

IMPROVEMENTS ON A NEWTON-KRYLOV BASED SOLVER FOR CFD MODELS USING FINITE RATE NOX CHEMISTRY

Qing Tang Martin Denison Mike Maguire
Mike Bockelie
Reaction Engineering International
Salt Lake City, UT, USA
www.reaction-eng.com

Jyh-Yuan Chen
Department of Mechanical Engineering
University of California at Berkeley
Berkeley, CA, USA

ABSTRACT

In this paper, we describe our progress on improving the performance of a newly developed Computational Fluid Dynamics (CFD) modeling tool, which uses reduced chemical kinetics mechanisms to model the finite rate chemistry effects and solves the resulting system of stiff partial differential equations with a matrix-free Newton-Krylov method. A multi-grid based preconditioner and a Newton iteration scheme have been implemented in the Newton-Krylov solver and the reduced mechanism module, respectively, to replace the original Picard based preconditioner and the point iteration scheme for steady state species evaluation. Preliminary tests of the improved modeling tool have been conducted using simple hotbox and a full-scale, coal fired electric utility boiler, and shown very promising results in terms of the accuracy, robustness, and efficiency of the new tool.

INTRODUCTION

Computer simulation plays an important role in US industry. However, CFD codes aimed at solving practical engineering problems involving chemically reacting flow have typically incorporated only simplified descriptions of the chemical processes involved. For example, chemical kinetic descriptions of hydrocarbon oxidation may require the tracking of hundreds of chemical species and thousands of reaction steps. CPU and memory limitations prohibit implementation of full detailed chemistry into CFD simulations.

Techniques are now available to create simplified chemical schemes that faithfully represent detailed chemical descriptions over an appropriate range of conditions using many fewer species or progress variables. Among them are: intrinsic low-dimensional manifold methods (ILDM) [1], rate-controlled constrained equilibrium (RCCE) [2], repro modeling [3], flame generated manifold methods (FGM) [4], and Roussel & Fraser

algorithm (RF) [5]. While each of these approaches can claim some success, and several have been extensively applied, none of these have achieved the level of applicability and universality of reduced mechanism methods based on quasi-steady-state assumption for a number of the chemical species (QSSA) [6-7]. In a paper by Chen [8], an automated mechanism reduction process was described that has been implemented into a computer code called CARM (Computer Assisted Reduction Method). CARM uses a set of input test problems to evaluate the error in the steady state assumption for each species in the detailed mechanism under the conditions of interest to determine the species for which the steady state approximation is applied. Thus, the CARM-generated reduced mechanisms retain a great deal of information from the original detailed mechanism while requiring the tracking of far few species. However, in the reduced mechanisms generated by CARM a set of nonlinear algebraic equations must be solved to obtain the steady state species concentrations, a computationally expensive operation. Also, reduced mechanisms may add significantly to the stiffness of the system, causing a substantial reduction in the convergence rate (or even failure) when incorporated into a CFD analysis as opposed to a global mechanism [9].

Recently a Newton-Krylov based CFD modeling tool has been developed by REI for performing engineering calculations of turbulent reacting flow [10], where the finite rate chemistry is represented by reduced mechanisms. The new tool employs a Newton-Krylov iteration scheme, which is suitable for solving large-scale systems containing severe non-linearity [11-14]. Moreover, the method only approximately solves the system of linear equations, and does not need to form the Jacobian matrix explicitly. Instead, the Jacobian matrix-vector products are estimated analytically or with finite differences of function values. Although on a per iteration basis Newton-Krylov

methods are considerably more expensive than traditional Picard iteration schemes, Newton-Krylov methods hold the promise for requiring fewer global iterations and allowing for addressing problems with more complex non-linearity where Picard iteration methods fail. This has been demonstrated by using the new modeling tool with NOx simulations for full scale, coal fired utility boilers [15]. Compared to the traditional solution scheme, the Newton-Krylov based solver provides up to a 60% reduction in CPU time, substantial improvement in robustness and only a modest increase in memory usage.

The objective of this paper is to report our progress on further improving the performance of the Newton-Krylov based CFD modeling tool. We describe, in order, an overview of the matrix free Newton-Krylov solver and its implementation in the CFD code, a Newton iteration scheme in the CARM generated reduced mechanism module, an Algebraic Multi-Grid (AMG) based preconditioner in the Newton-Krylov solver, and test results for the improved solver. The test cases for the impact of these modifications will be highlighted with NOx simulations for a hotbox test furnace and a full-scale coal-fired electric utility boiler.

