An Inviscid Decoupled Method for the Roe FDS Scheme in the Reacting Gas Path of FUN3D

Kyle B. Thompson* and Peter A. Gnoffo†

NASA Langley Research Center, Hampton, VA, 23681-2199

An approach is described to decouple the species continuity equations from the mixture continuity, momentum, and total energy equations for the Roe flux difference splitting scheme. This decoupling simplifies the implicit system, so that the flow solver can be made significantly more efficient, with very little penalty on overall scheme robustness. Most importantly, the computational cost of the point implicit relaxation is shown to scale linearly with the number of species for the decoupled system, whereas the fully coupled approach scales quadratically. Also, the decoupled method significantly reduces the cost in wall time and memory in comparison to the fully coupled approach. This work lays the foundation for development of an efficient adjoint solution procedure for high speed reacting flow.

Nomenclature

$A,\ A_d,\ A_m$	Jacobian Matrices
a	Speed of sound, m/s
$\mathbf{b}$	Residual vector
c_s	Species s mass fraction
C	Decoupled scheme chemical source term Jacobian
D	Decomposed diagonal Jacobian matrix
$dv_1,\ dv_2,\ dv_3$	Eigenvector components
e_s	Internal energy of species s
E	Total energy per unit mass J/kg^3
F'_ρ	Decoupled scheme mixture mass flux
F'_{ρ_s}	Decoupled scheme mass flux of species s
$\mathbf{F},\ \mathbf{F}',\ \hat{\mathbf{F}}$	Flux vectors
ns	Number of species
N_{nodes}	Number of nodes
N_{nz}	Number of non-zero off-diagonal entries in Jacobian
N_{nb}	Number of neighbors around local node
O	Decomposed off-diagonal Jacobian matrix
p	Pressure, N/m^2
$\mathbf{R},\ \mathbf{L}$	Right and left eigenvectors
R_ρ	Decoupled scheme constraint
$\mathbf{S}$	Face normal vector, m^2
$\mathbf{U},\ \mathbf{U}',\ \hat{\mathbf{U}}$	Conservative variable vectors
$\overline{U}$	Normal velocity, m/s^2
$u,\ v,\ w$	Components of velocity, m/s
V	Cell volume, m^3
w	Roe scheme weighting factor
ω	Chemical source term scaling factor

*NASA Pathways Intern, Aerothermodynamics Branch, AIAA Student Member.

†Senior Research Engineer, Aerothermodynamics Branch, AIAA Fellow.

$\hat{\mathbf{V}}$	Decoupled variables vector
$\mathbf{W}$	Chemical source term vector
ρ	Mixture density, kg/m^3
ρ_s	Species s density, kg/m^3
$\lambda_1,\ \lambda_2,\ \lambda_3$	Acoustic and convective eigenvalues
$\lambda^-,\ \lambda^+$	Species flux effective eigenvalues
Λ	Diagonal eigenvalue matrix
Superscript	
$n, n+1$	Time level
$R,\ L$	Right and left state quantities
f	Face

I. Introduction

The usability of hypersonic solvers on complex geometries is often limited by the extreme problem size associated with high energy physics. The additional equations required in reacting gas simulations lead to large Jacobians that scale quadratically in size to the number of governing equations. This leads to a significant increase in the memory required to store the flux linearizations and the computational cost of the point solver. As reacting gas CFD solvers are used to solve increasingly more complex problems, this onerous quadratic scaling of computational cost and Jacobian size will ultimately surpass the current limits of hardware and time constraints on achieving a flow solution.[1]

The proposed method is based heavily upon the work of Candler et al.[2] In that work, it was shown that quadratic scaling between the cost of solving the implicit system and adding species mass equations can be reduced from quadratic to linear scaling by decoupling the species mass equations from the mixture mass, momentum, and energy equations and solving the two systems sequentially. In the aforementioned work, the scheme was derived for a modified form of the Steger-Warming flux vector splitting method,[3] whereas the work presented here is derived for the Roe flux difference splitting (FDS) scheme.[4]

