5th US Combustion Meeting
Organized by the Western States Section of the Combustion Institute
and Hosted by the University of California at San Diego
March 25-28, 2007.

An adaptive approach for coupling detailed chemical kinetics and multidimensional CFD

L. Liang, J. G. Stevens, and J. T. Farrell

Corporate Strategic Research Laboratories, ExxonMobil Research and Engineering Co., Annandale, New Jersey 08801, USA

P. T. Huynh, I. P. Androulakis, and M. Ierapetritou

Department of Biomedical Engineering, Rutgers University, New Brunswick, New Jersey 08854, USA

An adaptive approach for coupling detailed kinetics with computational fluid dynamics (CFD) has been formulated to permit reactive flow simulations with improved computational efficiency. The approach makes use of a set of reduced mechanisms, each of which has been developed to reproduce the mixture kinetics under well-characterized thermo-chemical conditions. In the present implementation, the KIVA-3V code is used as the CFD framework, and CHEMKIN is used to formulate the detailed kinetics that arise in the source terms of the energy and species transport equations. An efficient two-stage mechanism reduction methodology was developed that combines elemental flux analysis and the Mixed Integer Nonlinear Programming method. A high-dimensional, nearest-neighbor search (NNS) method was implemented in the adaptive chemistry code to determine the appropriate reduced mechanism to use under a given set of thermo-chemical conditions. As a preliminary test of the adaptive chemistry concept, three reduced versions of a 385-species, 1895-reaction pentane oxidation mechanism were developed to reproduce the mixture kinetics in three contiguous temperature regimes. By using temperature alone as the query criterion, the adaptive approach shows great potential to improve computational efficiency while retaining acceptable accuracy. To test the more generally applicable NNS method, a total of 18 distinct reduced versions of the GRI 3.0 mechanism were generated that capture the essential methane oxidation chemistry over a wide range of equivalence ratio and temperature. The adaptive approach was evaluated by modeling compression ignition of a stoichiometric methane/air mixture in an HCCI engine. The accuracy and computational efficiency benefits of this approach have been assessed through comparisons to baseline cases involving the use of the full mechanism. Additional computational savings are realized by redimensioning the species array prior to integration to include only active species.

1. Introduction

The desirability of incorporating detailed chemistry in computational fluid dynamics (CFD) calculations has been widely recognized in the combustion science and engineering community in recent years. A commonly-employed approach to this coupling is to utilize detailed chemical kinetics to formulate the source terms in the species and internal energy transport equations of the CFD calculation. The use of detailed chemical kinetics gives rise to a large system of nonlinearly coupled stiff ordinary differential equations (ODE) for each CFD computational cell.

The numerical solution of the large number of such systems within the CFD framework results in dramatically increased requirement for CPU times. Previous studies have demonstrated that even with skeletal kinetic mechanisms, e.g., with the number of species less than 50, over 90% of the CPU time was consumed in solving the ODE systems in reactive flow simulations [1, 2]. This obstacle has led many researchers to explore alternatives that can drastically reduce CPU times and still maintain high solution accuracy.

Several methodologies have been proposed that generally fall in three categories: mechanism reduction, storage/retrieval-based methods, and adaptive chemistry strategies. Mechanism reduction approaches include the quasi-steady-state-assumption (QSSA) for some intermediate species [3], the partial equilibrium approximations to particular reactions, the computational singular perturbation (CSP) method [4], time integrated elemental flux analysis [5], the intrinsic low-dimensional manifolds (ILDM) method [6], and the direct-relation-graph (DRG) method [7]. Storage/retrieval approaches usually have very intensive computer memory requirements. Two typical examples are the *in situ* adaptive tabulation (ISAT) method [8], and the piecewise reusable implementation of solution mapping (PRISM) procedure [1]. The adaptive chemistry strategy has been explored in several previous publications [9, 10]. The basic idea of adaptive chemistry is to apply a set of locally accurate, reduced mechanisms, instead of the full mechanism or a single skeletal mechanism, based on the instantaneous, local thermo-chemical conditions in reactive flow simulations. Compared to single skeletal mechanism applications in which the reduced model must capture the relevant chemistry over vastly different conditions, it is relatively easier for the submechanisms in an adaptive chemistry library to entail a smaller number of species and yet to achieve superior accuracy, because a given submechanism is used only within a narrow ranges of thermo-chemical conditions. However, accurate and efficient mechanism reduction and the query technique used to determine the choice of the appropriate submechanism are still two main challenges to successful application of the adaptive chemistry to multidimensional CFD calculations.

Here we present a general discussion of our current research effort on adaptive chemistry, including a brief review of the general roadmap, the two-stage mechanism reduction methodology, a high dimensional nearest neighbor search scheme for adaptive query, and issues related to the implementation of the adaptive scheme in a CFD code. In the present work, CHEMKIN [11] was used to solve the chemical kinetic systems, and the KIVA-3V code [12] was used as the CFD framework. Two 2D computational meshes with moving boundaries were created to simulate homogeneous charge compression ignition (HCCI) combustion. A pentane oxidation mechanism with 385 species and 1895 reactions [13] and the GRI-Mech 3.0 for methane combustion [14] were used for testing the integrated models.

2. Review of General Roadmap

2.1. Generation of Adaptive Kinetic Mechanism Library and Query Scheme

A detailed description of the reduction technique and the feasibility analysis scheme used in this study has been presented in previous publications [10]. Following that approach, a pair-wise mixed stirred reactor (PMSR) model is used as a reactor framework for generation of the library of reduced models. First, the detailed reaction mechanism is employed in the PMSR model to provide an estimate of the range of conditions accessed by a reactive flow model. The conditions in the state space of thermo-chemical parameters (typically, species mass fractions, temperature,

and pressure) constituting the accessed region are then clustered in order to capture patterns in the data set and identify representative points in the pattern. Reduced mechanisms are then generated at each of the cluster centers. The next step is the determination of the range of validity of the resulting reduced model, i.e., the region in state space within which it achieves the specified accuracy. In order to determine its extent, the points in the cluster used for reduction are first checked to see whether they are feasible, i.e., whether they are in this region. These feasible points set the limits of the region of validity for the given reduced mechanism. They are removed and the updated data set is re-clustered to obtain another nominal point for reduction. This procedure is repeated until enough reduced mechanisms have been generated to cover the desired portion of the accessed region. In previous work [10], a so called 'point-in-polygon' searching procedure was employed for adaptive determination of the valid reduced model. This procedure was demonstrated by simulating methane combustion in PMSR, where fuel-oxygen-temperature space was efficiently used to represent the space for the combustion process. However, that approach is difficult to extend to higher dimensions, and consequently an alternative approach has been developed (described later).

