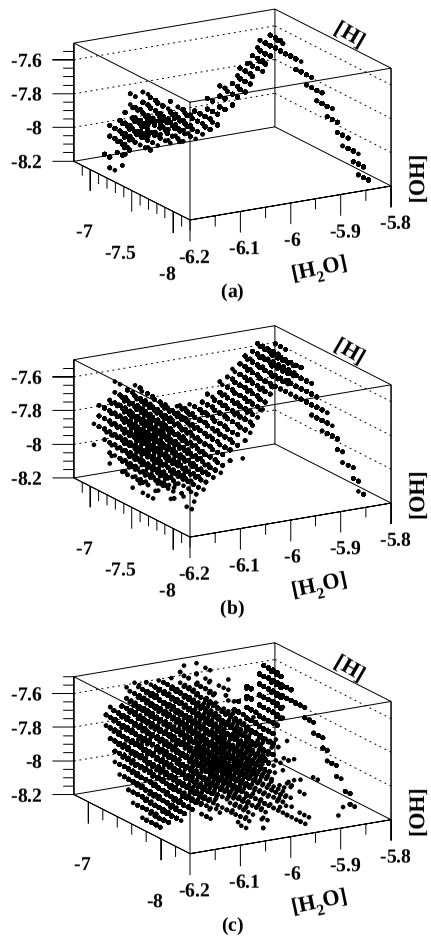

Fig. 4: The active portion of chemical composition space for: (a) turbulent intensity = 1 cm/s, (b) turbulent intensity = 10 cm/s, and (c) turbulent intensity = 100 cm/s. The plots are $\log_{10}$ (concentration) of OH, H, and H_2O.

The utility of PRISM for these types of simulations is related to the accuracy associated with the surrogate model within each hypercube. We apply a conservative criterion for local accuracy that was chosen to preserve a high level of global accuracy. Specifically, we require that each species within a hypercube has a relative error in molar concentration of less than 5×10^{-3} and an absolute error of less than 10^{-4} times the sum of all the molar concentrations. This level of local error tolerance maintained global concentration errors of 0.1–0.5% in the present computations. Using these criteria we find that for turbulence intensities of 1 and 10 approximately 14-15% of the full-size hypercubes are not sufficiently accurate. For the high intensity case the number of hypercubes that fail to meet these criteria drops to 0.3%. (The improved behavior at high intensity is because the time step used by the fluids code is smaller, placing less stringent demands on the surrogate.)

In the present implementation, when a hypercube fit is rejected by the error criteria, the ODE solver is called instead of using the surrogate polynomial. Therefore, an inability to meet the accuracy requirements has implications only for the economy of the method. The cost of construct polynomials has been wasted, but the accuracy of the simulations was not sacrificed. A small fraction of "inaccurate" hypercubes is actually informative, as it indicates that our hypercubes are not too small. For the half-size cubes, all the surrogates met the present accuracy requirements for all the intensities. Thus, the half-size hypercubes provide a lower limit for the "accurate" cube size. Considering that the chemistry solver is called twice for the advancement of each CFD computational cell, we obtain an average reuse of half-size hypercubes ranging from approximately 2000 in the strong turbulence case to about 2700 in the low intensity case. Thus, even for the stringent error requirement, PRISM effectively reduces the computational time by more than an order of magnitude.

Conclusions

Computational techniques that construct surrogates for the evolution of chemical kinetics equations offer the potential for dramatically reducing computational costs without sacrificing chemical fidelity. The utility of a given approach depends on the range of chemical compositions that are encountered in a simulation and the number of local models for the kinetics that are needed to cover that range with a desired accuracy. We have examined this issue using the PRISM methodology applied to modeling of two-dimensional premixed hydrogen flames. The results indicate that the range of chemical compositions is restricted to a low-dimensional subset of chemical compositions and that the region can be effectively represented by a modest number of hypercubes. Conservative choice of hypercube size allows high-fidelity simulation with more than an order of magnitude reduction in computational cost.

In future work we plan to extend the study presented here to non-premixed flames and

to three dimensions. Another question we plan to address is the universality of the regions of composition that are active in a given computation. In particular, we intend to explore under what circumstances are hypercubes generated for one problem applicable to another problem.