A primary motivator for the presented work is to prepare for the implementation of an adjoint capability in a reacting gas solver that employs the Roe scheme. Developing an adjoint implementation for non-equilibrium flows is an extremely challenging problem, because deriving exact Jacobians for a reacting gas system is particularly difficult. Decoupling the system also simplifies the derivation of exact Jacobians. If the species mass equations are decoupled, the mixture mass, momentum, and energy flux Jacobians can be easily derived,[5] and a weighting scheme can be used to correct the non-uniqueness of the pressure linearization.[6] For the decoupled species fluxes, we derive an exact linearization in the presented work. Future work will include an adjoint-based error estimation capability that leverages these exact linearizations.

II. Background: Fully-Coupled Point Implicit Method

All work presented here is for the inviscid conservation equations, but it can be extended to include viscous terms. For an inviscid, multi-species mixture, the governing equations in vector form are:

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{F} = \mathbf{W} \tag{1}$$

or, in semi-discrete form,

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{1}{V}\sum_f (\mathbf{F}\cdot\mathbf{S})^f = \mathbf{W} \tag{2}$$

summing over all faces, f, in the domain, where V is the cell volume, $\mathbf{W}$ is the chemical source term vector, and $\mathbf{S}$ is the face outward normal vector. The vectors of conserved variables and fluxes are:

$$\mathbf{U} = \begin{pmatrix} \rho_1 \\ \vdots \\ \rho_{ns} \\ \rho u \\ \rho v \\ \rho w \\ \rho E \end{pmatrix}, \quad \mathbf{F} = \begin{pmatrix} \rho_1 \overline{U} \\ \vdots \\ \rho_{ns} \overline{U} \\ \rho u \overline{U} + p s_x \\ \rho u \overline{U} + p s_y \\ \rho u \overline{U} + p s_z \\ (\rho E + p) \overline{U} \end{pmatrix} \tag{3}$$

where $\overline{U}$ is the outward pointing normal velocity, and E is the total energy of the mixture per unit mass, defined as:

$$E = \sum c_s e_s + \frac{u^2 + v^2 + w^2}{2} \tag{4}$$

where e_s is the internal energy of species s. By using the Roe FDS scheme,

$$\begin{aligned} \mathbf{F}^{n+1} &\approx \mathbf{F}^n + \tfrac{\partial \mathbf{F}}{\partial \mathbf{U}} \delta \mathbf{U}^n \\ \mathbf{W}^{n+1} &\approx \mathbf{W}^n + \tfrac{\partial \mathbf{W}}{\partial \mathbf{U}} \delta \mathbf{U}^n \end{aligned} \tag{5}$$

where $\delta \mathbf{U}^n = \mathbf{U}^{n+1} - \mathbf{U}^n$. By using an implicit time integration, the implicit scheme becomes:

$$\frac{\delta \mathbf{U}^n}{\Delta t} + \frac{1}{V} \sum_f \left(\frac{\partial \mathbf{F}^f}{\partial \mathbf{U}^{\mathbf{L}}} \delta \mathbf{U}^L + \frac{\partial \mathbf{F}^f}{\partial \mathbf{U}^{\mathbf{R}}} \delta \mathbf{U}^R \right)^n \mathbf{S}^f - \frac{\partial \mathbf{W}}{\partial \mathbf{U}} \delta \mathbf{U}^n = -\frac{1}{V} \sum_f (\mathbf{F}^f \cdot \mathbf{S}^f)^n + \mathbf{W}^n \tag{6}$$

or, put more simply:

$$A \delta \mathbf{U}^n = \mathbf{b} \tag{7}$$

where A is the Jacobian matrix of the fully coupled system, and $\mathbf{b}$ is the residual vector. For a point implicit relaxation scheme, the Jacobian matrix can be split into its diagonal and off-diagonal elements, with the latter moved to the RHS:

$$A = O + D \tag{8}$$

Each matrix element is a square $(ns+4) \times (ns+4)$ matrix. One method of solving this system is a Red-Black Gauss-Seidel scheme,[7] where matrix coefficients with even indices are updated first and, subsequently, the coefficients with odd indices are updated. This red-black ordering enables better vectorization in solving the linear system. The computational work for the Gauss-Seidel scheme is dominated by matrix-vector multiplications of elements of O with $\delta \mathbf{U}$, which are $O(N^2)$ operations, where $N = ns + 4$. In the next section, it is shown that decoupling the system reduces these matrix-vector multiplications to $O(N^2 + M)$ operations, where $N = ns$ and $M = ns$.

III. Background: Decoupled Point Implicit Method

If the species mass equations are replaced by a single mixture mass equation, the mixture equations can be separated from the species mass equations and the conserved variables become

$$\mathbf{U}' = \begin{pmatrix} \rho \\ \rho u \\ \rho v \\ \rho w \\ \rho E \end{pmatrix} \quad \hat{\mathbf{U}} = \begin{pmatrix} \rho_1 \\ \vdots \\ \rho_{ns} \end{pmatrix} \tag{9}$$

Solving the flux vector is performed in two sequential steps. The mixture fluxes are first solved as

$$\frac{\partial \mathbf{U}'}{\partial t} + \frac{1}{V} \sum_f (\mathbf{F}' \cdot \mathbf{S})^f = 0 \tag{10}$$

followed by the species fluxes as

$$\frac{\partial \hat{\mathbf{U}}}{\partial t} + \frac{1}{V}\sum_f (\hat{\mathbf{F}} \cdot \mathbf{S})^f = \hat{\mathbf{W}} \tag{11}$$

Point relaxation uses Red-Black Gauss-Seidel to update the conserved variables in $\mathbf{U}'$ and all associated auxiliary variables, such as temperature, pressure, speed of sound, etc. This is done by holding the thermo-chemical state constant, and will always result in the relaxation of a five-equation system. This does trade an implicit relationship between the mixture and species equations for an explicit one; thus, this decoupling can have an impact on the stability of the scheme, especially due to the non-linearity of the chemical source term.[8]

The solution of the species mass equations takes a different form. Based on the work of Candler et al.,[2] the decoupled variables can be rewritten in terms of mass fraction, as follows:

$$\delta \hat{\mathbf{U}}^n = \rho^{n+1}\hat{\mathbf{V}}^{n+1} - \rho^n \hat{\mathbf{V}}^n = \rho^{n+1}\delta\hat{\mathbf{V}}^n + \hat{\mathbf{V}}^n \delta\rho^n \tag{12}$$

where $\hat{\mathbf{V}} = (c_1, \ldots, c_{ns})^T$, and $c_s = \rho_s/\rho$ the mass fraction of species s. While the derivation of the species mass equations is different for the Roe FVS scheme from that of Steger-Warming proposed by Candler et al.,[2] the final result takes a similar form:

$$\hat{F}_{\rho_s} = c_s F'_\rho + (c_s^L - \tilde{c}_s)\rho^L\lambda^+ + (c_s^R - \tilde{c}_s)\rho^R\lambda^- \tag{13}$$

where F'_ρ is the total mass flux computed previously using all $\mathbf{U}'$ variables, and $\tilde{}$ denotes a Roe-averaged quantity. Likewise, linearizing the species mass fluxes with respect to the $\hat{\mathbf{V}}$ variables yields

$$\hat{\mathbf{F}}^{n+1} = \hat{\mathbf{F}}^n + \frac{\partial \hat{\mathbf{F}}}{\partial \hat{\mathbf{V}}^L}\delta\hat{\mathbf{V}}^L + \frac{\partial \hat{\mathbf{F}}}{\partial \hat{\mathbf{V}}^R}\delta\hat{\mathbf{V}}^R \tag{14}$$

$$\frac{\partial \hat{\mathbf{F}}}{\partial \hat{\mathbf{V}}^L} = wF_\rho + (1-w)\rho^L\lambda^+ - w\rho^R\lambda^- \tag{15}$$

$$\frac{\partial \hat{\mathbf{F}}}{\partial \hat{\mathbf{V}}^R} = (1-w)F_\rho + (w-1)\rho^L\lambda^+ + w\rho^R\lambda^- \tag{16}$$