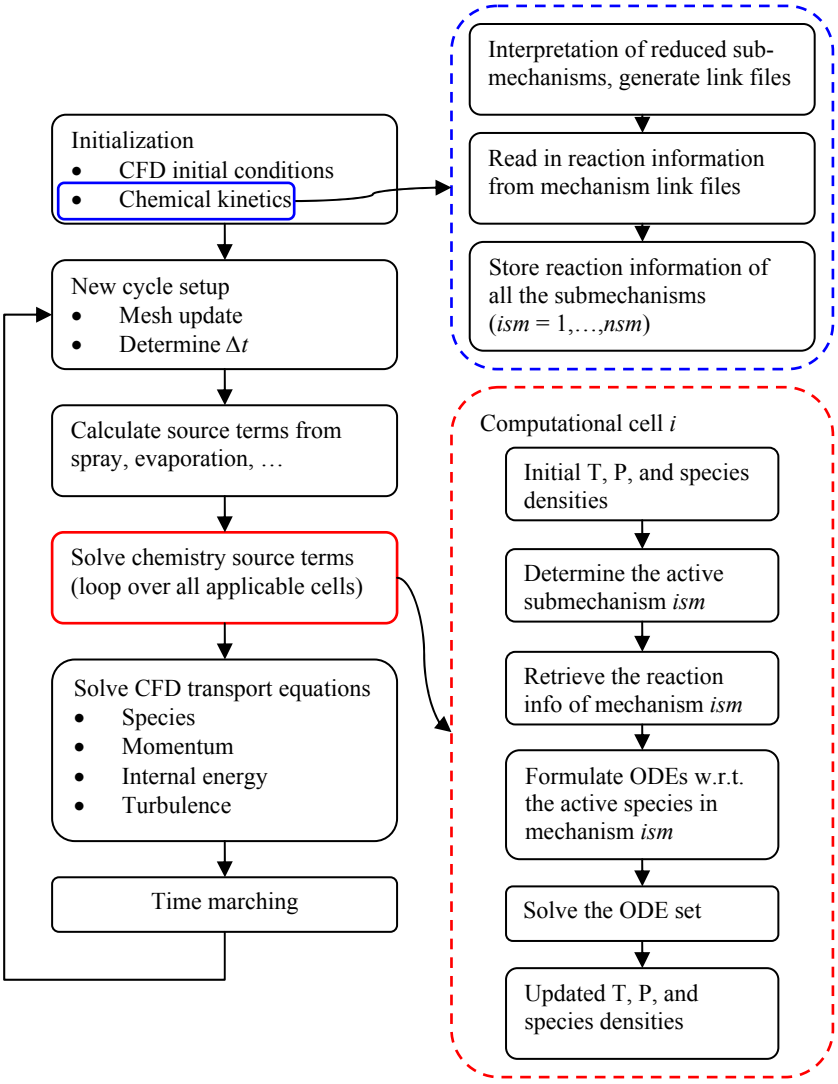

Figure 1: Flowchart for coupling adaptive chemistry with multidimensional CFD.

2.2. Implementation of the Adaptive Scheme into CFD code

Figure 1 is a flow chart that describes the coupling of adaptive chemistry with the KIVA-3V CFD code. All the reduced submechanisms were interpreted to produce link files using CHEMKIN. These files were read in and stored in the CFD code during initialization. Subsequently, in the time marching calculation, the choice of the submechanism to be used is based on the results of the specific local query.

Because the reduced subsets typically have different numbers of species and/or reactions, there are two options for choosing the numbers of species of the reduced sets when mechanism interpretation is performed. One choice is to keep the union of the active species of all the reduced sets, so that the number of ODEs formulated is based on the number of species in the union, no matter which subset is being invoked. This choice is straightforward. However, in cases for which the number of species in the union is much larger than the number of species in an individual subset, including the inert species will result in a larger system of ODEs thereby unnecessarily increasing CPU time, as will be shown below. The other choice is to retain only the active species for each individual subset, so that number of ODEs will be based on the number of active species in that submechanism. However, in this case, only the active species concentrations are taken into account in reaction rate calculations. As a result, third-body effects from the inert species are also neglected. This omission can result in dramatic errors under certain conditions. In the present implementation, the species array of the reduced mechanisms is redimensionalized so that ODEs are formulated with respect to only the active species of an invoked subset, while both the inert and active species are considered in reaction rate evaluations.

3. Two-Stage Mechanism Reduction

The adaptive chemistry methodology reported previously [10] is based on the utilization of mathematical programming techniques, in which the objective is to minimize the dimension of the system while retaining sufficient accuracy in the prediction of specific species profiles. This scheme does give accurate and flexible reduced kinetic mechanisms, but is itself an expensive operation, because it requires multiple integrations of the stiff, nonlinear ODEs. Also, the number of binary variables handled by the Mixed Integer Nonlinear Programming (MINLP) formulation can be restrictive. Thus direct application of this procedure for very large mechanisms may prove to be inefficient and in some cases infeasible. Thus we developed a two-step procedure in which the first level reduction is performed using the very effective methodology based on flux-analysis [5], followed by the more detailed reduction using the MINLP approach.