Acknowledgements

The work of JB, MD, and RP was supported by the Laboratory Directed Research and Development Program of Lawrence Berkeley National Laboratory under the U.S. Department of Energy, and that of NB, ST, and MF by the Director, Office of Energy Research, Office of Basic Energy Sciences, Chemical Sciences Division of the U.S. Department of Energy, all under contract No. DE-AC03-76SF00098. JG's work was supported by Sandia National Laboratories, a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract DE-AC04-94AL85000.

References

[1] Frenklach, M., *Complex Chemical Reaction Systems, Mathematical Modelling and Simulation* (J. Warnatz and W. Jäger, eds.), Volume 47 of *Springer Series in Chemical Physics*, Springer-Verlag, Berlin, 1987, pp. 2–16.

[2] Frenklach, M., *Numerical Approaches to Combustion Modeling* (E. S. Oran and J. P. Boris, eds.), American Institute of Aeronautics and Astronautics, Washington, D.C., 1991, pp. 129–154.

[3] Peters, N. and Williams, F. A., *Complex Chemical Reaction Systems, Mathematical Modelling and Simulation* (J. Warnatz and W. Jäger, eds.), Volume 47 of *Springer Series in Chemical Physics*, Springer-Verlag, Berlin, 1987, pp. 310–317.

[4] Ramshaw, J. D., *Phys. Fluid.*, 23:675 (1980).

[5] Brown, N. J., Li, G., and Koszykowski, M. L., *Int. J. Chem. Kinet.*, 29:393–414 (1997).

[6] Maas, U. and Pope, S. B., *Combust. Flame*, 88:239–264 (1992).

[7] Frenklach, M., *Combustion Chemistry* , chapter 7, Springer-Verlag, 1984.

[8] A. R. Marsden Jr., Frenklach, M., and Reible, D. D., *J. Air Pollut. Control Assoc.*, 37:370–376 (1987).

[9] Tonse, S. R., Moriarty, N. M., Brown, N. J., and Frenklach, M., *Israel J. Chem.*, 39:97–106 (1999).

[10] Turanyi, T., *Comp. Chem.*, 18:45–54 (1994).

[11] Pope, S. B., *Combust. Theory Modelling*, 1:41–63 (1997).

[12] Yang, B. and Pope, S. B., *Combust. Flame*, 112:85–112 (1998).

[13] Kee, R. J., Ruply, F. M., Meeks, E., and Miller, J. A., "CHEMKIN-III: A FORTRAN chemical kinetics package for the analysis of gas-phase chemical and plasma kinetics," Sandia National Laboratories Report No. SAND96-8216.

[14] Kee, R. J., Dixon-Lewis, G., Warnatz, J., Coltrin, M. E., and Miler, J. A., "A FORTRAN computer code package for the evaluation of gas-phase multicomponent transport properties," Sandia National Laboratories Report No. SAND86-8246.

[15] Day, M. S. and Bell, J. B., *Combust. Theory Modelling*, (1999), submitted for publication.

[16] Almgren, A. S., Bell, J. B., Colella, P., Howell, L. H., and Welcome, M., *J. Comput. Phys.*, 142:1–46 (1998).

[17] Pember, R. B., Howell, L. H., Bell, J. B., Colella, P., Crutchfield, W. Y., Fiveland, W. A., and Jessee, J. P., *Comb. Sci. Tech.*, 140:123–168 (1998).

[18] Brown, P. N., Byrne, G. D., and Hindmarsh, A. C., *SIAM J. Sci. Stat. Comput.*, 10:1038–1051 (1989).

[19] Box, G. E. P. and Draper, N. R., *Empirical Model Building and Response Surfaces*, John Wiley and Sons, New York, 1987.

[20] Frenklach, M., Wang, H., Goldenberg, M., Smith, G. P., Golden, D. M., Bowman, C. T., Hanson, R. K., Gardiner, W. C., and Lissianski, V., "GRI-Mech—an optimized detailed chemical reaction mechanism for methane combustion," Technical Report GRI-95/0058, http://www.me.berkeley.edu/gri_mech/.

[21] Kee, R. J., Grcar, J. F., Smooke, M. D., and Miller, J. A., "A FORTRAN program for modelling steady laminar one-dimensional premixed flames," Sandia National Laboratories Report SAND85-8240.