JOURNAL OF COMPUTATIONAL PHYSICS **123,** 84–95 (1996)
ARTICLE NO. 0007

An Implicit Multigrid Algorithm for Computing Hypersonic, Chemically Reacting Viscous Flows

JACK R. EDWARDS

Department of Mechanical and Aerospace Engineering, North Carolina State University, Raleigh, North Carolina

Received November 17, 1993; revised January 23, 1995

An implicit algorithm for computing viscous flows in chemical nonequilibrium is presented. Emphasis is placed on the numerical efficiency of the time integration scheme, both in terms of per-iteration workload and overall convergence rate. In this context, several techniques are introduced, including a stable, $\mathbb{O}(m^2)$ approximate factorization of the chemical source Jacobian and implementations of V-cycle and filtered multigrid acceleration methods. A five species–seventeen reaction air model is used to calculate hypersonic viscous flow over a cylinder at conditions corresponding to flight at 5 km/s, 60 km altitude and at 11.36 km/s, 76.42 km altitude. Inviscid calculations using an eleven-species reaction mechanism including ionization are presented for a case involving 11.37 km/s flow at an altitude of 84.6 km. Comparisons among various options for the implicit treatment of the chemical source terms and among different multilevel approaches for convergence acceleration are presented for all simulations.

1. INTRODUCTION

General combustion mechanisms and hypersonic flows at atmospheric re-entry conditions are examples of situations in which the fluid of interest must be considered in terms of its molecular constituents. The finite-rate chemical and thermodynamic processes that provide for production/depletion of each species and for transfer of energy among internal modes must be explicitly accounted for in the computation. In the most general continuum approach, an expanded Navier–Stokes set, containing separate transport equations for each species density and each internal energy mode, must be solved. By assuming equilibrium among all or some of the energy modes, one can reduce the number of equations, but the resulting system is still numerically ''stiff''due to the presence of widely varying time scales and can still be very large, depending on the complexity of the reaction mechanisms. With the advent of high-speed computers, fully coupled solutions of flowfields in chemical and thermal nonequilibrium are now possible (see, for example, [1–4]). However, well-resolved computations of complex reacting-gas flow phenomena still tax the limits of the available supercomputer technology, especially if three-dimensional simulations are considered. Although large strides have been made, it is clear that further improvements in algorithmic technology are needed before solutions of large-scale chemically reacting, viscous flowfields can be considered routine enough to be used in other than a research environment.

This article represents the beginning of a research effort designed to focus on ways to improve the numerical efficiency of algorithms for computing chemically reacting, viscous flow. The paper is outlined as follows. In Section 2, the governing equations are briefly outlined and the discretization approach is described. The basic implicit time integration scheme, a point Gauss–Seidel approach, is presented in Section 3. Several methods for the implicit treatment of the chemical source term are outlined in Section 4, including a new $\mathbb{O}(m^2)$ approximate factorization approach (m is the number of chemical species). Section 5 describes the implementation of two multigrid schemes, a V-cycle approach and a parallelizable ''filtered'' method, designed to accelerate the convergence of the baseline reacting-gas solver. Several computations of two-dimensional inviscid and viscous reacting flows at atmospheric re-entry conditions are presented in Section 6 to illustrate the robustness and efficiency of the developed algorithm.

2. GOVERNING EQUATIONS AND DISCRETIZATION

In terms of a curvilinear coordinate system defined by the steady transformation $\xi = \xi(x, y)$, $\eta = \eta(x, y)$ the expanded Navier–Stokes set governing flows in thermo-chemical nonequilibrium can be written as

$$\frac{\partial}{\partial t}\left(\frac{U}{J}\right) + \frac{\partial (F_\xi - F_\xi^v)}{\partial \xi} + \frac{\partial (F_\eta - F_\eta^v)}{\partial \eta} - \frac{\dot{\omega}}{J} = 0, \qquad (1)$$

where U is the vector of conserved variables, $F_{\xi,\eta}$ are the inviscid fluxes, $F_{\xi,\eta}^v$ are the viscous fluxes, and $\dot{\omega}$ is the source vector. In the above, J is the Jacobian of the transformation; $1/J$ is a measure of the cell volume. Derivations of the terms within the vectors have been presented in other references [5, 6]. A detailed description will therefore be omitted, and only the highlights of the current approach will be outlined.

Two reacting-air mechanisms are considered in this study. In the first, used for the viscous computations, separate species

0021-9991/96 $12.00

density equations are solved for O_2, O, N_2, N, and NO. The law of mass action is used to formulate appropriate source terms describing the production or destruction of a particular species. As is commonly done, an Arrhenius form is adopted for the forward reaction rates. Curve-fit coefficients for the forward rates and the associated equilibrium constants are taken from Park [7]. Expressions for the species viscosities are taken from Blottner [8], and thermal conductivities for each species are found by assuming an Eckert relation. The mixture viscosity and thermal conductivity are determined using Wilke's law [9]. The diffusion velocities and the associated heat flux terms are modeled by using Fick's law. One diffusion coefficient, related to the viscosity and density via the assumption of a constant Schmidt number, is used for all species. Thermal equilibrium at the translational temperature is assumed, and a harmonic oscillator model is used to characterize the vibrational energy modes.

In the second reaction set, additional equations for the ionic species O_2^+, O^+, N_2^+, N^+, and NO^+ are solved, as is an equation for the electron density. Charge exchange, associative ionization, and electron-impact ionization reactions are also considered in the formation of the source terms [10]. The energy for electronic excitation, assumed to be significant for O, O_2, and N, is modeled using a Boltzmann distribution of electronic states again characterized by the translational temperature [3]. Viscous effects are neglected in the 11-species calculations.

The governing equations are discretized using a control volume approach. Viscous and diffusion terms are central-differenced to second-order accuracy, while first-order upwinding based on flux vector splitting is used for the inviscid fluxes. The AUSM [12, 13] flux-splitting technique is utilized for the 5-species viscous calculations. For reasons outlined later, an extension of the more dissipative Van Leer (Hänel) approach [14] is employed for the 11-species inviscid calculations. To highlight the similarities and differences between the two flux-splitting formulations, we first write the inviscid flux vector in the kth coordinate direction ($k = \xi, \eta$) as a sum of convective and pressure contributions:

$$F_k \equiv F^c + F^p = \frac{|\nabla k|}{J} M\tilde{F}^c + \frac{|\nabla k|}{J} p\tilde{F}^p, \tag{2}$$

where

$$\tilde{F}^c = \begin{bmatrix} \rho_1 a \\ \vdots \\ \rho_m a \\ \rho a u \\ \rho a v \\ aH \end{bmatrix}, \quad \tilde{F}^p = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \tilde{k}_x \\ \tilde{k}_y \\ 0 \end{bmatrix}, \tag{3}$$

$$\tilde{k}_{x,y} = \frac{k_{x,y}}{|\nabla k|}, \tag{4}$$

$$H \equiv (E_t + p). \tag{5}$$

The contravariant Mach number M is given as

$$M = \frac{1}{a}[\tilde{k}_x u + \tilde{k}_y v],$$

where a is the frozen speed of sound.