NOMENCLATURE

A_p	coefficient in discretized equation
A_{nb}	coefficient in discretized equation
D_e	turbulent diffusivity
\mathbf{J}	Jacobian matrix
\mathbf{M}	preconditioning matrix
n_c	number of non-steady state species
n_{max}	maximum number of iteration
n_p	number of point relaxation iteration
n_s	number of chemical species
n_{ss}	number of steady state species
u	velocity
V_{cell}	cell volume
w_i	volumetric net reaction rate of species i
Y_i	mass fraction of species i

Greek Symbols

δ	small perturbation
ϵ_j	error in species j
ϵ_{AtoI}	absolute error tolerance
ϵ_{RtoI}	relative error tolerance
η	convergence tolerance
ρ	density

APPROACH OVERVIEW

REI's proprietary CFD codes have been used to simulate reacting and non-reacting flows of gases and particles, including gases diffusion flames, pulverized-coal flames, liquid sprays, coal slurries, reacting two-phase flow etc. One of the applications of this code is to predict trace levels of pollutants such as (NOx), which is of great interest in many industrial combustion processes. The heat release associated with these reactions is insignificant in comparison to that associated with the main combustion reactions. As a result, the solution of the pollutant related transport equations can be decoupled from the

solution of the main combustion process and carried out in a "post-process" fashion. The equations to be solved is

$$\nabla(\overline{\rho u Y_i}) + \nabla(D_e \nabla \bar{Y}_i) = \bar{w}_i, \quad (1)$$

with \bar{Y}_i being the Favre-mean mass fraction of species i and \bar{w}_i the time mean volumetric net reaction rate. The density $\bar{\rho}$, velocity \bar{u} , and turbulent diffusivity D_e are fixed according to the solution of the main combustion calculation. In discrete form, the transport equations take the form:

$$A_p \cdot Y_{pi} = \sum_{\substack{\text{neighbour} \\ \text{cell}}} A_{nb} \cdot Y_{nbi} + \bar{w}_i \cdot V_{cell} \quad (2)$$

where A_p , A_{nb} are coefficients due to convection and diffusion, and V_{cell} is cell volume.

The traditional solution scheme of the REI code for solving the species transport equations [9] is based on a Picard iteration scheme in which the governing partial differential equations are solved separately within a large iteration loop. For inner solves, the linear system is solved using a standard Alternating Direction Implicit (ADI) method. The chemical source terms calculated from the reduced mechanism are a highly nonlinear function of the species mass fraction. Hence, a very small under-relaxation factor for the chemical source terms has to be used to make the overall solution scheme successful.

The discretized species transport equation (2) can be rewritten in a more general form:

$$\mathbf{F}(\mathbf{x}) = [f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_n(\mathbf{x})]^T = 0, \quad (3)$$

where \mathbf{x} is an n -vector of species mass fraction, and n the total number of unknowns in the calculation domain. Newton's method can be applied to solve equation (3) by making a first-order Taylor series expansion, as shown in equation (4)

$$\mathbf{F}(\mathbf{x}^{k+1}) = \mathbf{F}(\mathbf{x}^k) + \mathbf{J} \times \mathbf{s}^k = 0, \quad (4)$$

where $\mathbf{J} = \mathbf{F}'(\mathbf{x}^k)$ is the Jacobian matrix of the system at the k^{th} iteration step; $\mathbf{F}(\mathbf{x}^k)$ is the non-linear system residual at the k^{th} step; and \mathbf{s}^k is the solution of the linear Newton equation (4) and is used to determine the current approximate of the solution of the non-linear system $\mathbf{x}^{k+1} = \mathbf{x}^k + \mathbf{s}^k$. In general, this linearized approximation is not very accurate except near the solution. Hence, it is not efficient to solve equation (2) exactly. In our implementation of the Newton-Krylov solver, an inexact Newton step \mathbf{s}^k is solved that satisfies the inexact Newton condition

$$\|\mathbf{F}(\mathbf{x}^k) + \mathbf{J} \times \mathbf{s}^k\| \leq \eta^k \|\mathbf{F}(\mathbf{x}^k)\| \quad (5)$$

where η^k is the convergence tolerance (forcing term) with a value of less than 1. The above process is continued until a

satisfactory solution is found by making $\|\mathbf{F}(\mathbf{x}^k)\|$ or $\|\mathbf{s}^k\|$ sufficiently small. We use Krylov subspace methods, the restarted generalize minimal residual (GMRES) algorithm [16-17], to solve the inexact linear Newton equation (5), in which only the matrix-vector product, $\mathbf{J} \times \mathbf{s}^k$, is required, and is calculated using a first-order finite differences approximation:

$$\mathbf{J} \times \mathbf{s}^k \approx \frac{\mathbf{F}(\mathbf{x}^k + \delta \mathbf{s}^k) - \mathbf{F}(\mathbf{x}^k)}{\delta}, \quad (6)$$

where δ is a small perturbation, and is chosen based on standard arguments to achieve accuracy by balancing estimates of floating point error and truncation error. It should be noted that only one extra function evaluation is required in equation (6). However, since the function evaluation involves reduced mechanism calculation, it is still a time consuming operation.

IMPROVE REDUCED MECHANISM MODULE

The reduced mechanism is constructed using the quasi-steady-state assumption (QSSA) introduced by Bodenstein [18], where given a detailed mechanism with n_s chemical species, a small number of n_c species are retained to represent the dynamics of the detailed system, while the other $n_{ss} = n_s - n_c$ species that are associated with fast processes are assumed to be in steady state with their net chemical production rates being set to zero. Then, in the turbulent combustion calculation, the relevant equations are solved for the n_c unsteady-state species instead of for the n_s species. Since the computational cost increases at least linearly with the number of species represented, substantial gains can be achieved for $n_c \ll n_s$.