A full derivation of Eqs. (13-16), along with the definition of w, is included in Appendix A. The chemical source term is linearized in the same manner as the fully coupled scheme; however, the updated $\mathbf{U}'$ variables are used to evaluate the Jacobian, and the chain rule is applied to linearize $\hat{\mathbf{W}}$ with respect to the species mass fractions:

$$\hat{\mathbf{W}}^{n+1} = \hat{\mathbf{W}}^n + \left.\frac{\partial \hat{\mathbf{W}}}{\partial \mathbf{U}}\right|_{\mathbf{U}'} \frac{\partial \mathbf{U}}{\partial \hat{\mathbf{V}}} \tag{17}$$

For simplicity of notation, we define

$$C = \left.\frac{\partial \hat{\mathbf{W}}}{\partial \mathbf{U}}\right|_{\mathbf{U}'} \frac{\partial \mathbf{U}}{\partial \hat{\mathbf{V}}} \tag{18}$$

The decoupled system to be solved becomes:

$$\begin{aligned} \rho^{n+1}\frac{\delta\hat{\mathbf{V}}^n}{\Delta t} + \frac{1}{V}\sum_f \left(\frac{\partial \hat{\mathbf{F}}^f}{\partial \hat{\mathbf{V}}^L}\delta\hat{\mathbf{V}}^L + \frac{\partial \hat{\mathbf{F}}^f}{\partial \hat{\mathbf{V}}^R}\delta\hat{\mathbf{V}}^R\right)^{n,n+1}\mathbf{S}^f - C^{n,n+1}\delta\mathbf{V}^n \\ = -\frac{1}{V}\sum_f (\hat{\mathbf{F}}^{n,n+1}\cdot\mathbf{S})^f + \mathbf{W}^{n,n+1} - \hat{\mathbf{V}}^n\frac{\delta\rho^n}{\Delta t} - R_\rho \end{aligned} \tag{19}$$

$$R_\rho = -\frac{1}{V}\sum_f\sum_s (\hat{F}^{n,n+1}_{\rho_s}\cdot\mathbf{S}) \tag{20}$$

where R_ρ is included to preserve the constraint that the mass fractions sum to unity, i.e., $\sum_s c_s = 1$, $\sum_s \delta c_s = 0$.

IV. Cost and Memory Savings of the Decoupled Implicit Problem

In decoupling the species equations, the most significant savings comes from the source term linearization being purely node-based.[9] Solving the mean flow equations is conducted in the same manner as the fully coupled system. All entries in the Jacobian A_m are linearizations of the mixture equation fluxes, which results in 5×5 matrices. All entries in the Jacobian A_d are linearizations of the species mass fluxes, which results in $ns \times ns$ matrices. Because there is no interdependence of species, except through the chemical source term, all contributions due to linearizing the convective flux are purely diagonal $ns \times ns$ matrices. Via Eq. (8), we decompose A_d into its diagonal and off-diagonal elements, resulting in the following linear system:

$$\begin{pmatrix} \square & & & \\ & \ddots & & \\ & & \square & \\ & & & \ddots \\ & & & & \square \end{pmatrix} \begin{pmatrix} \delta\hat{\mathbf{V}}_1 \\ \vdots \\ \delta\hat{\mathbf{V}}_i \\ \vdots \\ \delta\hat{\mathbf{V}}_{nodes} \end{pmatrix} = \begin{pmatrix} \hat{b}_1 \\ \vdots \\ \hat{b}_i \\ \vdots \\ \hat{b}_{nodes} \end{pmatrix} - \begin{pmatrix} (\sum_{j=1}^{N_{nb}} [\diagdown]\delta\hat{\mathbf{V}}_j)_1 \\ \vdots \\ (\sum_{j=1}^{N_{nb}} [\diagdown]\delta\hat{\mathbf{V}}_j)_i \\ \vdots \\ (\sum_{j=1}^{N_{nb}} [\diagdown]\delta\hat{\mathbf{V}}_j)_{nodes} \end{pmatrix} \tag{21}$$

where $\square$ represents a dense $ns \times ns$ matrix, $[\diagdown]$ represents a diagonal matrix, and $\delta\hat{\mathbf{V}}_j$ is the decoupled variable update on the node j that neighbors node i, where N_{nb} is the number of nodes neighboring node i. Thus, the non-zero entries in the off-diagonal matrix can be reduced from diagonal matrices to vectors. This results in significant savings in both computational cost and memory, as the only quadratic operation left in solving the implicit system is dealing with the diagonal entries in the Jacobian. Because the off-diagonal entries significantly outnumber the diagonal entries, we can expect nearly linear scaling in cost with the number of species. If compressed row storage[10] is used to only store non-zero off-diagonal entries, the relative memory savings in the limit of a large number of species for the Jacobian is given by

$$\begin{aligned} \textit{Relative Memory Cost} &= \frac{size(A_d)}{size(A)} \\ &= \lim_{ns\to\infty} \frac{(ns^2 + 5^2)(N_{nodes}) + (ns + 5^2)(N_{nz})}{(ns+4)^2(N_{nodes} + N_{nz})} \\ &= \frac{N_{nodes}}{N_{nodes} + N_{nz}} \end{aligned} \tag{22}$$

where N_{nodes} is the number of nodes, and N_{nz} is the number of non-zero off-diagonal entries stored using compressed row storage. For a structured grid, each node has six neighbors in 3D, i.e., $N_{nz} = 6N_{nodes}$; therefore, we can expect the Jacobian memory required to decrease by a factor of seven using this decoupled scheme. Interestingly, for a grid that is not purely hexahedral, $N_{nz} > 6N_{nodes}$; thus, this decoupled scheme provides higher relative memory savings on unstructured grids than structured grids when using compressed row storage.

V. Results

A. Cylinder

Both the proposed decoupled scheme and the traditional, fully coupled approach have been implemented using FUN3D.[11] Demonstrating the improved efficiency in cost and memory required to utilize the decoupled scheme, and that both the fully coupled and decoupled approaches converge to the same result, a grid convergence study was conducted on a simple cylinder geometry (radius 0.5 m). Due to the presence of strong shocks in blunt body flows, it was advantagous to generate structured-type grids to preserve grid alignment with the bow shock. A 50×50, 100×100, and 200×200 family of grids were adapted using the adaptation capability in FUN3D[12] to produce shock-aligned grids. These grids serve as a surrogate for conducting a grid convergence study, in that differences observed between the decoupled and fully coupled schemes decrease as the average mesh spacing decreases. These grids are unstructured, consisting totally of hexahedral elements with a single cell in the spanwise direction, and the 50×50 grid is shown in Figure 1. The cell elements of these grids were also subdivided into tetrahedral elements, and it was verified that there

are no issues with a true unstructured grid topology. The free stream conditions used were $V_\infty = 5000\ m/s$, $\rho_\infty = 0.001\ kg/m^3$, and $T_\infty = 200\ K$. Several chemical kinetics models were used, including a 5-species model with 5 reactions, an 11-species model with 22 reactions, and an 18-species model with 29 reactions. All cases were run in thermodynamic equilibrium, with a one-temperature model.