Following Refs. [12, 15], we define a compact form for the interface flux $F_{1/2} \equiv F^c_{1/2} + F^p_{1/2}$ by considering the behavior of the convective and pressure components separately. The convective portion of the interface flux $F^c_{1/2}$ is given by

$$F^c_{1/2} = \frac{|\nabla k|}{J}[C^+\tilde{F}^c_L + C^-\tilde{F}^c_R], \tag{7}$$

where for the Van Leer (Hänel) scheme,

$$C^+ \equiv C^+_{\mathrm{VL}} = \alpha^+_{\mathrm{L}}(1 + \beta_{\mathrm{L}})M_{\mathrm{L}} - \frac{\beta_{\mathrm{L}}}{4}(M_{\mathrm{L}} + 1)^2 \tag{8}$$

and

$$C^- \equiv C^-_{\mathrm{VL}} = \alpha^-_{\mathrm{R}}(1 + \beta_{\mathrm{R}})M_{\mathrm{R}} + \frac{\beta_{\mathrm{R}}}{4}(M_{\mathrm{R}} - 1)^2. \tag{9}$$

From these definitions, the AUSM representations for $C^\pm$ are given as

$$C^\pm = \tfrac{1}{2}(M_{1/2} \pm |M_{1/2}|), \tag{10}$$

where

$$M_{1/2} = C^+_{\mathrm{VL}} + C^-_{\mathrm{VL}}. \tag{11}$$

The functions $\alpha^\pm$ and β provide the correct sonic-point transition behavior [15]:

$$\alpha^\pm_{\mathrm{L,R}} = \tfrac{1}{2}[1 \pm \mathrm{sgn}(M_{\mathrm{L,R}})], \tag{12}$$

$$\beta_{\mathrm{L,R}} = -\max[0, 1 - \mathrm{int}(|M_{\mathrm{L,R}}|)]. \tag{13}$$

For both schemes, the pressure contribution to the interface flux is defined as

$$F^p_{1/2} = \frac{|\nabla k|}{J}\tilde{F}^p[D^+_{\mathrm{L}}p_{\mathrm{L}} + D^-_{\mathrm{R}}p_{\mathrm{R}}], \tag{14}$$

where

$$D^\pm_{\mathrm{L,R}} = \alpha^\pm_{\mathrm{L,R}}(1 + \beta_{\mathrm{L,R}}) - \frac{\beta_{\mathrm{L,R}}}{2}(1 \pm M_{\mathrm{L,R}}). \tag{15}$$

This expression corresponds to the simple pressure splitting defined in Ref. [12].

In each of the expressions above, the L and R notations represent the dependence of the quantity on the fluid properties at left or right states about the interface. Unlike that of the Van Leer (Hänel) approach, the convective portion of the AUSM interface flux vanishes as M_L and M_R both approach zero, allowing the AUSM to capture shear layers quite accurately [12]. A side effect of the AUSM convective-flux formulation (illustrated later) is the possibility of a nonmonotone capturing of strong shock waves.

3. DIAGONAL IMPLICIT APPROACH

The basic time integration scheme used in this investigation is the LU-SGS algorithm of Yoon and Jameson [16]. In its perfect gas form, a diagonally dominant implicit operator based on a point Gauss–Seidel matrix splitting is constructed so that each lower–upper sweep only requires scalar diagonal matrix inversions. The CPU time per iteration for the LU-SGS approach is on the order of that of explicit schemes, and the technique is unconditionally stable for many perfect-gas problems. The requirement for diagonal dominance can lead to slow convergence rates for viscous computations involving a high degree of mesh clustering. Nevertheless, the algorithm has been successfully applied to low speed and high speed flows in two and three dimensions [17, 18]. Only a brief description of the basic LU-SGS approach, modified to include the implicit treatment of the chemical source terms, is presented herein. More detailed information regarding the construction of the implicit operator may be found in [17, 19, 20].

Given that $\Delta\xi$ and $\Delta\eta$ are equal to unity, the LU-SGS procedure for advancing the solution vector U from time levels k to $k + 1$ is as follows:

Step 1. Lower sweep:

$$\delta U^*_{i,j} = (d_s I + W)^{-1}_{i,j}[-\mathscr{F}_{i,j} + A^+_{i-1,j}\delta U^*_{i-1,j} + B^+_{i,j-1}\delta U^*_{i,j-1}]$$

Step 2. Upper sweep:

$$\delta U_{i,j} = \delta U^*_{i,j} - (d_s I + W)^{-1}_{i,j}[A^-_{i+1,j}\delta U_{i+1,j} + B^-_{i,j+1}\delta U_{i,j+1}] \quad (16)$$

Step 3. Update:

$$U^{k+1}_{i,j} = U^k_{i,j} + \delta U_{i,j}.$$

In the above, $\mathscr{F}_{i,j}$ represents the discretized residual (or collected defect error for the multigrid procedures) evaluated at point i, j. The split Jacobian matrices $A^\pm$ and $B^\pm$ are defined as

$$A^\pm = \tfrac{1}{2}(A \pm \sigma_A I)$$

and

$$B^\pm = \tfrac{1}{2}(B \pm \sigma_B I),$$

where A and B are the inviscid flux Jacobian matrices corresponding to the ξ and η directions and $\sigma_{A,B}$ are the associated spectral radii. The scalar d_s is given as

$$d_s = \frac{1}{J\,\Delta t} + \sigma_A + \sigma_B, \quad (17)$$

and W represents the Jacobian of the chemical source vector

$$W = -\frac{1}{J}\frac{\partial \dot{\omega}}{\partial U} \quad (18)$$

Although not indicated in the above, ''viscous'' spectral radii are added to their appropriate places within the implicit operator [21]. Algorithm (16) is used as the smoother for the multigrid procedures to be discussed later. Boundary conditions are explicitly updated on the finest grid following Step 3 of (16).

4. SOURCE JACOBIAN APPROXIMATIONS

It is obvious that its use of scalar diagonal matrix inversions, as opposed to block inversions, makes the LU-SGS approach potentially attractive for solving the expanded reacting-gas Navier–Stokes set. The key to the successful extension of the LU-SGS scheme lies in the implicit treatment of the chemical source terms, which account for the mass and energy transfer rates. Several methods for considering these terms have appeared in the recent literature. Shuen and Yoon [20] and Park and Yoon [22] linearize the source terms corresponding to the species continuity equations with respect to the species densities (or mass fractions). The resulting $m \times m$ Jacobian matrix (W in Eq. (18)) is inverted directly at every grid point to yield corrections for the species densities. Scalar diagonal inversions are used for the remainder of the equation set. The work per iteration for the LU-SGS algorithms of [20, 22] is further reduced by only performing the matrix inversion within the lower sweep. For this situation, $(d_s I + W)^{-1}$ in the upper sweep of (16) would simply be replaced by the scalar inversion $(d_s I)^{-1}$. Most of the calculations presented herein employ this convention; cases requiring source Jacobian inversions in both sweeps will be especially noted in Section 6.

A simpler approach for ''diagonalizing'' the source term Jacobian within the LU-SGS framework has been implemented and tested by Eberhardt and Imlay [4]. In this method, a ''time scale'' for each reaction is defined by taking the L_2 norm of the appropriate row elements within the source Jacobian. In particular, if $W_{k,l}$ represents an element of the source Jacobian described above, the Eberhardt and Imlay scalar diagonal matrix is given by

$$\text{diag}\left[d_s + \frac{1}{\tau_1}, d_s + \frac{1}{\tau_2}, ..., d_s + \frac{1}{\tau_m}, d_s, d_s, d_s\right], \tag{19}$$

where

$$\frac{1}{\tau_k} = \beta \sqrt{\sum_{l=1}^{m} (W_{k,l})^2} \tag{20}$$

This formulation always adds to the diagonal dominance of the LU-SGS implicit operator and thus will always underrelax the effects of the chemical source term. Candler and Olynick [23] have observed that element nonconservation and generally poor convergence behavior can result if β is not chosen properly. They propose an alternative linearization that maintains elemental conservation and leads to better numerical behavior.