In this research, the reduced mechanism module (a set of Fortran 77 subroutines) is generated automatically by CARM. The module, which calculates the chemical source term (i.e., the net production rates of unsteady-state species) defined by the reduced mechanism, can be linked easily to the Newton-Krylov CFD code. Timing studies have shown that approximately 50% of the total CPU time to perform a simulation with the Newton-Krylov solver is consumed in the evaluation of the chemical source terms for the reduced mechanism. When computing the chemical source terms more than 95% of the CPU time is spent in solving the algebraic non-linear system for the steady-state species concentrations.

Previously, a point relaxation scheme was used to solve the non-linear equation system for the steady-state species. It was found that the scheme fails for about 10% (on average) of the grid cells per global iteration in the finite volume solver. When the iteration fails, the latest estimate of the steady-state species mass fractions was used to compute the chemical source terms used in the transport equations for the non-steady state species. It has been hypothesized that the incomplete solves for the steady-state species is contributing to the severe "stiffness" encountered when using reduced mechanisms [19].

The point relaxation algorithm is used to solve a set of non-linear equations for the steady-state species. The equation system can be expressed as

$$\begin{cases} f_1(Y_1, Y_2, \dots, Y_{n_s}) = 0 \\ f_2(Y_1, Y_2, \dots, Y_{n_s}) = 0 \\ \vdots \\ f_{n_s}(Y_1, Y_2, \dots, Y_{n_s}) = 0 \end{cases}, \quad (7)$$

where f_i and Y_i are the net production rate and the mass fraction of the steady-state species i , respectively. Due to the special formation of the expression of f_i which is a non-linear function of the steady-state species mass fractions, equation (7) can be re-written as

$$\begin{cases} g_1(Y_2, Y_3, \dots, Y_{n_s}) \cdot Y_1 = h_1(Y_2, Y_3, \dots, Y_{n_s}) \\ g_2(Y_1, Y_3, \dots, Y_{n_s}) \cdot Y_2 = h_2(Y_1, Y_3, \dots, Y_{n_s}) \\ \vdots \\ g_{n_s}(Y_1, Y_2, \dots, Y_{n_s-1}) \cdot Y_{n_s} = h_{n_s}(Y_1, Y_2, \dots, Y_{n_s-1}) \end{cases}, \quad (8)$$

in which g_i and h_i are also non-linear functions of steady-state species mass fractions. In the point relaxation algorithm, the steady-state species mass fractions are solved using an iteration method, in which, for the $k+1$ iteration, the species mass fractions are calculated as

$$\begin{cases} Y_1^{k+1} = h_1(Y_2^k, Y_3^k, \dots, Y_{n_s}^k) / g_1(Y_2^k, Y_3^k, \dots, Y_{n_s}^k) \\ Y_2^{k+1} = h_2(Y_1^{k+1}, Y_3^k, \dots, Y_{n_s}^k) / g_2(Y_1^{k+1}, Y_3^k, \dots, Y_{n_s}^k) \\ \vdots \\ Y_{n_s}^{k+1} = h_{n_s}(Y_1^{k+1}, Y_2^{k+1}, \dots, Y_{n_s-1}^{k+1}) / g_{n_s}(Y_1^{k+1}, Y_2^{k+1}, \dots, Y_{n_s-1}^{k+1}) \end{cases} \quad (9)$$

The initial values of the steady-state species mass fraction are set to be zero.

The point relaxation algorithm is similar to the Gauss-Seidel algorithm for linear equation system. The characteristics of the Gauss-Seidel iterations are well studied in the literatures (e.g., [20]). However, when using the method to solve the non-linear problems encountered in the reduced mechanism, the behavior of the algorithm becomes very difficult to predict. It has been observed that even after hundreds of iterations, converged solutions (i.e., iteration errors are below given error tolerance) still cannot be achieved in the test calculations using this method. To overcome this problem, two major changes have been made in the reduced mechanism module. First, a new stopping criteria is implemented so that the iteration will stop at iteration $k+1$ ($k+1 \leq n_{max}$, n_{max} the maximum number of iteration allowed), if the following condition

$$\varepsilon_j \equiv |Y_j^{k+1} - Y_j^k| \leq \max(Y_j^{k+1} \cdot \varepsilon_{Rtol}, \varepsilon_{Atol}), \quad (10)$$

is satisfied for all steady-state species (i.e., $j=1, \dots, n_{ss}$). In equation (10) ε_{Rtol} and ε_{Atol} are user specified relative error tolerance and absolute error tolerance, respectively. Compared to the old stopping condition which reads

$$\varepsilon \equiv \max_{j=1, \dots, n_s} \left(\frac{Y_j^{k+1} - Y_j^k}{Y_j^k} \right) \leq \varepsilon_{Rot}, \quad (11)$$

the new stopping condition is more robust especially in the situation where species mass fraction is in the range of round-off error (which is very common in chemical kinetics calculations).

The second modification is the implementation of a Newton-Raphson based method to replace the old point relaxation scheme. The new algorithm is divided into two parts, and is referred to as the Newton scheme:

- First, the point relaxation iteration is performed for n_p times ($n_p = 5$ in current implementation).
- Second, a Newton-Raphson method is employed to solve the equation system using the results from the point relaxation iterations as the initial guess.

If the initial guess is close to the solution, the Newton-Raphson method converges to the solution quadratically. The stopping criteria for the Newton iterations are the same as those defined in equation (10).