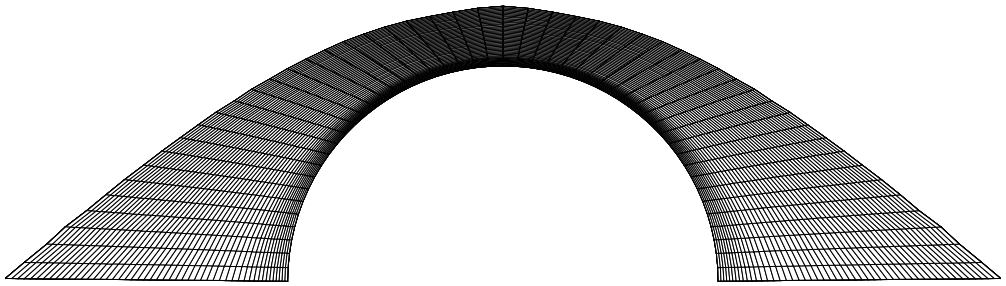

Figure 1: 50×50 cylinder grid.

1. *Verification of Implementation*

In order to be valid, the decoupled scheme must yield converged solutions that are nearly identical to those of the fully coupled system. To quantitatively assess this, we compare the predicted surface pressure, surface temperature, and the species composition on the stagnation line for both schemes. Figure 2 shows the predicted quantities on the 100×100 grid, for species mixture of N, N_2, O, O_2, and NO with five reactions. All results are indeed nearly identical, with temperature and pressure matching discretely to eight digits and the species mass fractions on the stagnation line matching to four digits. This difference was further reduced on the finest grid level of 200×200, suggesting that both schemes converge to the same solution with grid refinement.

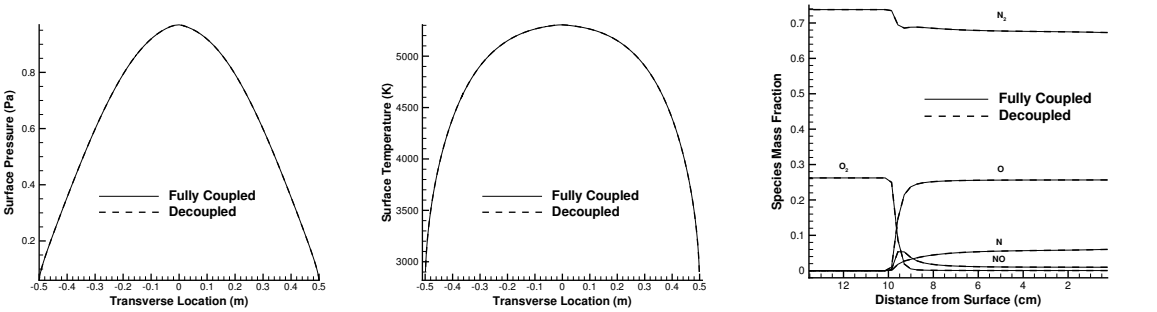

(a) Surface pressure (b) Surface temperature (c) Mass fractions on stagnation line

Figure 2: Cylinder predicted quantities.

2. *Memory Cost*

In order to determine the required memory of the decoupled scheme compared to the fully coupled scheme, a convergence study was conducted using Valgrind[13] to determine the memory actually allocated by FUN3D for an increasing number of species. Figure 3 shows that the relative memory cost converges asymptotically to ∼1/4, which is nearly twice the predicted value of 1/7. For the implementation of FUN3D, this is correct because the off-diagonal entries are reduced from double to single precision. Each structured grid node has

six neighboring nodes, with the exception of those at the boundary. Because each of these six neighboring nodes yields single precision, off-diagonal Jacobian elements,

$$N_{nz} = \frac{6N_{nodes}}{2} = 3N_{nodes} \tag{23}$$

Substituting Eq. (23) into Eq. (22), the relative memory cost is:

$$Relative\ Memory\ Cost = \frac{N_{nodes}}{N_{nodes} + N_{nz}} = \frac{N_{nodes}}{N_{nodes} + (3N_{nodes})} = \frac{1}{4} \tag{24}$$

thus, the relative memory saved by using the decoupled scheme correctly approaches a factor of 1/4.