An alternative approach for approximating the source Jacobian inversion, somewhere between the Shuen–Yoon and Eberhardt–Imlay methods in complexity, can be derived by first realizing that the following steps must be taken (or approximated in some fashion) to provide a solution update at every grid point (**b** represents the terms within the square brackets of algorithm (16)):

Step 1.

$$\begin{bmatrix} \Delta\rho_1 \\ \Delta\rho_2 \\ \vdots \\ \Delta\rho_m \end{bmatrix} = (d_s I + W)^{-1} \begin{bmatrix} b_1 \\ b_2 \\ \vdots \\ b_m \end{bmatrix},$$

Step 2.

$$\begin{bmatrix} \Delta\rho u \\ \Delta\rho v \\ \Delta E_t \end{bmatrix} = \frac{1}{d_s} \begin{bmatrix} b_{m+1} \\ b_{m+2} \\ b_{m+3} \end{bmatrix}. \tag{21}$$

The second step of (21) is trivial; we attempt, therefore, to find a less-costly way of obtaining the matrix–vector product $(d_s I + W)^{-1}\mathbf{b}$. This can be accomplished by approximately factoring this system in terms of the diagonal matrix

$$\tilde{d} = \text{diag}[d_s + W_{1,1}, d_s + W_{2,2}, ..., d_s + W_{m,m}] \tag{22}$$

and the off-diagonal *columns* of the source Jacobian W.
The result, obtained by applying the Sherman–Morrison theorem [24] to each factor, is given by the following product of rank-one matrices,

$$(d_s I + W)^{-1} \approx \left[\prod_{i=1}^{m} (I - \alpha \tilde{d}^{-1} \mathbf{u}_i \mathbf{v}_i^T)\right] \tilde{d}^{-1}, \tag{23}$$

where

$$\mathbf{u}_i^T = [W_{1,i}, ..., W_{i-1,i}, 0, W_{i+1,i} ..., W_{m,i}] \tag{24}$$

and

$$\mathbf{v}_i^T = [0, ..., 0, 1(i\text{th place}), 0, ..., 0]. \tag{25}$$

The product $(d_s I + W)^{-1}\mathbf{b}$ is then approximated by successive multiplications of the rank-one matrices defined in (23). Each multiplication is $\mathbb{O}(m)$ since $\tilde{d}$ is a scalar diagonal matrix, making the overall cost $\mathbb{O}(m^2)$. For large m, this scheme should provide savings over the direct LU decomposition approach ($\mathbb{O}(m^3)$). The stability of the approach will depend strongly upon the magnitudes of the elements of W, the ordering of the species continuity equations (O_2, O, N_2, N, NO, $[O_2^+, O^+, N_2^+, N^+, NO^+, e^-]$ for this work), and the magnitude of d_s. It is clear that this approach will not necessarily maintain diagonal dominance, as does the Eberhardt–Imlay approach, but it will provide a closer approximation to the actual matrix inversion. For the cases considered herein, computational experiments have shown that the use of an appropriately defined underrelaxation parameter α renders the scheme as stable as the direct inversion approach. The choice of α used for the results presented in this work is

$$\alpha = 1.0 - 0.4 \min(1.0, \Psi), \tag{26}$$

where the local Damköhler number Ψ is defined by

$$\Psi = \frac{L_{\text{ref}} \max(|\dot{\omega}_k|, k = 1, m)}{\rho_\infty U_\infty}. \tag{27}$$

The results presented later using this approach are preliminary but are very encouraging. Incidentally, a similar ''approximate inverse'' can be constructed by considering each *row* of W as a rank-one perturbation. This approach was tried but was found to be almost unconditionally unstable.

5. MULTIGRID APPROACH

The second part of this article involves an evaluation of the effectiveness of multigrid methods as means of accelerating the convergence of the implicit reacting-gas solver. The multigrid ''philosophy'' is based on the fact that most commonly used iterative schemes for solving grid-discretized partial differential equations are adept at smoothing or eliminating only certain ranges of high frequency error components. Low frequency errors are generally not damped, resulting in slow convergence rates for fine grids. On a coarser grid, formed perhaps from the removal of every other fine-grid line in all coordinate directions, the fine-grid low frequency errors are resolved as higher frequencies. Smoothing passes on the coarse grid can effectively damp these frequencies, generating a ''coarse-grid correction'' to the low frequency spectrum. The recursive application of this

technique on successively coarser grids, known as a ''multigrid cycle,'' can in principle eliminate all error modes, dramatically accelerating convergence. The effectiveness of multigrid is directly related to how well the basic iterative scheme attenuates the high-frequency errors resulting from the interpolation of the coarse-grid corrections to the finer grids.

The full approximation-storage (FAS) concept of Brandt [25] provides a means of generalizing the multigrid idea to nonlinear, nonelliptic problems. The utility of FAS-based multigrid methods in solving the Navier–Stokes equations for subsonic and transonic flow problems has been well-documented [26, 27]. Several recent efforts, [28–30] for example, have focused on developing multigrid techniques for computing supersonic and hypersonic viscous flowfields, with varying degrees of success having been reported. In general, though, it appears that most current multigrid approaches for computing high-speed flows are more labor-intensive than their low-speed counterparts. This is primarily the result of additional smoothing needed to reduce high-frequency errors in the vicinity of strong shock waves and the use of complicated cycling strategies and grid transfer operators. To this author's knowledge, only one multigrid-based algorithm for computing hypersonic, chemically reacting flows has been documented in the literature [31]. The primary conclusion from the obtained inviscid results was that the performance of multigrid is relatively unaffected by chemical reactions but is strongly dependent on the hyperbolic nature of the problem.

In this investigation, the implicit algorithms outlined earlier are used as smoothers for a V-cycle FAS multigrid approach and for a filtered multigrid scheme [32, 33]. The underlying coarse grids (three in this investigation) are determined by successively removing each fine grid line in both coordinate directions (full coarsening). Fairly standard grid-transfer operators are employed for both approaches. As the basic discretization approximates a cell-vertex scheme in 2D, the restriction of the fluid properties to the coarser grids is accomplished via simple injection. A conservative restriction of the residual error is achieved by a weighted summation of the fine-grid values:

$$\begin{aligned}\tilde{r}[\mathscr{F}_{ic,jc}] = \mathscr{F}_{i,j} &+ \tfrac{1}{2}(\mathscr{F}_{i,j+1} + \mathscr{F}_{i,j-1} + \mathscr{F}_{i+1,j} + \mathscr{F}_{i-1,j}) \\ &+ \tfrac{1}{4}(\mathscr{F}_{i+1,j+1} + \mathscr{F}_{i+1,j-1} + \mathscr{F}_{i-1,j+1} + \mathscr{F}_{i-1,j-1}).\end{aligned} \tag{28}$$

Values at the fine grid nodes i, j, etc. can be accessed from the coarse grid nodes ic, jc by the rule

$$i, j = 2ic - 1, 2jc - 1.$$

The prolongation operator for transferring the coarse-grid corrections to the next finest grid is simple bilinear interpolation. Only one implicit smoothing pass per coarse-grid level is performed for the V-cycle approach, and post-smoothing of residuals or corrections is not utilized. The source Jacobian $\partial\dot{\omega}/\partial U$ is computed on the finest level and is restricted to the coarser levels by injection.