AMG PRECONDITIONER

In the matrix free Newton-Krylov solver, a GMRES method is used to solve for the inexact Newton step which satisfies the inexact linear Newton equation. The choice of preconditioner used in the linear solves within the GMRES method plays a major role in the overall robustness and efficiency of the Newton-Krylov methods. Previous research has shown that, comparing to the test cases using a preconditioner, the cases without using a preconditioner took about twice as much as CPU time when applied with a global mechanism, and failed when applied with reduced mechanisms.

Solving the Newton equations (4) with a preconditioned iterative method requires computing the action of the inverse of a preconditioner \mathbf{M}

$$\mathbf{z} = \mathbf{M}^{-1} \mathbf{r} \quad (12)$$

at every iteration of the linear solver. In this, \mathbf{r} is some vector in the Krylov subspace generated by the Jacobian $\mathbf{F}'(\mathbf{x}^k)$, \mathbf{M} , and the initial residual $\mathbf{F}(\mathbf{x}^k)$. In a matrix-free Newton Krylov method, constructing \mathbf{M} can be a challenge since the matrix entries of the Jacobian are not available. In our previous implementation, a “frozen coefficient” approach was used where the preconditioner was formed from the system of PDEs being solved. The discretization schemes used in the REI models are very compact, resulting in a sparse matrix system. The required coefficients for the convection diffusion portion of the transport equations are available from the outer, non-linear iterations and thus readily available for re-use. The derivatives of the chemical sources from the reduced mechanism are approximated with a linearization. The (approximate) inverse of the “frozen coefficient” preconditioner is obtained through an incomplete LU (ILU) decomposition method by neglecting the strictly lower triangular part of the preconditioner \mathbf{M} . This approach for

handling the preconditioner, which we refer to as Picard preconditioner approach, was simple and was demonstrated to be quite effective [10]. However, as noted in [14], the solver performance could potentially be improved by using a multi-grid based preconditioner because the calculation of the ILU preconditioner can still be computationally expensive due to the enormous number of grid cells in typical simulations of industrial applications, and this cost is multiplied by the number of times the preconditioner is updated during the nonlinear solution process.

While multi-grid methods are highly efficient in solving linear systems on their own, they also serve as excellent preconditioners. Previously, an Algebraic Multi-Grid solver has been developed in REI for an Adaptive Mesh Refinement flow solver which has been successfully used for combustion applications. In this research, we re-use the AMG solver and integrated it into the preconditioning calls of the Newton Krylov solver to replace the original Picard type preconditioner. In contrast to the so-called “geometric” multi-grid methods, the AMG can be applied to certain types of (sparse) matrix equation, $\mathbf{A}\mathbf{u} = \mathbf{f}$, without explicitly referring to geometry and without requiring any predefined hierarchy, and therefore is more robust and can be applied in complex geometric situations such as the full scale utility boiler with embedded boundaries (i.e., inlets and no-slip walls located within the furnace interior).

In this work, both recursive V-cycle and F-cycle multi-grid algorithms have been implemented in the AMG solver to form the preconditioner. Figure 1 shows the sketch of a V-cycle. This implementation allows arbitrary levels of coarse cell. The coarse cells are created by grouping contiguous fine cells where maximum coupling (based on neighbor coefficients in the

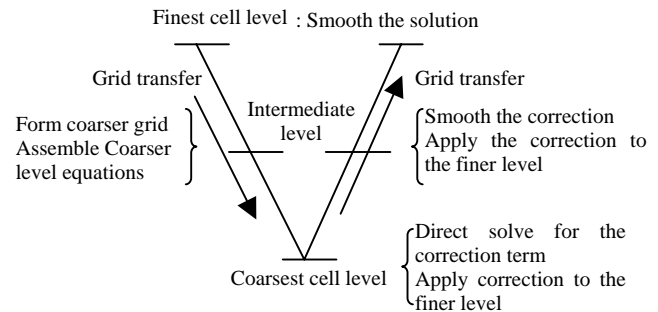


Figure 1. Sketch of a recursive multi-grid V-cycle

discretized partial differential equations) can be achieved. A minimum number of cells in each group is preset (8 for 3D). Figure 2 shows a sketch of the cell agglomeration strategy. It can be seen that the shape of the coarse level cells can be very “odd” based on the “strength of equation coupling” criterion.

In this implementation, defining the operator (i.e., matrix \mathbf{M}) on the coarsened grid is the key component. Because our CFD solver employs a mesh with embedded boundaries (i.e., grids are not body fitted), simply re-computing operators on the coarser grid can lead to difficulties in defining appropriate

boundary conditions. To circumvent this, we create the coarse grid operators using an AMG strategy with which REI has




Figure 2. Sketch of two-level multi grid system using AMG

experience. Briefly, using piecewise constant restriction and volume-average prolongation, the coarse grid operator can be constructed as a Galerkin operator in a finite volume agglomeration strategy, eliminating the need to estimate the Jacobian on coarse grid.