3.1. Stage I: Time-Integrated Element Flux Analysis

The concept of element flux analysis, introduced by Revel et al. [15] allows one to develop time-dependent flux diagrams that identify key reaction pathways with minimal effort. The main idea is to study principal reaction pathways along the reaction process coordinate. The atomic fluxes for each atom (C, H, O, N, …) at different reaction times are calculated based on reaction rates and the major sources and sinks for each element are identified. The instantaneous chemical flux of atom A from species j to species k through reaction i is defined as [15]:

$$A_{ijk} = q_i \frac{n_{A,j} n_{A,k}}{N_{A,i}}, \quad (1)$$

where q_i is the net production rate of reaction i (mol-s^{-1}), $n_{A,j}$ is the number of atoms A in species j, $n_{A,k}$ is the number of atoms A in species k, and $N_{A,i}$ is the total number of atoms of element A in reaction i. For any given atom (A), we can then consider the transformations that occur from species to species via all reactions in the network. The total transfer of element mass for any pair of species can then be defined as:

$$\overline{\dot{A}}_{FROM,TO}(t) = \sum_{i=1}^{N_R} A_{i,FROM,TO}(t). \quad (2)$$

The fluxes defined in Eqs. (1) and (2) are functions of time. Thus the analysis must be repeated at preselected times during the reaction process. In the present work, we overcome this limitation by using the concept of a time-integrated flux indicator. The main idea is to derive an indicator for a source-sink combination based on integration of the quantity defined in Eq. (2). The quantity is then normalized as a means of representing the results in a unified way. Therefore, the following quantity is used:

$$\hat{A}_{FROM,TO} = \frac{\int_{t=0}^{\tau} \overline{\dot{A}}_{FROM,TO}(t)dt}{\max \int_{t=0}^{\tau} \overline{\dot{A}}_{FROM,TO}(t)dt}. \quad (3)$$

The integrals in Eq. (3) are estimated numerically. In this way, the entire reaction trajectory is taken into account with the consequence that there is no need to select the location of the snapshots *a priori*. Moreover, the quantity defined in Eq. (3) assigns a unique, overall value to each source-sink pair that is representative of the entire reaction period. The normalization as defined in Eq. (3) ensures that pathways activated at different points in time receive appropriate weighting.

Time-integrated element flux analysis examines the values of the quantities defined in Eq. (3) and allows one to characterize the main conversions that take place during the reaction period. The analysis reveals key pathways in terms of source-sink relationships. However, this analysis does not reveal any information related to specific reaction path by which these transformations take place. The purpose of the time integrated element flux analysis is to establish 'global' insight into the reaction pathways. It facilitates a comprehensive, macroscopic view of the reactant transformation to intermediates and finally to products.

3.2. Stage II: MINLP Optimization Reduction

Once the skeletal mechanism has been derived using the flux analysis, the MINLP reduction procedure can be applied for further reduction. The formulation of the optimization problem is given by [10]:

$$\min \sum_{k=1}^{Ns} \lambda_k \qquad \text{subject to } \chi \le \delta, \quad (4)$$

$$\chi = \left(\sum_{k \in \kappa} \int_{t_i}^{t_f} \left(\frac{y_k^{reduced}(t) - y_k^{full}(t)}{y_k^{full}(t)} \right)^2 dt + \int_{t_i}^{t_f} \left(\frac{T^{reduced}(t) - T^{full}(t)}{T^{full}(t)} \right)^2 dt \right)^{1/2}, \quad (5)$$

where λ_k is a binary variable used to denote the presence (λ_k=1) or absence (λ_k=0) of a particular species. The objective function $\sum_{k=1}^{Ns} \lambda_k$ represents the total number of species in the reduced set, which has to be minimized, subject to a specified accuracy, δ. The integral error measure, χ in Eq. (5), defines the approximation error of the trajectories of the key observable quantities (species mass fractions y_k and temperature T in the present case) for the interval of interest. In the above formulation, the evolution of y_k and T is governed by the detailed chemical kinetic system coupled with the mass and energy balances for the reactor. Optimization of the number of species also indirectly reduces the number of reactions. When a particular species is taken out of the system, all the reactions in which that species participates are also eliminated from the system. The above described formulation corresponds to an MINLP problem, which was solved using genetic algorithm (GA) in Ref. [10].

4. A Preliminary Test

Before discussing rigorous adaptive query methods, we set up a preliminary test to explore the potential benefit of the adaptive chemistry scheme. It was performed by applying the two-stage reduction methodology and a simple empirical query method to HCCI like combustion simulation of stoichiometric pentane/air mixtures. Following the two-stage reduction procedure, the original detailed pentane oxidation mechanism (385 species and 1895 reactions based on a modified version of an n-heptane mechanism [13]) was reduced into three subsets for low temperature (LT), intermediate temperature (referred to as negative temperature coefficient, NTC, in this paper), and high temperature (HT) regimes, respectively. The degree of reduction is summarized in Table 1. A temperature-based query criterion was employed in this test. Specifically, an empirical value T_S is used as the transition temperature from the LT regime to the NTC regime; the transition point from the NTC regime to the HT regime was chosen to be 3 °CA after the onset of the main heat release as determined by monitoring the changing rate of temperature with crank angle, *dT/dCA*. This 3 °CA delay was chosen to prevent certain intermediate species in the NTC mechanism from being 'frozen' too early during the transition.

Because the size of the full mechanism is very large, a simple one-column-cell numerical mesh was used to create an HCCI environment that could utilize the full mechanism with reasonable computational time. As shown in Fig. 2, the mesh contains 16 cells at bottom dead center (BDC) position. The bottom boundary is a moving boundary, resulting in a compression ratio 16:1. Adiabatic boundary conditions were used in the simulation. The test conditions include different initial pressures and 'engine' speeds, as listed in Table 2. Calculations were carried out from -40 ATDC to 20 ATDC.

Table 1: Numbers of species and reactions of the n-pentane mechanisms.

Mechanisms	Full	LT	NTC	HT
Number of Species, *Ns*	385	103	83	40
Number of Reactions, *Nr*	1895	387	372	135

Table 2: Operating conditions of the n-pentane combustion simulations.