To enhance the performance of both schemes in computing flows with strong shock waves, a slight modification to the error calculation procedure at each grid level is made. Given that $\tilde{r}$ represents a restriction operator, N_k represents the implicit smoothing operation on grid level k, and that $\mathscr{R}(U_k)$ is the Navier–Stokes residual at level k, the following sequence of operations is performed within the FAS V-cycle: ($\tilde{p}$ represents the prolongation operator):

Step 1. Restrict solution values from levels $k - 1$ to k,

$$U_k = \tilde{r}[U_{k-1}]$$

Step 2. Smooth to determine correction ($\mathscr{F}_{k-1}$ is the collected error at the previous grid level),

$$\text{solve: } N_k\, \delta U_k = \tilde{r}[\mathscr{F}_{k-1}]$$

Step 3. Calculate residual error at level k,

$$\mathscr{F}_k = \tilde{r}[\mathscr{F}_{k-1}] - \frac{1}{\varepsilon}(\mathscr{R}(U_k + \varepsilon\, \delta U_k) - \mathscr{R}(U_k)) \tag{29}$$

Step 4. Restrict residual error $\mathscr{F}_k$ to next coarsest level and continue cycling process

Step 5. Sum corrections at grid level k,

$$\delta U_k = \delta U_k + \tilde{p}[\delta U_{k+1}].$$

The term in parentheses within Step 3 represents a Frechét derivative approximation to the matrix–vector product $(\partial\mathscr{R}/\partial U_k)\delta U_k$. As ε becomes small, the accuracy of the approximation increases and the residual error calculation approaches that of a *linear* multigrid method. As ε approaches unity, the standard FAS scheme is recovered. The effect of $\varepsilon < 1$ is to minimize the direct effects of large coarse-grid solution changes, which result from the combination of strong discontinuities and low resolution levels. This simple correction seems to obviate the need for the extra smoothing steps and complicated initialization procedures found in some ''hyperbolic'' multigrid algorithms. For the blunt body calculations presented herein, it has been found that the choice of $\varepsilon = 1.0$ is clearly nonoptimal in the transient stages of the iteration sequence. As convergence is approached, however, the performance of the schemes varies little for ε ranging from 1.0 to 10^{-6}.

The filtered multigrid algorithm was developed by Chan and Tuminaro [31] as a means of efficiently utilizing processors of a distributed memory architecture. At each grid level of a standard multigrid approach implemented on such a machine, the number of grid points decreases, leading to an unavoidable reduction in parallel efficiency as processors become idle. The filtered multigrid method avoids this deficiency by constructing multi-

ple subproblems at each grid level via a frequency decomposition. Previously idle processors can then be used to solve the subproblems concurrently.

In the current approach, two subproblems per grid level are created by splitting the residual error ($\mathcal{F}_k$ in (29)) into ''oscillatory'' and ''smooth'' components. The ''oscillatory'' component is smoothed again on the current grid level, while the ''smooth'' error component is restricted to coarser grids. Like the standard V-cycle approach, the process is recursive, and the smoothing and filtering operations are repeated on each coarser grid. The total coarse-grid correction at each grid level is found by summing the corrections generated by the solutions of the ''oscillatory'' and ''smooth'' subproblems. The filtered multigrid analogue of algorithm (29) is as follows:

Step 1. Restrict solution values from levels $k-1$ to k,

$$U_k = \tilde{r}[U_{k-1}]$$

Step 2. Smooth to determine correction ($\mathcal{F}^s_{k-1}$ is the collected ''smooth'' error at the previous grid level),

$$\text{solve: } N_k\, \delta U_k = \tilde{r}[\mathcal{F}^s_{k-1}]$$

Step 3. Calculate residual error at level k,

$$\mathcal{F}_k = \tilde{r}[\mathcal{F}^s_{k-1}] - \frac{1}{\varepsilon}(\mathcal{R}(U_k + \varepsilon\, \delta U_k) - \mathcal{R}(U_k)) \tag{30}$$

Step 4. Filter $\mathcal{F}_k$ into ''oscillatory'' and ''smooth'' subproblems $\mathcal{F}^o_k$ and $\mathcal{F}^s_k$

Step 5. Smooth $\mathcal{F}^o_k$ again on the current level:

$$\text{solve: } N_k\, \delta U^o_k = \mathcal{F}^o_k$$

Step 6. Restrict ''smooth'' residual error $\mathcal{F}^s_k$ to next coarsest level $k+1$ and continue cycling process

Step 7. Sum corrections at grid level k:

$$\delta U_k = \delta U_k + \delta U^o_k + \tilde{p}[\delta U_{k+1}].$$

In addition to the V-cycle workload, the filtered approach requires one additional implicit smoothing step per grid level. A suitable filtering operation must also be performed to split the residual error. The key to the performance of the method is the behavior of the filter, which should separate out high-frequency modes but not damp the underlying low frequency components. In Tuminaro's approach [32], the ''smooth'' component of the error on each grid level is obtained by first restricting the residual error to the next coarsest mesh and then bilinearly interpolating the restricted values to the finer grid:

$$\mathcal{F}^s_k = \tfrac{1}{4}\tilde{p}[\tilde{r}[\mathcal{F}_k]]. \tag{31}$$

The ''oscillatory'' component is obtained by subtracting the ''smooth'' component from the original residual error:

$$\mathcal{F}^o_k = \mathcal{F}_k - \mathcal{F}^s_k \tag{32}$$

In general, the filtered multigrid approach with Tuminaro's frequency decomposition converges better than the simple V-cycle strategy. However, as illustrated later, it is also subject to a degradation in convergence when highly clustered meshes are used. Quite by accident, it was discovered that the performance of the filtered multigrid scheme could be substantially improved by weighting the frequency decomposition so as to overrelax the smooth component of the residual error:

$$\mathcal{F}^s_k = \frac{\omega}{4}\tilde{p}[\tilde{r}[\mathcal{F}_k]] \tag{33}$$

$$\mathcal{F}^o_k = \frac{1}{2-\omega}(\mathcal{F}_k - \mathcal{F}^s_k), \tag{34}$$

$$1 \le \omega < 2$$

The net effect of this procedure is to ''overcorrect'' lower frequency error modes that are either not damped by the smoother or are approximated poorly on the coarser grids. In a recent article, Brandt and Yavneh [34] discuss the use of residual overrelaxation to alleviate multigrid convergence degradation associated with anisotropic advection–diffusion operators. Among other things, they conclude that extra smoothing passes are needed to prevent the overcorrection from adversely affecting error components that are resolved and attenuated effectively. Such extra smoothing is provided naturally by Step 5 of (30). Note that the Tuminaro decomposition is a special case ($\omega = 1$) of the ''overrelaxed'' decomposition presented above. Except for the $\omega = 1$ case, $\mathcal{F}^s_k$ and $\mathcal{F}^o_k$ do not sum to produce the original residual error. Both ''smooth'' and ''oscillatory'' error components are always present on the finer grid after the decomposition and are damped accordingly.

Both of the multilevel algorithms are implemented in the context of a ''full multigrid'' (FMG) initialization strategy. The FMG procedure starts on the coarsest mesh, where the basic diagonal implicit scheme is used to solve the problem to a specified tolerance. The solution is then interpolated to the next finest mesh and is improved by the application of a number of two-level multigrid cycles. The process is repeated, each time increasing the number of grid levels used in the multigrid cycling. After the interpolation to the finest level, the procedure reverts to the standard cycling strategy. The use of the FMG procedure has proved essential for the rapid computation of the blunt-body flowfields, as high-frequency errors due to the movement and formation of the normal shock are rapidly expelled by iterations on the coarser meshes.