An aggregation AMG strategy is applied in the preconditioner. The idea is to apply a piecewise constant correction to each group of cells such that the sum of residuals over the group becomes zero. We consider a two-level simple system. Let h and H denote the fine and coarse levels, respectively. The task of the preconditioner is to find the solution of a linear system $\mathbf{A}^h \mathbf{x}^h = \mathbf{b}^h$, where \mathbf{A}^h and \mathbf{b}^h are coefficients defined at the fine cell level. Let \mathbf{x}^{h*} be an approximated solution, a piecewise constant correction $\mathbf{d}^h = \mathbf{I}_H^h \mathbf{d}^H$ is to be found such that the sum of residuals over a group of fine cells that form the coarse cell is zero:

$$\mathbf{I}_h^H (\mathbf{b}^h - \mathbf{A}^h (\mathbf{x}^{h*} + \mathbf{d}^h)) = 0. \quad (13)$$

In the above equations, \mathbf{I}_h^H is a Piecewise Constant operator and \mathbf{I}_h^H is an Addition operator. An aggregated equation that is solved at the coarse level can be derived from equation (13)

$$(\mathbf{I}_h^H \mathbf{A}^h) \mathbf{d}^H = \mathbf{I}_h^H \mathbf{r}^h, \quad (14)$$

where $\mathbf{r}^h = \mathbf{b}^h - \mathbf{A}^h \mathbf{x}^{h*}$ is the residual at the fine level cells. In our implementation, a smooth operator is applied on the intermediate cell levels to solve equation (14) in which one sweep of back substitution is performed by neglecting the strictly lower triangular part of matrix \mathbf{A} , and is followed by one sweep of line Gauss-Seidel iteration on the full matrix. At the coarsest cell level, a direct solve (LU decomposition) is used to solve equation (14) accurately.

TEST PROBLEM AND DISCUSSION

The test cases used in this study include a hotbox and a full-scale, coal-fired utility boiler. The hotbox case is industrial relevant and has been used in REI to test new models and concepts routinely. Figure 3 shows the sketch of the hotbox furnace where urea droplets are injected upstream to remove NOx from the combustion products. The modeled utility boiler

is a 138 MWe cyclone fired boiler with an overfire air (OFA) system. A furnace model for this cyclone boiler with a




Figure 3. Hotbox test case




Figure 4. Geometry of the modeled cyclone fired furnace

computational grid of $109 \times 55 \times 71$ (~426K cells) is shown in Figure 4. Also shown is the location of the monitor point where calculation results using different models and parameters are compared with each other.

The reduced mechanism used in this study is intended for simulating Selective Non-Catalytic Reduction (SNCR) of NOx [21] which contains ten non-steady state species (CO_2 , CO , O_2 , OH , H_2O , N_2 , NO , N_2O , HNCO , NH_3) and sixteen steady-state species.

TEST REDUCED MECHANISM MODULE

The test cases are listed in Table 1 Case 1 and 2 are the same except a larger n_{max} is used in Case 2 to avoid incomplete solves, and hence is referred to as the baseline case. Case 3 and Case 9 are essentially the same (with the new solver and the tightest error tolerances among all the cases) except the former one runs for more iterations. The old and new solvers refer to the point relaxation scheme and the Newton-Raphson scheme, respectively.

Figure 5 shows the history of the concentration of NOx at the outlet monitor point (see Figure 4) in all test runs in the utility boiler test problem. It can be seen that the concentrations converge to the same value in all cases except in Case 1 and Case 4. A 10% relative error in NO concentration is found at

Table 1. Test cases of reduced mechanism module

Case index	Stopping Criteria	Solver *
1	old, $n_{max} = 25$, $\epsilon_{Rel} = 10^{-5}$	old
2	old, $n_{max} = 200$, $\epsilon_{Rel} = 10^{-5}$	old
3	new, $n_{max} = 200$, $\epsilon_{Atol} = 10^{-15}$, $\epsilon_{Rel} = 10^{-9}$	new
4	new, $n_{max} = 200$, $\epsilon_{Atol} = 10^{-9}$, $\epsilon_{Rel} = 10^{-3}$	old
5	new, $n_{max} = 200$, $\epsilon_{Atol} = 10^{-12}$, $\epsilon_{Rel} = 10^{-6}$	old
6	new, $n_{max} = 200$, $\epsilon_{Atol} = 10^{-15}$, $\epsilon_{Rel} = 10^{-9}$	old
7	new, $n_{max} = 200$, $\epsilon_{Atol} = 10^{-9}$, $\epsilon_{Rel} = 10^{-3}$	new
8	new, $n_{max} = 200$, $\epsilon_{Atol} = 10^{-12}$, $\epsilon_{Rel} = 10^{-6}$	new
9	new, $n_{max} = 200$, $\epsilon_{Atol} = 10^{-15}$, $\epsilon_{Rel} = 10^{-9}$	new

* old: point relaxation scheme; new: Newton scheme

the monitor point in Case 1. The deviation is caused by the relatively small n_{max} used in the old reduced mechanism module that leads to many incomplete solves of the steady state species during the calculation. By increasing n_{max} by a factor of 10 (roughly) in Case 2, the accurate solution is achieved with the cost of much longer CPU time per global iteration. We have also noticed that the average number of iteration required by the old solver is about 40 for this test furnace.

Case 4 uses the same old solver as Case 1 and 2 but has new stopping criteria. A 20% relative error in NO concentration is observed at the monitor point. This is caused by the relatively large error tolerances adopted. By tightening the stopping criteria in Case 5 and 6, the error becomes negligible.