Cases	1 (Baseline)	2	3
Initial Temperature, T_0 (K)	600	600	600
Initial Pressure, p_0 (MPa)	0.5	1	0.5
Engine Speed (rev/min)	1500	1500	2500

Before investigating the adaptive approach, we first used each individual reduced set throughout the entire time period with the baseline conditions (Case 1). The predicted 'in-cylinder' average temperature and pressure are compared to those calculated from the full mechanism in Fig. 3. The LT mechanism predicts earlier ignition timing compared to the full mechanism, most likely because the NTC reactions that 'turn off' the ignition-accelerating low temperature reactions (and thereby reduce the heat release and delay ignition) are not included in this reduced mechanism. Conversely the ignition time predicted by the NTC mechanism is too long since the reactions that promote heat release at low temperature are not included. When only the HT mechanism was applied, misfire was predicted. This phenomenon is not surprising because the HT model does not contain proper low temperature pathways that are necessary to accumulate heat and certain intermediates. These results demonstrate that none of the reduced mechanisms is sufficiently representative of the full mechanism throughout the combustion process and that an adaptive approach is necessary.

Figure 2: The one-column-cell 'HCCI' mesh at BDC.

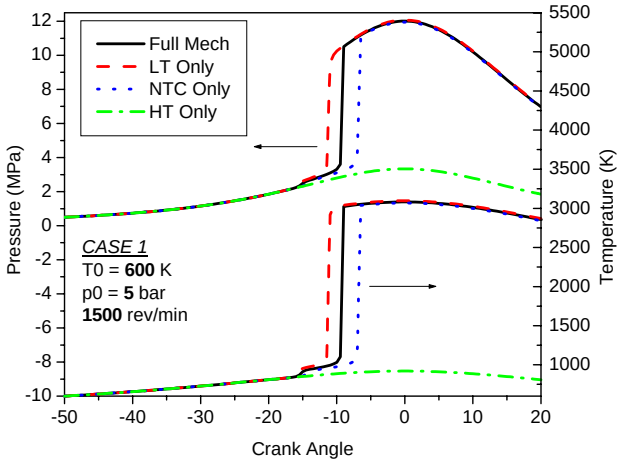

Figure 3: Simulated pressure and temperature using individual mechanism (Full, LT, NTC and HT).

A 'manual' adaptive scheme was first applied to Case 1. By trial and error, we found that T_S = 880K was an optimal value for the transition from LT to NTC, while the simulation results were not sensitive to the threshold value of *dT/dCA* for the second mechanism switching. Figure 4 shows that the behavior of the full mechanism, as reflected by the in-cylinder averaged pressure and temperature traces, can be accurately reproduced by the adaptive combination of the three reduced models. The in-cylinder species mass fractions, shown in Fig. 5, also show encouraging agreement, which further corroborates the effectiveness of the adaptive approach.

For further testing, the same adaptive criteria used in Case 1 were applied to simulations of Cases 2 and 3. Compared to Case 1, Case 2 has a higher initial pressure and Case 3 has a higher engine speed. Predicted pressure and temperature traces are compared between the full mechanism and the adaptive approach in Fig. 6 and Fig. 7, respectively. In both cases, discrepancies in ignition timing of approximately two crank angle degrees were observed. These differences are not unexpected given the simplicity of the adaptive query criteria. Nevertheless, it is encouraging that such straightforward criteria are able to produce such reasonable agreement. Also, the performance of the adaptive scheme could be further improved by fine tuning T_S for each case separately.

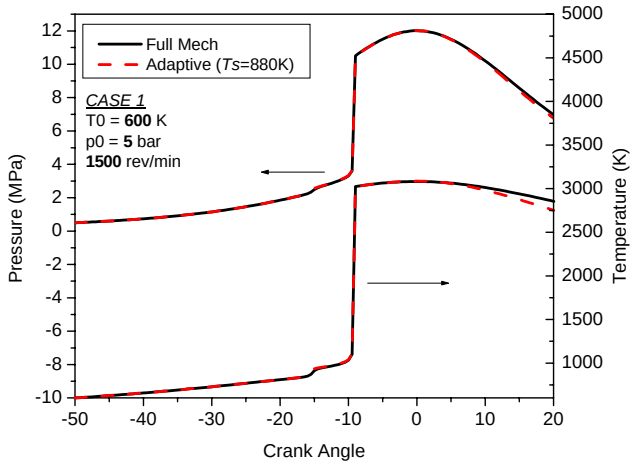

Figure 4: Comparison of pressure and temperature predicted by the full mechanism and by using adaptive scheme.

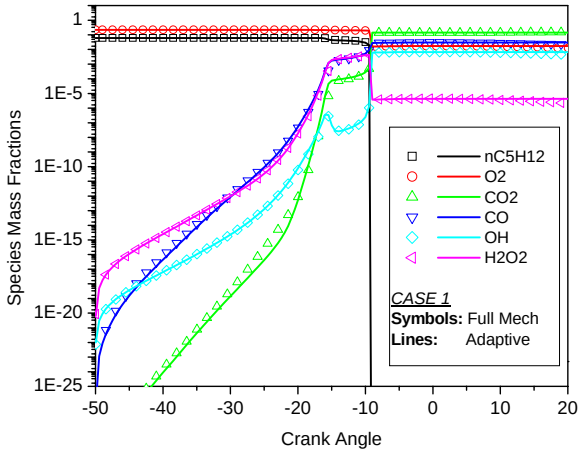

Figure 5: Comparison of in-cylinder species mass fractions predicted by the full mechanism and by using adaptive scheme.

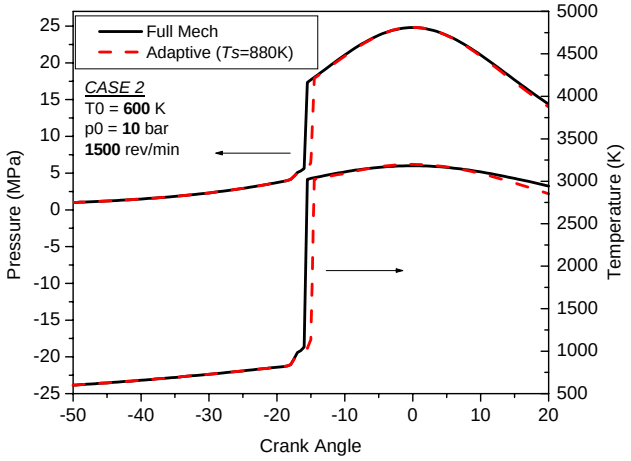

Figure 6: Comparison of pressure and temperature predicted by the full mechanism and by using adaptive scheme for Case 2.