6. RESULTS

Three test cases involving hypersonic flow over a one meter radius 2D cylinder are considered as means of evaluating the performance of the implicit reacting-gas flow solver. Two 65×113 grids, one uniformly spaced in the radial direction for Euler calculations and one with clustering to the cylinder surface ($\Delta r_{\min} = 8 \times 10^{-5}\ m$) for Navier–Stokes calculations, are considered (see Fig. 1). All computations are initialized by fixing the freestream fluid properties in the interior of the domain to their postnormal shock (perfect gas assumption) values. With this approach, the shock forms at the outer boundary and moves toward the cylinder until a steady position is reached.

To indicate convergence, the $\mathscr{L}_2$ norm of the discretized fine-grid residual is plotted versus Cray Y-MP single-processor CPU time. The residual $\mathscr{L}_2$ is normalized with respect to the maximum of its initial value and the value obtained after the coarse-grid steps of the FMG procedure. In this way, the level of fine-grid error reduction corresponding to the FMG initialization can be directly ascertained. The time for the FMG initialization is indicated by a shift in the fine-grid residual curves. At each coarse-grid stage of the FMG procedure, a two order-of-magnitude residual reduction is required before proceeding to the next level. The number of fine-grid cycles required for convergence is also listed on the convergence plots. These do not represent an equal amount of work—one four-level V-cycle is roughly 2.15 times as expensive as a single-grid iteration, and one filtered multigrid cycle is about 1.18 times as expensive as a V-cycle.

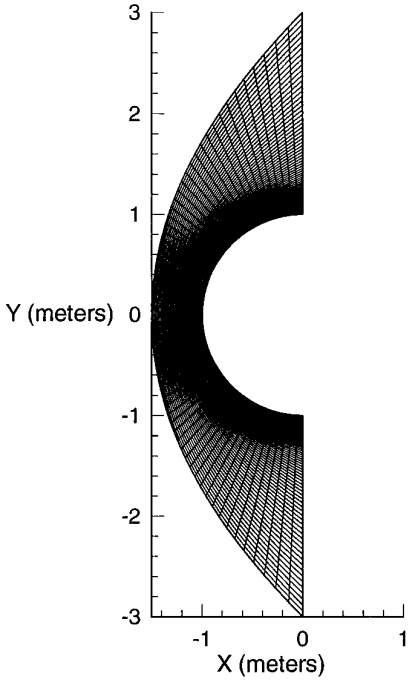

FIG. 1. Grid for Navier–Stokes computations (65×113).

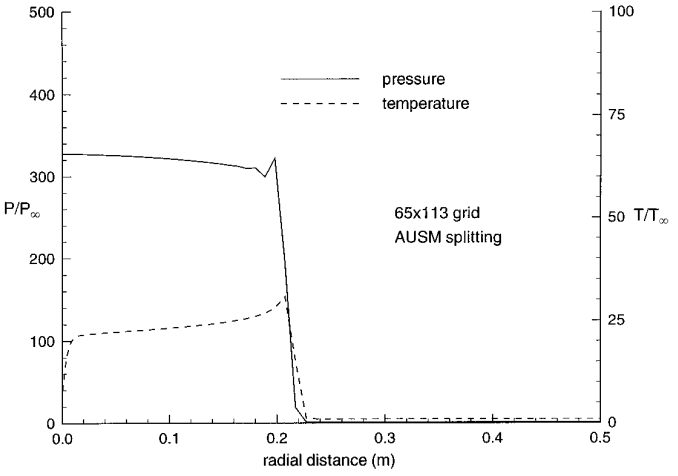

FIG. 2. Stagnation line pressure and temperature distributions (5-species Navier–Stokes; 5 km/s, 60 km).

6.1. *Five-Species Viscous Calculations* (5 km/s, 60 km)

Conditions for the first test case correspond to a 5 km/s flight at an altitude of 60 km. Freestream mass fractions are given as follows: $C_{N_2} = 0.78$, $C_{O_2} = 0.22$, $C_O = C_N = C_{NO} = 1 \times 10^{-5}$. A fixed wall temperature of 1273°K is enforced, and the wall is assumed to be noncatalytic. For all simulations, the nondimensional time step Δt is allowed to increase from an initial value of 0.01 to a final value of 0.0667 in proportion to a decrease in the normalized residual. All calculations (including the single-grid one) are initialized using the FMG procedure to provide a fair comparison.

Stagnation-line pressure and temperature distributions provided by the AUSM discretization are illustrated in Fig. 2. The AUSM approach, even in its first-order implementation, converges to a nonmonotone pressure solution. Other features of note include the temperature relaxation behind the strong normal shock wave and the rapid decrease in temperature within the thin thermal boundary layer. Species mass fraction distributions along the stagnation line are shown in Fig. 3. At these conditions, the dissociation of N_2 is incomplete, with some recombination occurring within the boundary layer as a result of the rapid temperature drop. Oxygen dissociation proceeds almost to completion. Comparisons of the AUSM stagnation-line results with those of other discretizations may be found in Ref. [35].

The effects of the choice of cycling strategy on the convergence of the 5 km/s Navier-Stokes simulation are shown in Fig. 4. For this particular comparison, the direct inversion (DI) approach for the source Jacobian is utilized exclusively. As evidenced, the convergence rates for all approaches, multilevel or single-grid, start to degrade after about a three order-of-magnitude reduction in the residual. This behavior is probably the result of inadequate damping or expulsion of high frequency errors, both within the well-resolved viscous layer and in the

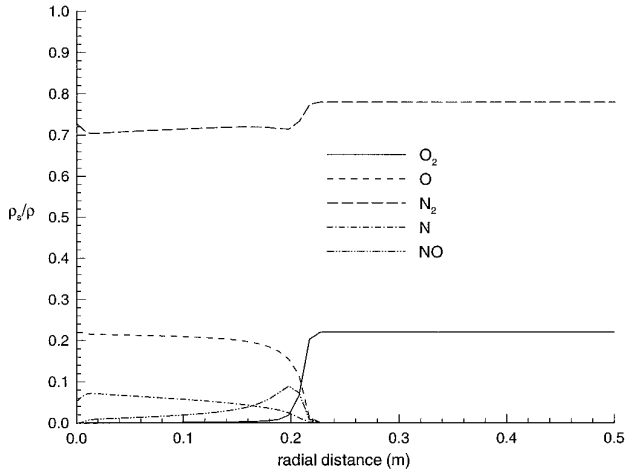

FIG. 3. Stagnation line mass fractions (5-species Navier–Stokes; 5 km/s, 60 km).

vicinity of the normal shock wave. In regions of heavy mesh clustering, the inviscid and viscous spectral radii within (16) and (17) remain large, acting to decrease the effective time step and thus to underrelax the obtained corrections. The nonmonotone behavior of the AUSM scheme near the normal shock wave may also influence the convergence behavior by providing a source for high-frequency errors that remain undamped throughout the iteration. The V-cycle convergence rate suffers the worst from these factors, leveling out at the single grid rate after about a 3.5 order-of-magnitude residual reduction. The filtered multigrid method using Tuminaro's decomposition ($\omega = 1$) performs significantly better, converging approximately 2.5 times faster than the single-grid iteration. Further improvements in the convergence rate are evidenced as the overrelaxation parameter ω is increased, with the optimal performance occurring around $\omega = 1.5$. Again, this improvement appears to result from the combination of additional high-frequency smoothing steps with an ''overcorrection'' procedure for poorly resolved error components. Above the $\omega = 1.5$ point, the implicit smoothing steps are unable to effectively attenuate the high-frequency errors resulting from the overrelaxation, and oscillatory behavior in the residual norm is observed. This process is reminiscent of SOR acceleration of a Gauss–Seidel procedure in that an ''optimal'' $\omega > 1$ usually exists for a particular problem, but the actual value must be determined by trial and error.