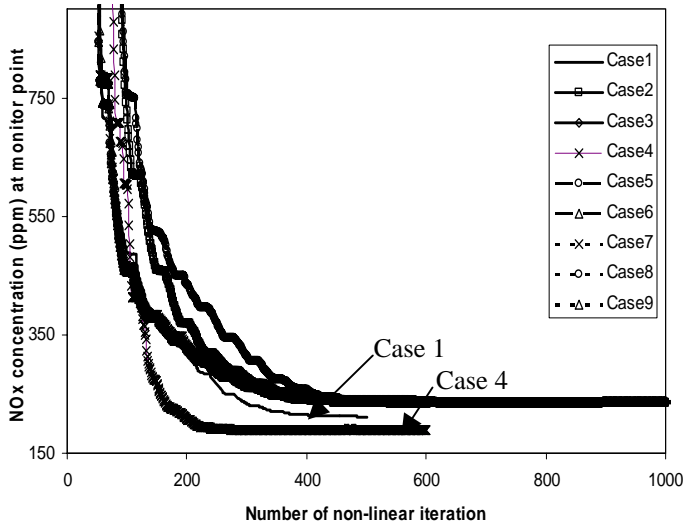


Figure 5. History of NOx concentrations at monitor point calculated using different reduced mechanism modules

All three cases using the new solver (Case 7, 8 and 9) achieve very good accuracy. It is interesting to note that using

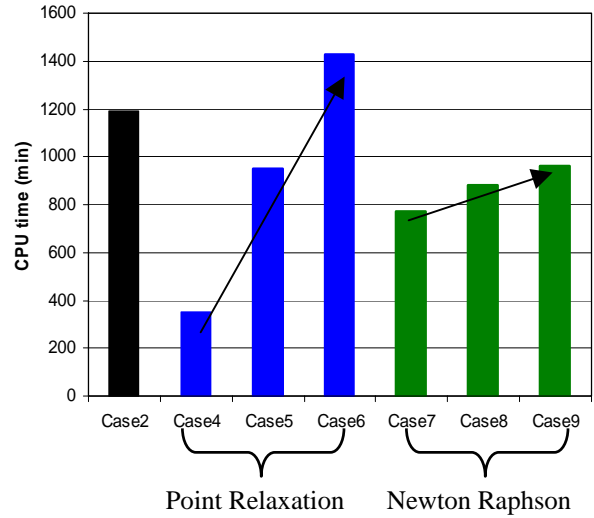


Figure 6. CPU times of the test cases to reach convergent solutions using different reduced mechanism modules

the same error tolerances as Case 4, Case 7 with the Newton method is very accurate as opposed to Case4.

Figure 6 shows the CPU times of Case 2 and Case 4 to 9. It is found that to reach the converged solutions, the numbers of Newton-Krylov global iterations required in all cases are about the same (around 400 iterations). However the CPU times required are quite different. First of all, Case 4 is excluded from the comparison due to its unacceptable accuracy although it requires the least amount of CPU time to reach global convergence. The case using Newton-Raphson scheme and the new stopping criteria with a moderate error tolerance (Case 8) saves about 25% CPU time compared to Case 2 (the baseline case), and saves about 10% CPU time compared to Case 5, which has the same error tolerance but uses the old point relaxation scheme. It can also be seen that when decreasing the error tolerance, the CPU times of the cases using point relaxation scheme increase dramatically, whereas the CPU times of the cases using Newton Raphson method only increase modestly.

Similar observations can be found in the results of the hotbox test case that are not shown here.

TEST AMG PRECONDITIONER

The AMG preconditioner has been tested against the original “Picard” type preconditioner in both hotbox case and the utility boiler case. The same boundary condition and reduced mechanism module are used in all test runs. We first define the species residual as

$$R_Y^w \equiv \left(\frac{\sum_{\text{Calculation domain}} R_Y^0}{\sum_{\text{Calculation domain}} \psi_Y} \right), \quad (15)$$

where the error residual R_Y^0 is defined using the finite-difference form of the species transport equation

$$R_Y^0 \equiv A_E Y_E + A_W Y_W + A_N Y_N + A_S Y_S + A_T Y_T + A_B Y_B + S_U - A_P Y_P \quad (16)$$

The definitions of the variables in equation 16 can be found in the reference [22]. The truncation term ψ_Y used in equation (15) to normalize the error residual is calculated as $\psi_Y = A_P Y_P$. The maximum of R_Y^w of all species finite difference equations is reported as the residual at current global non-linear iteration. This residual is normally used to determine the required level of convergence.