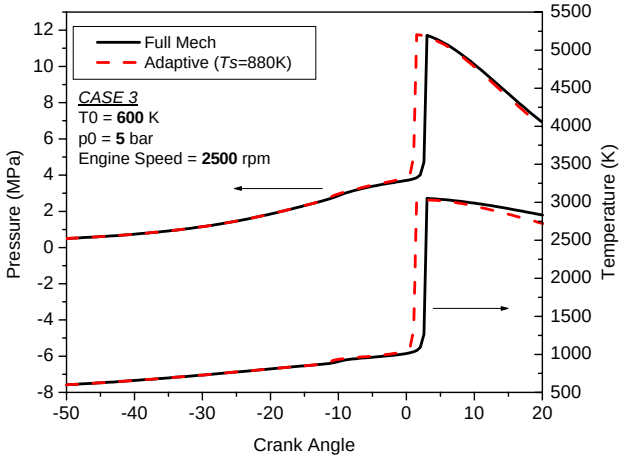

Figure 7: Comparison of pressure and temperature predicted by the full mechanism and by using adaptive scheme for Case 3.

Figure 8 shows the CPU time requirements for simulations using the full mechanism and those using the 'manual' adaptive scheme for the three cases listed in Table 2. It is readily seen that the adaptive scheme is roughly 30 times more efficient than the nonadaptive approach that uses the full mechanism.

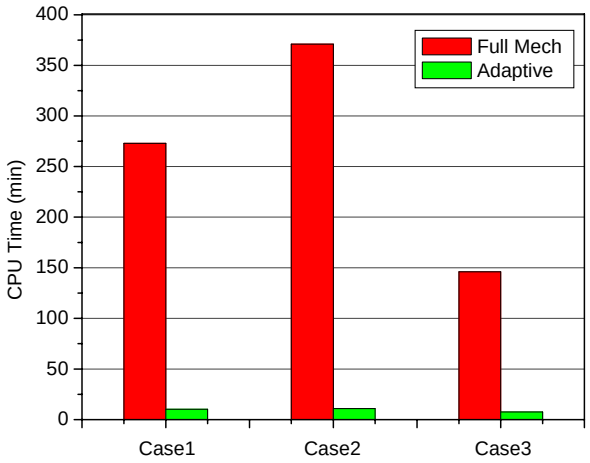

Figure 8: CPU time comparison of full mechanism and adaptive scheme for pentane system.

These results illustrate that reasonable accuracy combined with considerable CPU time savings can be obtained by using a query scheme based on temperature. However, it is necessary to adjust a small number of parameters empirically to obtain the desired agreement. For more complicated scenarios, e.g., when mixture stratification effects are considered, such an empirical approach may be onerous or, at the very least, unappealing. Therefore, it is desirable to pursue a more fundamental approach. A nearest neighbor search method is discussed in the next section as one such alternative.

5. A High Dimensional Nearest Neighbor Search Scheme

During PMSR-based mechanism reduction process, all the state vectors, that have Ns+1 dimensions (species mass fractions plus temperature), were associated with the indices that enumerate the reduced mechanisms. These state vectors are referred to as *nominal points* in this paper. The correspondence between the nominal points and the submechanism indices forms the search database. Given a query point (P_Q), the task is to efficiently locate the point(s) in the database that are closest to P_Q in order to assign the corresponding index. This problem has some interesting complexities associated with it because the database is not a structured one. The raw data points are the ones defining the trajectories of the evolution of the original training system. Therefore, it is very difficult, if at all possible, to identify a structured representation of the data that would enable a very efficient search. Thus we need to identify search and retrieve algorithms that can efficiently operate on highly unstructured data sets. A number of approaches can be considered. For example, a straightforward (if inelegant) method denoted the Brute Force method can be used in which the (Euclidean) distance of all the nominal points in the database from the query point is estimated. The points that are within a predefined distance ε from the query points, i.e., $d(P_Q,P) < \varepsilon$, are retained, or alternatively only the closest point to P_Q, i.e.,

min($d(P_Q,P)$), can be retained. This approach scales as $O(N^2)$, where N is the number of nominal points in the database, and thus is clearly very time consuming.

The algorithm that we have implemented is referred to as nearest-neighbor-search (NNS) method in the present work and was proposed by Nene and Nayer [16]. The central idea is to bound the query point by successively trimming a list of candidate points. The algorithm attempts to identify a hyper-cube of size 2ε around the query point. The algorithm operates on unstructured data. The only provision is that the original database is transformed and is stored as a collection of 1D arrays such that the jth array contains the j coordinate of the points sorted in j. This is a process that takes place only once, and the associated overhead is negligible.

Once the data are sorted two maps are constructed. The backward one maps a coordinate in the ordered set to the corresponding coordinate in the original set and the forward one maps a coordinate in the point set to a coordinate in the ordered set. By this means, the coordinates in the ordered set that lie between the parallel hyper-planes positioned at $Q_1-\varepsilon$ and $Q_1+\varepsilon$ are efficiently identified with the search steps on the order of $O(\log_2(N))$. Using the backward map, we find the corresponding points in the point set and add appropriate points to the candidate list. Next we trim the candidate list by iterating through j=2, 3, ..., Ns+1, by checking every point in the candidate list with reference to the forward map, i.e., to see if its jth coordinate lies within the limits $Q_1-\varepsilon$ and $Q_1+\varepsilon$. More detailed discussions on this method can be found in Ref. [16]. We now dicuss the major steps of this method in the present implementation.