Figure 5 illustrates the effect of the source Jacobian approximations on the convergence of the filtered multigrid method. An overrelaxation factor of $\omega = 1.5$, corresponding to the ''optimal'' value for the DI approach, is utilized for all simulations. As shown, the columnwise approximate factorization (AF) procedure converges slightly slower than the direct inversion technique in terms of number of cycles. For this case, the per-iteration CPU improvement provided by the approximate factorization procedure is minimal, an effect due primarily to the small size ($m = 5$) of the source Jacobian and to the fact that the direct factorization is performed only *once* per grid level and is then stored. The Eberhardt–Imlay (EI) procedure converges much more slowly but in a more monotone fashion. This behavior is consistent with the nature of the approximation, which will always maintain diagonal dominance and will always underrelax the effect of the chemical source terms. It is possible that the poor relative performance of this method may again be related to the small size of the source Jacobian. For larger reaction sets, the $\mathbb{O}(m)$ cost of the EI method may somewhat offset its larger factorization error. It is also possible that optimal filtered multigrid performance for the EI and AF procedures occurs at values of ω different from that of the DI approach.

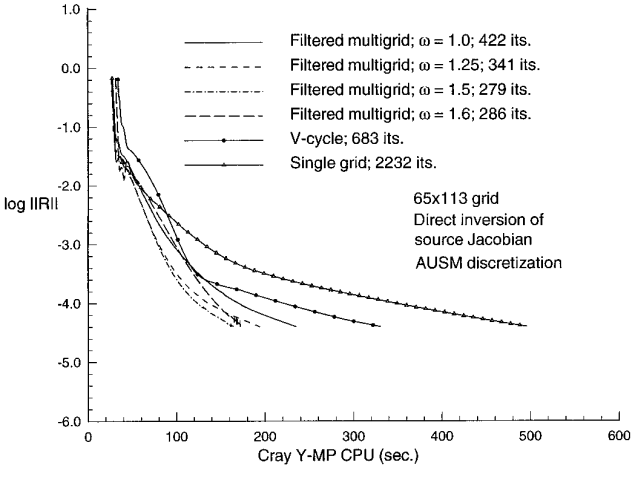

FIG. 4. Effect of cycling strategy on convergence (5-species Navier–Stokes; 5 km/s, 60 km),

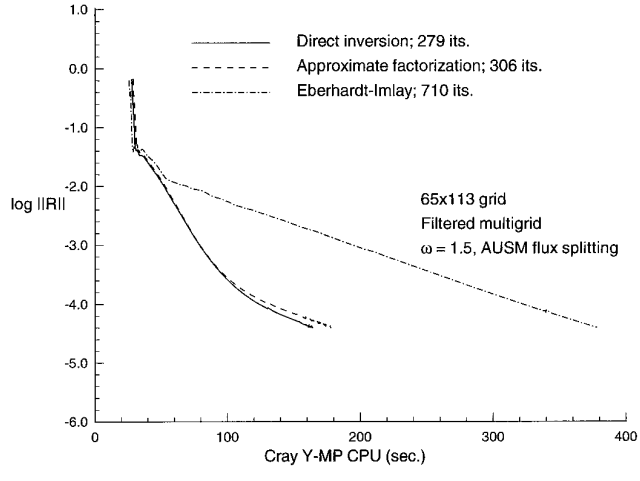

FIG. 5. Effect of source Jacobian approximations on convergence (5-species Navier–Stokes; 5 km/s, 60 km).

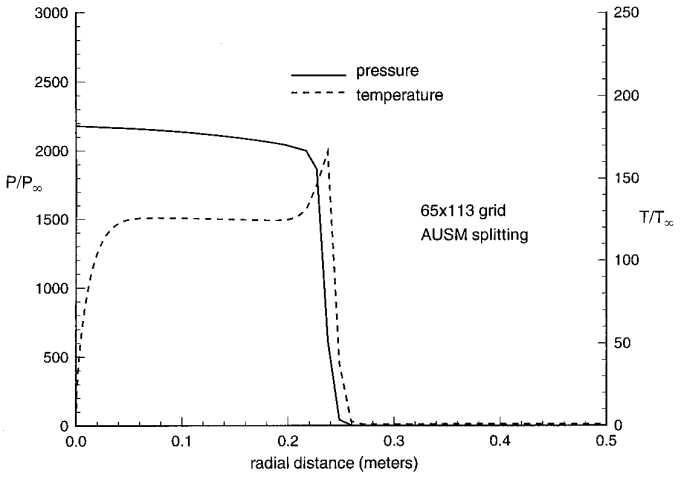

FIG. 6. Stagnation line pressure and temperature distributions (5-species Navier–Stokes; 11.37 km/s, 76 km).

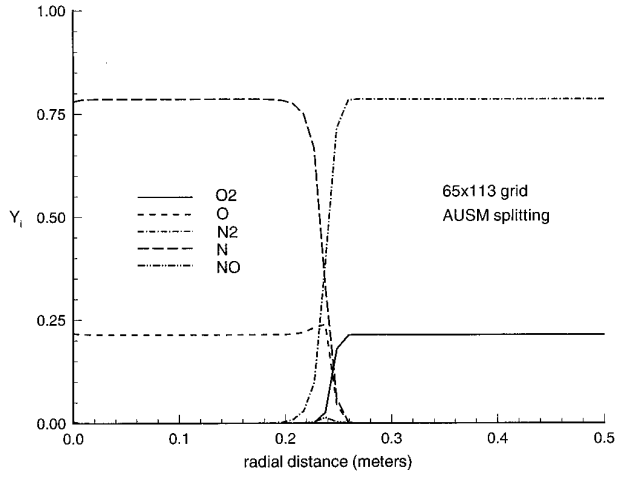

FIG. 7. Stagnation line mole fractions (5-species Navier–Stokes; 11.37 km/s, 76 km).

6.2. *Five-Species Viscous Calculation* (11.37 km/s, 76.42 km)

The second test case considered in this investigation involves 11.37 km/s flow over the same cylinder at an altitude of 76.42 km. These conditions correspond directly to those experienced by the Project Fire II re-entry vehicle 1634 s into its launch [36]. For this case, the wall temperature is set to 615°K, a noncatalytic assumption is again enforced, and the AUSM discretization is employed. The assumption of thermal equilibrium is clearly invalid at the Fire II conditions; the speed of the vehicle is also high enough to result in significant ionization. For the purposes of this paper, the inexactness of the physical modeling is not seen to be a significant weakness; the extreme conditions of the Fire II reentry serve only to test the robustness and efficiency of the ideas presented herein.

Computed stagnation line pressure, temperature, and mole fraction distributions for this case are illustrated in Figs. 6 and 7. The extreme increase in temperature through the normal shock (Fig. 6) results in the complete dissociation of O_2 and N_2. Only a slight amount of NO is present, as most quickly dissociates into its atomic components. Some recombination of N into N_2 occurs within the boundary layer (Fig. 7). For this case, the source Jacobian inversion procedure is performed within both sweeps of the LU-SGS algorithm. With this modification, the nondimensional time step could be increased to a final value of 0.2.