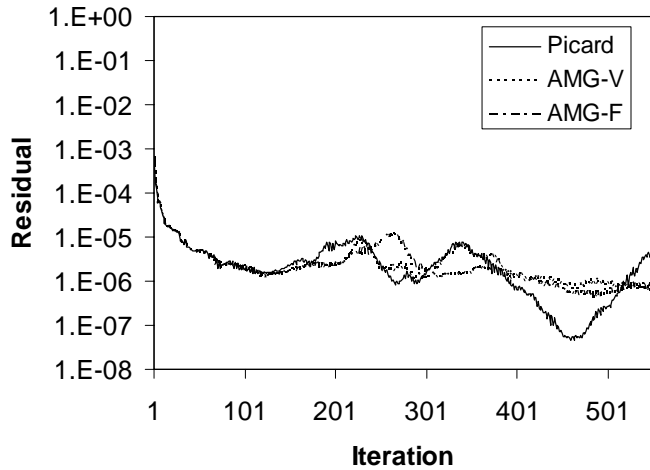


Figure 7. Residuals vs. number of non-linear iterations (hotbox test case)

Figure 7 shows the residue histories of three test runs of the hotbox problem using different preconditioners - a Picard preconditioner and two AMG preconditioners (V-cycle and F-cycle) - in the Newton-Krylov solver. It can be seen that for the first 100 non-linear iterations, the trends and the values of the residuals from both calculations are about the same. Substantial differences among the three runs are found when the calculations approach convergence. Large fluctuations in residual can be found in the case using the Picard preconditioner while the residuals in the other two cases are relatively stable. The CPU times (for 500 iterations) of the three runs are comparable. For this case, the AMG preconditioner has not been able to improve the computational efficiency of the Newton-Krylov solver significantly.

The investigation of the AMG preconditioner incorporated in the Newton-Krylov solver is still on-going. One area we are focus on is the agglomeration strategy used in the preconditioner to form coarse level cells. In the previous implementation of the AMG preconditioner (results shown in Figure 7), convection and diffusion terms from the discretized equations were agglomerated into a single term in the coarsened grid problem. It is expected that by separating the convective and diffusive fluxes and utilizes both within our “strength of equation coupling” criterion, improved convergence behavior can be achieved because unlike the previous implementation, the balance of the contributions of the convective and diffusive fluxes into the coarse grid problem is preserved.

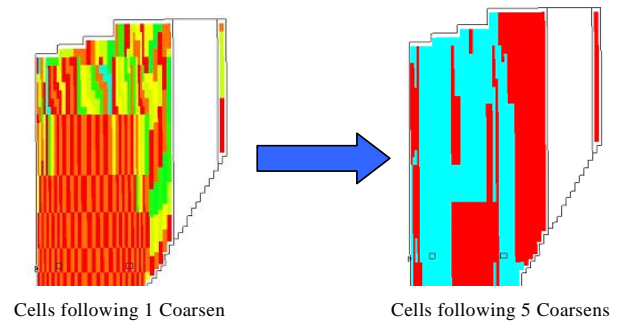


Figure 8. Fine grid cell agglomerations to create coarse grid (old method)

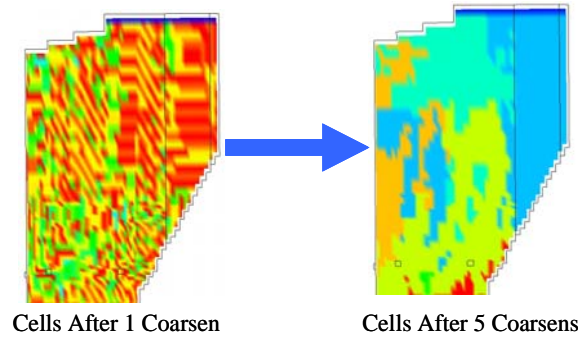


Figure 9. Fine grid cell agglomerations to create coarse grid (new method)

The impact of the new method on cell agglomeration is evident from Figure 8 and Figure 9. The same upper furnace sections are shown in the figures. The new method leads to very different patterns in the shapes of the coarse level cells (represented by different grey scale) compared to the old method. Investigation of the new method using the hotbox test furnace and full-scale industrial furnace is on-going.

SUMMARY

A combined point relaxation and Newton-Raphson scheme for evaluating the steady state species concentration has been implemented into a reduced chemical mechanism module for NOx prediction. The iteration stopping criteria are also redefined. The improved reduced mechanism is incorporated into the newly developed Newton-Krylov CFD solver, and is tested by modeling the NOx emission in a full scale, coal fired utility boiler. The calculation results have shown significant improvement in the accuracy of the simulation, and a 25% save in CPU time compared to the baseline case using the old reduced mechanism module. This is mainly due to the facts that there is no incomplete solve of the steady state species concentration when the combined algorithm is in place, and it only requires less than 3 Newton iterations (in average) to get converged solutions, even with a very tight error tolerance (e.g., $\epsilon_{A_{tol}} = 10^{-15}$, $\epsilon_{R_{tol}} = 10^{-9}$), whereas the old pure point relaxation scheme takes tens or even hundreds of iterations, and fails often.

It should be pointed out that the reduced mechanism used in this study is fairly small (i.e., only 16 steady state species). Other projects on-going within REI have indicated that as the number of steady state species increase, the computational cost for every Newton iteration may increase significantly (as expected) due to the larger and often dense Jacobian matrix in the Newton equation. In this circumstance, an optimal number of point relaxation iterations before the Newton iteration has to be identified (currently 5 iterations are used) to balance the numerical accuracy and the CPU time required. The new version of CARM software is now available which creates automatically the reduced mechanism module with the combined point relaxation and Newton-Raphson scheme for steady state species evaluation, and the new iteration stopping criteria.