5.1. Generating non-empty sets of neighbors

It is apparent that the algorithm strongly depends on the choice of ε. Ref. [16] presented a number of approaches for estimating adequate bounds on ε so as to identify a non-empty set of neighbors. We are implementing the value that corresponds to the construction of the smallest hyper-cube [16]. More specifically the value of ε is estimated as:

$$\varepsilon = \frac{L}{2}\left(1-(1-p)^{\frac{1}{N}}\right)^{\frac{1}{D}}, \qquad (6)$$

where L is the range of the variable (L=|max-min|), N is the number of data points, D is the dimensionality (which equals $Ns+1$ in this particular problem), and p is the probability that the domain determined will be non-empty. In the current implementation, an iterative check is performed to guarantee that a specified minimum number of points are inside the neighborhood. If, for a given value of ε, the number of bounded points is smaller, the value of ε is increased by a factor α. Generally, the search scales linearly with the size of the database, N.

An important point has to be made at this point. The construction of the algorithm is based on the assumption that the L_∞ norm is used, as opposed to the conventional L_2 norm (Euclidean distance). This norm is chosen because the state variables can be different by orders of magnitude (for example, the difference between temperature and the mass fraction of a trace species). If L_2 norm is used, even with logarithmic transformation, the relative 'importance' of the state variable can still be obscured by variables that are of higher order magnitude. Appropriate scaling of the state variables should be explored in future work to promote more meaningful searching.

5.2. Assigning properties to the query point

Once a non-empty neighbor set is identified, properties should be assigned to the query point based on the properties of its neighbors. The simplest approach would be to augment the nearest neighbor search with an additional step that locates the point closest to the query point and assigns its properties to the query point. However, given that the final list of points is rather short, an improved approach is to linearly interpolate the properties of the points in the neighborhood of the query point and use the interpolation scheme to assign its properties. Linear interpolation is as good an approach as any other given fact that, by design, the neighborhood is quite restricted. Standard algorithms can be used for this purpose, resulting in an efficient computational step. It was empirically determined that averaging the property values in a neighborhood of the query point can potentially provide an advantage over simply assigning to the query point the property value of the point closest to it.

6. Test of the Adaptive Scheme

To test the integrated adaptive scheme, the GRI-Mech 3.0 without NO_x kinetics (36 species, 219 reactions) [14] was reduced to 18 submechanisms based on the two-stage reduction methodology. The dimensions of the reduced models are listed in Table 3. The reduced mechanism library, combined with the NNS method, was applied to the simulation of HCCI like combustion of stoichiometric methane/air mixture. A 2D numerical mesh of an HCCI engine was used in this test. The mesh has approximately 1000 cells at BDC position, as shown in Fig. 9. The calculation was carried out from -40 ATDC to 40 ATDC, where T_0 = 1000K and p_0 = 5 bar were used as initial conditions.

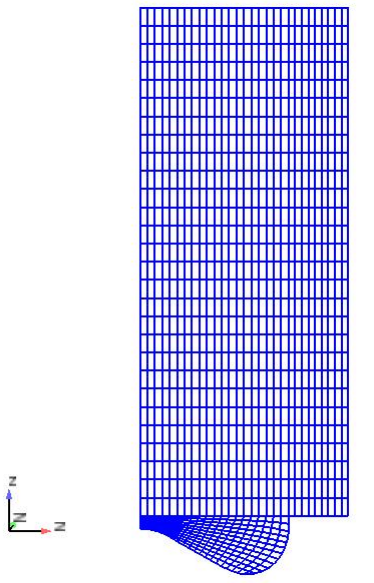

Figure 9: Numerical mesh of the HCCI engine.

Table 3: Dimensions of the CH_4 oxidation mechanisms.

Mechanism index	N_S	N_R
1	14	30
2	23	55
3	13	23
4	13	19
5	21	93
6	31	147
7	16	55
8	19	71
9	21	93
10	19	76
11	18	52
12	17	35
13	15	33
14	18	66
15	13	39
16	25	86
17	25	102
18	26	112
19 (the full mechanism)	36	219

Figure 10 compares the calculated in-cylinder pressure traces using the full mechanism and the adaptive scheme, respectively. The ignition timing predicted by the adaptive approach agrees well with that of the full mechanism simulation, however, the peak pressure and main heat release amount is much underpredicted. To diagnose this discrepancy, we examined the in-cylinder averaged mass fraction traces of several key species, including CH_4, O_2, CO_2, CO, OH, and H_2O_2, as shown in Fig. 12, where the symbols indicate full mechanism simulation and lines indicate adaptive chemistry results. It is clear that there is almost no accumulation of CO_2 and H_2O_2 before the major ignition in the adaptive simulations. Because H_2O_2 plays a crucial role in low temperature combustion, the omission of this species means that the reduced mechanism(s) invoked prior to the major heat release are 'ill-posed' for the low temperature scenario. Based on the mechanism utilization frequency histogram, as shown in Fig. 11, we identified that the responsible reduced mechanism was #16. This assessment was corroborated in a subsequent simulation in which the selection by the NNS algorithm of mechanism #16 was ignored, and the full mechanism used in its place. Results based on the modified adaptive scheme are compared with the full mechanism results in Fig. 13. The good agreement proves the effectiveness of the remedy. Meanwhile, as shown in Fig. 14, the pattern of mechanism utilization frequency is also changed dramatically, because the NNS code gave very different results under the updated thermo-chemical environments compared to the original calculation.

In Fig. 15, the CPU times of both the original calculation and the modified calculation were compared to that of the full mechanism simulation in which the NNS code was not invoked and therefore did not add in any CPU time overhead. As can be seen, the adaptive scheme achieves roughly 50% CPU time savings for this test case.

Although the modified calculation was generally successful, the failure of the original calculation reveals that the mechanism reduction process needs to be scrutinized more carefully to guarantee more rigorous feasibility.