Convergence histories for the 11.37 km/s simulation are shown in Fig. 8. Only the results from the filtered multigrid method are shown, as the V-cycle approach was found to be unstable at the time step considered. The Eberhardt–Imlay source Jacobian approximation also failed at the test conditions. Of particular note in the figure is the large amount of error reduction (∼1.25 orders-of-magnitude) corresponding directly to the FMG initialization. It seems that most of the rapid changes due to chemical reaction take place within the coarse-grid cycling process. From the end of the initialization, the same general trends found in the 5 km/s simulation are evidenced. Operating at its optimal overrelaxation factor ($\omega = 1.5$), the filtered multigrid approach using the DI procedure again provides a threefold improvement in time to convergence, relative to the single-grid iteration. In contrast to the 5 km/s simulation, the asymptotic multigrid convergence rate does not degrade significantly, a possible consequence of the nearly monotone capturing of the normal shock wave for these conditions (Fig. 6).

For this case, the columnwise AF procedure converges in fewer cycles than the DI approach. As before, however, the per-iteration CPU difference between the two methods is not substantive enough to result in a significant improvement in

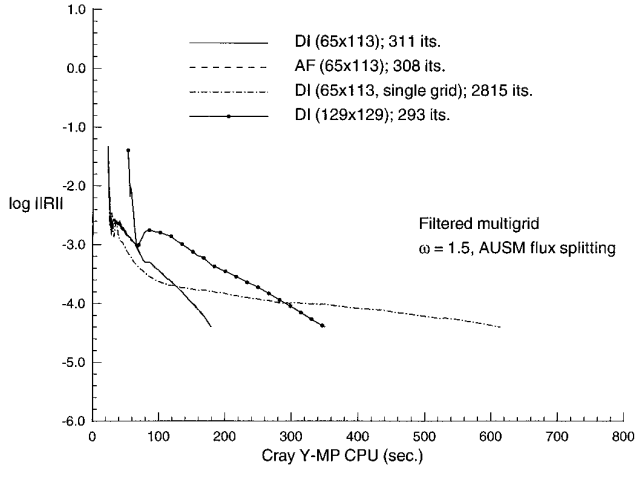

FIG. 8. Effect of source Jacobian approximations and grid refinement on convergence (5-species Navier–Stokes; 11.37 km/s, 76 km).

overall time to convergence. Also shown in the figure is a convergence history corresponding to a calculation on a 129 × 129 grid using the DI technique. The algorithm actually converges in fewer cycles on the finer grid than on the coarser grid, an effect possibly due to a more adequate resolution of the shock structure.

6.4. *Eleven-Species Inviscid Calculation* (11.37 km/s, 84.6 km)

The final test case considered in this study involves 11.37 km/s inviscid flow over the cylinder at an altitude of 84.6 km. These conditions correspond to the 1631 s point in the Fire II trajectory [36]. For this simulation, the 11-species reacting-air model with ionization is utilized. At this altitude, viscous and noncontinuum effects are important, and the translational, rotational, vibrational, and electron–electronic internal energy modes do not equilibrate at a common temperature, except very near the cylinder surface [37]. The assumption of inviscid flow characterized by one temperature is therefore quite inadequate, but again, the focus at this point is more toward examining the numerical performance of the developed algorithm for a large, stiff equation set than toward the physical modeling of reentry flows.

Figure 9 shows the computed temperature and pressure along the stagnation line. The effects of viscous layer displacement and a higher free-stream density notwithstanding, it is apparent from a comparison with Fig. 6 that the inclusion of ionization promotes a smaller shock-standoff distance and a lower temperature in the shock layer. The kinetic energy of the free-stream flow is absorbed in electronic and vibrational excitation as well as in the formation of neutral and ionic reaction products (Fig. 10).

Filtered multigrid convergence histories using the Van Leer (Hänel) splitting are shown in Fig. 11. Due perhaps to the severe thermal gradients in the vicinity of the shock, the less-dissipative AUSM discretization failed to give acceptable multigrid convergence for this problem. For this case, source term inversions are performed within both sweeps of the LU-SGS procedure, allowing a final nondimensional time step of 0.1 to be achieved. Compared with the direct approach for the source Jacobian inversion, the columnwise AF procedure provides only a 7.4% improvement in per-cycle efficiency, as the size of source Jacobian is still relatively small. However, the AF approach converges in fewer iterations (132 as compared with 154 for the direct method), leading to an overall improvement in time to convergence of over 18%. This behavior, consistent with that evidenced in the second test case, may indicate a tendency of the approximate procedure to underrelax the strong nonlinearities due to

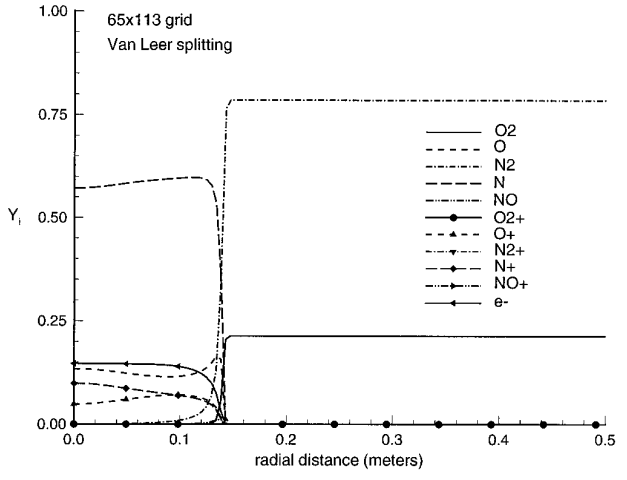

FIG. 10. Stagnation line mole fractions (11-species Euler; 11.37 km/s, 85 km).

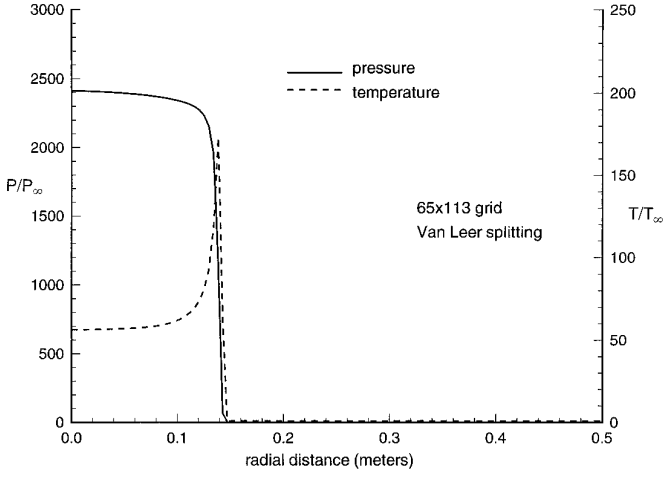

FIG. 9. Stagnation line pressure and temperature distributions (11-species Euler; 11.37 km/s, 85 km).