Preconditioner based on AMG method has been implemented in the Newton-Krylov flow solver. Two different cell agglomeration strategies in the AMG preconditioner have been incorporated. In the first one, convection and diffusion terms from the discretized equations are agglomerated into a single term in the coarsened grid problem. Compared to the original Picard type preconditioner, the AMG preconditioner using this agglomeration strategy has only slightly improved the convergence behavior of the Newton-Krylov solver in the test problem. We believe better performance can be achieved by separating the convective and diffusive fluxes and utilizing both within our "strength of equation coupling" criterion that is used to determine how cells are agglomerated. Preliminary results have shown that the new method leads to very different patterns in the shapes of the coarse level cells compared to the old method. Investigation of the new method using a hotbox test furnace and full-scale industrial furnace is on-going.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under SBIR Phase II grant no. DMI-0216590. The NSF Program Manager is Dr. Errol Arkilic. The authors would like to thank Dr. Michael Pernice, Los Alamos National Laboratory, for his input and guidance in this work.

REFERENCES

[1] Mass, U. and Pope, S. B., 1992, "Simplifying chemical-kinetics – intrinsic low-dimensional manifolds in composition space," *Combustion and Flame*, 88, pp. 239-264.
 [2] Keck, J. C., 1990, "Rate-controlled constrained equilibrium theory of chemical reactions in complex systems," *Progress in Energy and Combustion Science*, 16, pp. 125-154.
 [3] Turanyi, T., 1994, "Parameterization of reaction mechanisms using orthogonal polynomials," *Computational Chemistry*, 18, pp.45-54.
 [4] van Oijen, J. A. and de Goey, L. P. H., 2000, "Modelling of premixed laminar flames using flamelet-generated manifolds," *Combustion Science and Technology*, 161, pp. 113-137.
 [5] Roussel, M. R. and Fraser, S. J., 1993, "Geometry of the steady-state approximation: perturbation and accelerated convergence methods," *Journal of Physics Chemistry*, 97, pp. 8316.

[6] Peters, N. and Rogg, B., (Eds.), 1993, *Reduced Kinetic Mechanisms for Applications in Combustion Systems*, Springer-Verlag, Berlin.
 [7] Smooke, M. D. (Ed.), 1991, *Reduced Kinetic Mechanisms and Asymptotic Approximations for Methane-Air Flames*, Springer-Verlag, Berlin.
 [8] Chen, J-Y., 1997, *Workshop on Numerical Aspects of Reduction in Chemical Kinetics*, CERMICS-ENPC, Cite Descartes-Champus sur Marne, France, 1997.
 [9] Cremer, M. A., Wang, D. H., Montgomery, C. J., and Adams, B. R., 2001, "Utilization of Reduced Mechanism Methods in CFD Simulations for Improved NOx Prediction in Utility Boilers and Furnaces," Joint AFRC/JFRC/IEA International Combustion Symposium, Kauai, Hawaii.
 [10] Wang, D. H., Bockelie, M. J., Cremer, M. A., and Chen, J-Y., 2002, "A Newton-Krylov Based Solver for Modeling Finite Rate Chemistry," Proceedings, 5th International Bi-annual Symposium on Computational Technologies for Fluid /Thermal /Structural/Chemical Systems with Industrial Applications, 448-1, pp. 113-120.
 [11] Smooke, M. D., Mitchel, R., and Keyes, D., 1989, "Numerical solutions of two-dimensional axisymmetric laminar diffusion flames," *Combustion Science and Technology*, 67, pp. 85-122.
 [12] Xu, Y., 1991, "Numerical Calculations of an Axisymmetric Laminar Diffusion Flame with Detailed and Reduced Reaction Mechanisms," Ph.D. thesis, Yale University, New Haven, CT.
 [13] Ern, A., 1994, "Vorticity-Velocity Modeling of Chemically Reacting Flows," Ph.D. thesis, Yale University, New Haven, CT.
 [14] Pernice, M., 2000, "A hybrid multigrid method method for the steady-state incompressible Navier-Stokes equations," *Elec. Trans. Num. Anal.*, 10.
 [15] Bockelie, M., Wang, D. H., Cremer, M. A., and Chen, J-Y., 2002, "A Newton-Krylov Based Solver for Modeling Finite Rate Chemistry in Reacting Flows," NSF SBIR grantees conference.
 [16] Saad, Y., Schultz, M., 1986, "A generalized minimal residual method for solving non-symmetric linear systems," *SIAM J. Sci. Statist. Comput.*, 7, pp. 856-869.
 [17] Freund, R., Golub, G., Nachtigal, N., 1992, "Iterative solution of linear systems," *Acta Numerica*, 1, pp. 57-100.
 [18] Bodenstein, M., and Lind, S.C., 1906, "Geschwindigkeit der bildung des bromwasserstoffs aus seinen elementen," *Z. Phys. Chem.*, 57, pp.168.
 [19] Chen, J-Y., 2002, private communication.
 [20] Golub, G. H., and van Loan, C. F., 1996, *Matrix Computations*, The Johns Hopkins University Press, Baltimore and London.
 [21] Cremer, M. A., Montgomery, C. J., Wang, D. H., Heap, M. P., and Chen, J-Y., 2000, "Development and Implementation of Reduced Chemistry for CFD Modeling of Selective Noncatalytic Reduction," *Proceeding of the Combustion Institute*, 28, pp. 2427.
 [22] Patankar, S. A., 1980, *Numerical Heat Transfer and Fluid Flow*, Hemisphere Publishing Corporation.