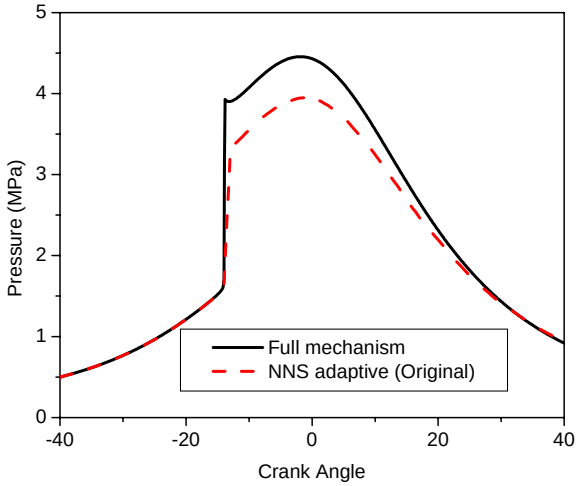

Figure 10: Comparison of pressure of full mechanism simulation and adaptive chemistry.

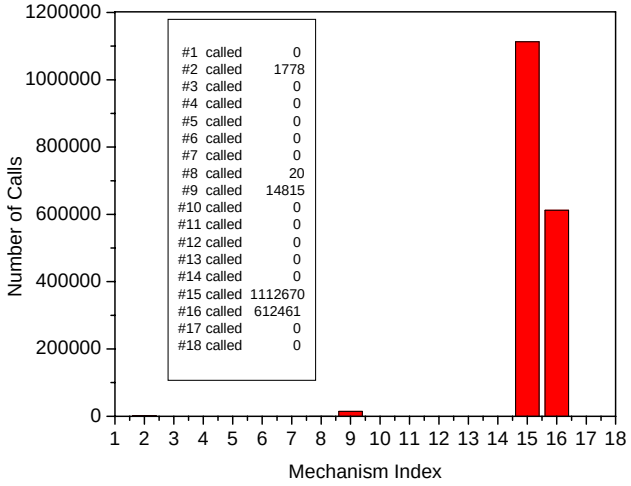

Figure 11: Histogram of mechanism utility frequency in adaptive scheme.

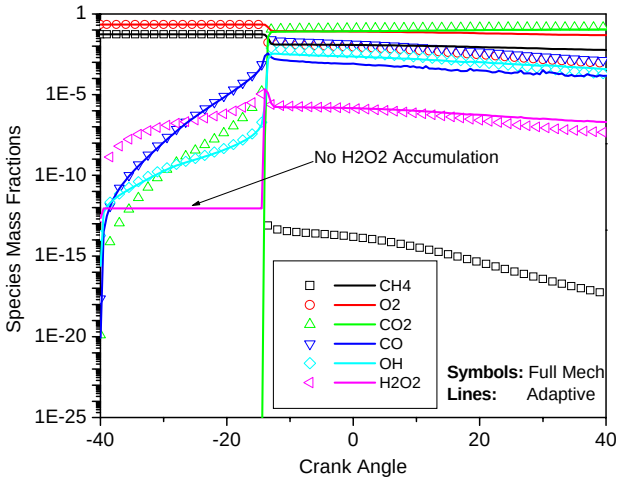

Figure 12: Species mass fractions of full mechanism simulation and adaptive chemistry.

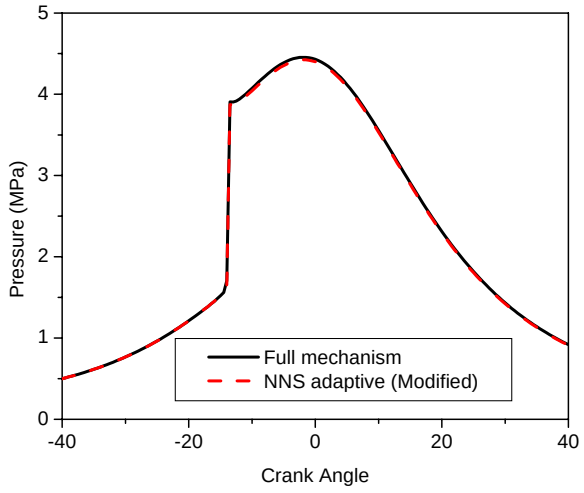

Figure 13: Comparison of pressure between full mechanism simulation and the modified adaptive chemistry calculation.

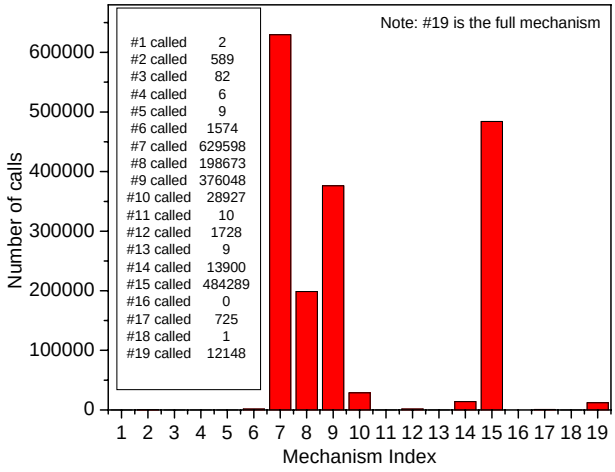

Figure 14: Histogram of mechanism utilization in the modified adaptive chemistry calculation.

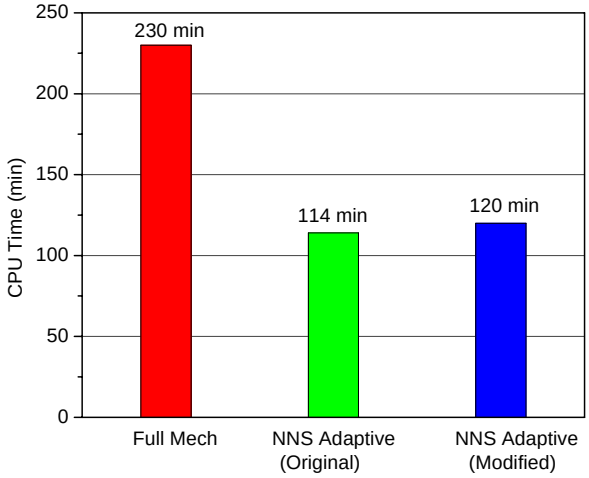

Figure 15: CPU time comparisons of the CH_4 system calculation.