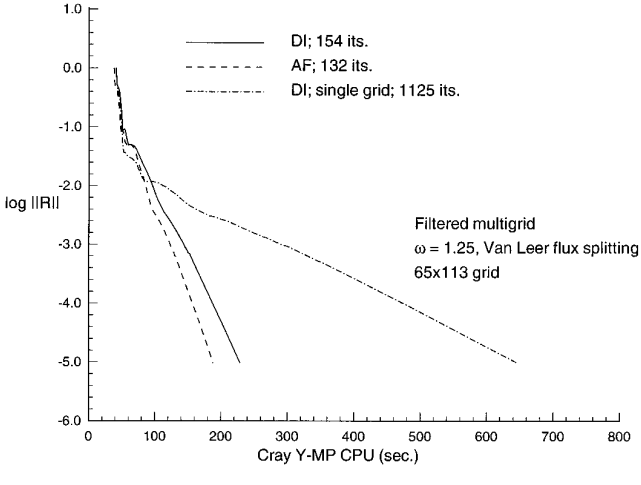

FIG. 11. Effect of source Jacobian approximations on convergence (11-species Euler; 11.37 km/s, 85 km).

chemical reaction so that a smoother path of convergence is achieved. The Eberhardt–Imlay ''diagonalization'' procedure again failed at the test conditions. Once again, the filtered multigrid method, operating at its optimal ω of 1.25, provides around a threefold improvement in overall run time compared with the single-grid calculation.

7. CONCLUSIONS

This work has outlined several techniques designed to improve the numerical efficiency of LU-SGS-based algorithms for computing chemically reacting, viscous flowfields. Current results indicate that the use of a columnwise approximate factorization procedure for the source vector Jacobian matrix can provide convergence behavior on a par with that using a direct factorization procedure, at least for the reaction sets and flight conditions considered in this study. For the five-species set, the use of *any* approximate procedure for the source Jacobian inversion appears a bit superfluous, but for the larger eleven-species reaction set, the approximate factorization procedure provides substantial benefit, converging in fewer iterations than the direct procedure while exhibiting a 7% improvement in per-iteration efficiency. The Eberhardt–Imlay ''diagonalization'' of the source Jacobian was found to be inferior to the other procedures, both in terms of overall rate of convergence and in terms of robustness.

Two multilevel convergence acceleration methods, a simple V-cycle and a parallelizable filtered approach, have been implemented and tested for inviscid and viscous reacting air simulations at hypersonic speeds. The filtered multigrid approach, enhanced by an overrelaxation procedure, provides significant improvement, normally converging three or more times faster than the single-grid approach. This method provides both additional high-frequency smoothing in the vicinity of strong shock waves and a mechanism by which lower-frequency error components can be resolved and attenuated effectively. A more theoretical study is needed to quantify the behavior of the overrelaxation and filtering processes, but the preliminary results are certainly encouraging.

The V-cycle strategy was found to be much less robust than the filtered approach for the extreme conditions encountered in this study. It is possible that the performance of the V-cycle could be improved through the use of additional smoothing passes on each grid level. One point of concern is the low per-iteration efficiency of the multigrid methods: a single filtered multigrid cycle is over 2.5 times as expensive as a single fine-grid iteration of the LU-SGS method. It should be possible to increase the efficiency of the cycling procedures through more careful coding and the use of compiler optimization. It is also possible that further improvements in the convergence rates of both multilevel approaches could be achieved through the use of upwind-biased grid-transfer operators and, possibly, semi-coarsening strategies.

ACKNOWLEDGMENTS

Part of this work was accomplished while the author was in residence at the Institute for Computational Mechanics in Propulsion (ICOMP) at NASA Lewis Research Center. This work was supported by the NASA Center of Research Excellence Program (NAGW-2924). Cray Y-MP computing time was provided by grants from the North Carolina Supercomputing Center and from NASA Lewis Research Center.

REFERENCES

1. G. V. Candler and R. W. MacCormack, AIAA Paper 88-0511 (unpublished).
2. P. A. Gnoffo, AIAA Paper 89-1972-CP (unpublished).
3. G. Palmer, AIAA Paper 90-1702 (unpublished).
4. S. Eberhardt and S. Imlay, AIAA Paper 90–1577 (unpublished).
5. J. H. Lee, ''Basic Governing Equations for the Flight Regimes of Aeroassisted Orbital Transfer Vehicles, in *Thermal Design of Aeroassisted Orbital Transfer Vehicles*, edited by H. F. Nelson, Progress in Astronautics and Aeronautics, Vol. 96 (AIAA, New York, 1985), p. 3.
6. P. A. Gnoffo, R. N. Gupta, and J. L. Shinn, NASA TP-2867, 1989 (unpublished).
7. C. Park, AIAA Paper 85-0247 (unpublished).
8. F. G. Blottner, M. Johnson, and M. Ellis, *Report No.* SC-RR-70-754, Sandia Laboratories, Albuquerque, NM, 1971 (unpublished).
9. C. R. Wilke, *J. Chem. Phys.* **18,** 517 (1950).
10. C. Park, AIAA Paper 87-1574 (unpublished).
11. Y. Liu and M. Vinokur, NASA CR-177489, 1988 (unpublished).
12. M. S. Liou and C. J. Steffen, NASA TM-104404, 1991 (unpublished).
13. L. Bergamini and P. Cinnella, AIAA Paper 93-0876 (unpublished).
14. D. Hänel, R. Schwane, and G. Seider, AIAA Paper 87-1105 (unpublished).
15. J. R. Edwards, *AIAA J.* **32,** 2120 (1994).
16. S. Yoon and A. Jameson, *AIAA J.* **26,** 1025 (1988).
17. S. Yoon and D. Kwak, AIAA Paper 91-1555-CP (unpublished).
18. H. Riegar and A. Jameson, AIAA Paper 88-0619 (unpublished).
19. B. Hassan, G. V. Candler, and D. R. Olynick, AIAA Paper 92-2877 (unpublished).
20. J. S. Shuen and S. Yoon, *AIAA J.* **27,** 1752 (1989).
21. T. L. Tysinger and D. A. Caughey, AIAA Paper 91-0242 (unpublished).
22. C. Park and S. Yoon, *AIAA J.* **30,** 999 (1992).
23. G. V. Candler and D. R. Olynick, ''Hypersonic Flow Simulations Using a Diagonal Implicit Method,'' in *Proceedings, 10th International Conference on Computing Methods in Applied Sciences and Engineering, Paris, France*, edited by R. Glowinski (Nova Science, New York, 1991), p. 29.
24. W. W. Hager, *SIAM Rev*. **31,** 221 (1989).
25. A. Brandt, *Math. Comput.* **31,** 333 (1977).
26. L. Martinelli, A. Jameson, and F. Grasso, AIAA Paper 86-0208 (unpublished).
27. V. N. Vatsa, and B. W. Wedan, AIAA Paper 89-1891 (unpublished).
28. E. Turkel, R. C. Swanson, V. N. Vatsa, and J. A. White, AIAA Paper 91-1572 (unpublished).
29. R. Radespiel and R.C. Swanson, NASA CR-189579 (ICASE Report 91-89, 1991 (unpublished).
30. F. Grasso and M. Marini, AIAA Paper 93-0771 (unpublished).
31. J. F. Slomski, J. D. Anderson, and J. J. Gorski, AIAA Paper 90-1575 (unpublished).
32. T. Chan and R. Tuminaro, ''Design and Implementation of Parallel

Multigrid Algorithms,'' in *Proceedings, Third Copper Mountain Conference on Multigrid Methods*, New York, 1987, edited by S. McCormick and M. Dekker, p. 101.

33. R. Tuminaro, *SIAM J. Sci. Statist. Comput.* **13,** 88 (1992).
34. A. Brandt and I. Yavneh, *SIAM J. of Sci. Statist. Comput.* **14,** 607 (1993).
35. J. R. Edwards, AIAA Paper 94-0762 (unpublished).
36. E. S. Cornette, NASA TM X-1305, 1966 (unpublished).
37. D. R. Olynick, Ph.D thesis, Department of Mechanical and Aerospace Engineering, North Carolina State University, 1992 (unpublished).