7. CPU Time Savings by Species Array Redimensionalization

To illustrate the effect of species array redimensionalization, discussed in Section 2.2, on CPU time, the methane/air HCCI simulation studied in the previous section was reexamined with the difference that only one reduced model was used throughout the calculation. Two cases were compared: in the first, all the 36 species from the full mechanism were included in the ODE set formulation; in the second, the species array was redimensionalized so that only active species are kept. Mechanisms #4 (*Ns*=13) and #5 (*Ns*=21) were arbitrarily chosen for this test. Figures 16 and 17 show the comparisons of simulated pressure traces and CPU time, respectively. As expected, the simulated pressure traces match perfectly despite the different ODE set formulations. However, a reduction of 25-30% in CPU time was obtained in both test cases. This demonstrates the effectiveness of the species array redimensionalization. In actual calculations, CPU time savings will depend on the number of species in the reduced mechanism compared to the size of the full mechanism.

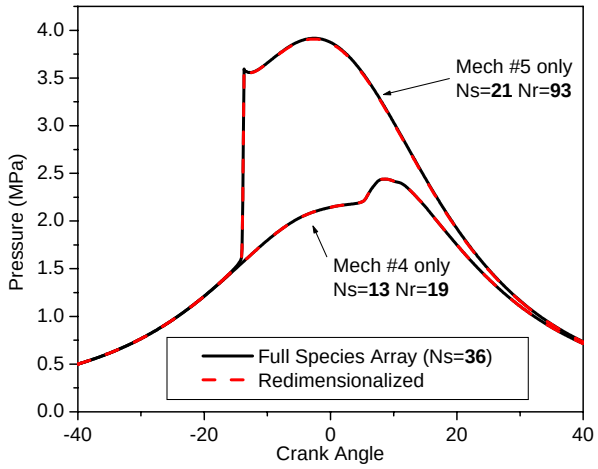

Figure 16: Species array redimensionalization effect on simulated pressure traces.

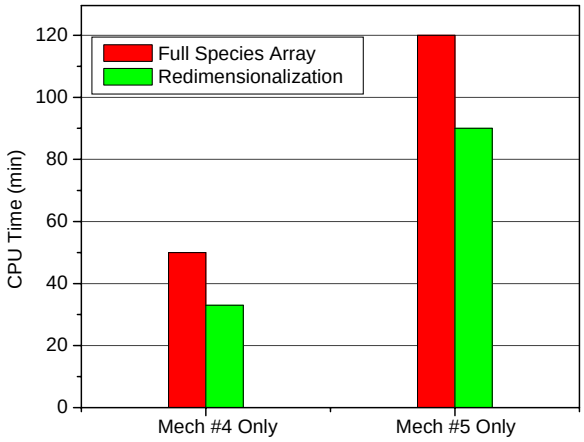

Figure 17: Species array redimensionalization effect on CPU times.

8. Summary

A methodology for adaptive coupling between chemical kinetics and CFD has been formulated. A two-stage mechanism reduction methodology and a nearest-neighbor-search approach were developed and implemented into the KIVA-3V CFD code. The integrated methodology was applied to a pentane oxidation mechanism and the GRI 3.0 methane oxidation mechanism. Compared to computations using full mechanisms, the adaptive chemistry approach shows great potential for significantly improving computational efficiency, while retaining reasonable accuracy for practical engineering calculations. In the present implementation, additional computational savings are realized by redimensioning the array of species mass fractions prior to integration, limiting the species included to those that occur in the given submechanism.

Based on the framework developed so far, the components of adaptive chemistry will be further investigated and refined in future work. For example, more rigorous feasibility analysis will be incorporated in both the mechanism reduction method and the query method. Additionally,

pressure will be included in the state vectors for query, and the nearest neighbor search approach will be optimized to enable higher query efficiency.

References

[1] S. R. Tonse, N. W. Moriarty, M. Frenklach, and N. Brown, *Int J. chem. Kinet* 35: 438-452, 2003.
[2] L. Liang, S. C. Kong, C. Jung, and R. D. Reitz, *J. Eng. for Gas Turbines and Power*, Vol. 129, 271-278, 2007.
[3] N. Peters, and F. A. Williams, *In Complex Chemical Reaction Systems*, Springer-Verlag: Berlin, 1987; Mathematical Modeling and Simulation, Vol. 47, 310-317.
[4] S. H. Lam, *Combust. Sci. Technol.*, Vol. 89, 375-404, 1993.
[5] I. P. Androulakis, J. M. Grenda, and J. W. Bozzelli, *AIChE J.* l50 (11), 2956, 2004.
[6] U. Maas, and S. B. Pope, *Combust. Flame*, 88: 239-264, 1992.
[7] T. Lu, and C. K. Law, *Proc. Combust. Inst.*, 30:1333-1341, 2005.
[8] S. B. Pope, *Combust. Theory Model.* 1 (1997) 41-63.
[9] I. Banerjee, and M.G. Ierapetritou, *Chem. Eng. Sci.*, 8, 4537, 2003.
[10] I. Banerjee and M. G. Ierapetritou, *Combust. Flame*, 144 (3), 619 (2006).
[11] R. Kee, F. Rupley, E. Meeks, J. A. Miller, CHEMKIN-III: A Fortran Chemical Kinetics Package for the Analysis of Gas-Phase Chemical and Plasma Kinetics, Sandia Report SAND 96-8216, 1996.
[12] A. A. Amsden, KIVA-3V: A Block-structured KIVA Program for Engines with Vertical or Canted Valves, Report LA-13313-MS, Los Alamos National Lab., 1997.
[13] H. J. Curran, P. Gaffuri, W. J. Pitz, and C. K. Westbrook, *Combust. Flame* 114:149-177, 1998.
[14] GRI-Mech 3.0, *http://www.me.berkeley.edu/gri-mech/*.
[15] J. Revel, J. C. Boettner, M. Cathonet, and J. S. Bachman, *J. Chem. Phys. Phys. Chem. Biol.,* 91: 365, 1994.
[16] S. A. Nene, and S. K. Nayar, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19: 989-1003, 1997.