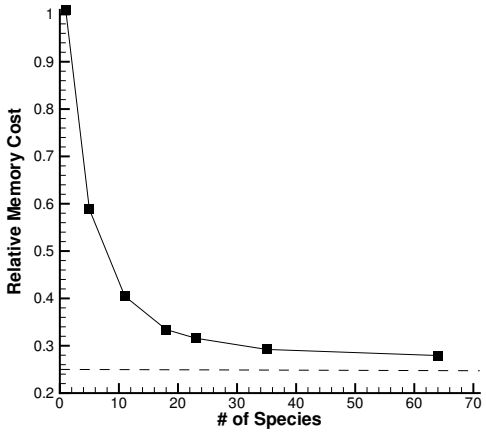

Figure 3: Memory required convergence study

3. Computational Cost

As stated before, the cost of solving the decoupled implicit system should scale approximately linearly with the number of species, whereas the fully coupled problem should scale quadratically; thus, the speedup of the implicit solve should be approximately linear when comparing the decoupled and fully coupled approaches. Figure 4 shows this to be true for the cylinder test case, and that the total speedup of the problem is less than that of just the linear solve. It is to be expected that the overall gains are not as large as those for the implicit solve, since there are many other factors that scale with the number of species, especially calculating the species source term and its linearization.

B. Sphere-Cone

To ensure that the decoupled scheme is robust and accurate at higher velocities, both the fully coupled and decoupled approaches were run on a sphere-cone geometry identical to that presented by Candler et. al.[2] (10 cm nose radius, 1.1 m length, 8^o cone angle). For this case, a simple 64×64 hexahedral grid was constructed, and freestream conditions were set as $V_\infty = 15000\ m/s$, $\rho_\infty = 0.001\ kg/m^3$, $T_\infty = 200\ K$. It was discovered that CFL limitations for the decoupled scheme were prohibitive, because of the stiffness of the chemical source term. In order to converge the scheme in a manner competitive with the fully coupled approach, it was necessary to scale the magnitude of the source term contribution to the flux balance by a value ω, such that $0 \leq \omega \leq 1$. To ensure that the decoupled and fully coupled approaches yielded the same result, a ramping scheme was implemented such that no scaling was performed on the source term when the solution was in a converged state.

1. Verification of Implementation

As with the cylinder test case, the surface pressure and temperature were used as metrics to determine that both the decoupled and fully coupled approaches give the same answer when converged to steady-state.

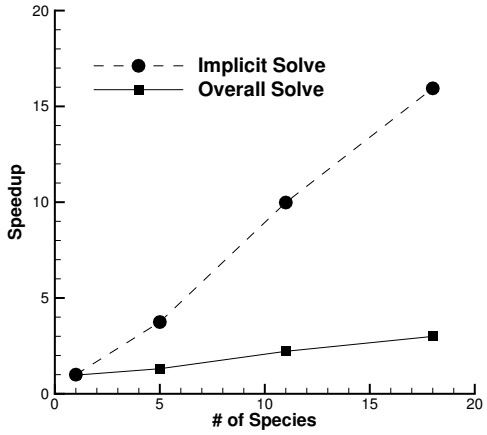

Figure 4: Relative speedup for the decoupled scheme vs. fully coupled scheme.

The species composition consisted of N, N_2, O, O_2, NO, N^+, N_2^+, O^+, O_2^+, NO^+, and electrons, with 22 possible reactions. Figure 5 shows that both methods again yield similar results, and the high stagnation temperature indicates that this is an inviscid, one-temperature simulation. This demonstrates that the decoupled approach is able to converge to the same solution as the fully coupled solution, in spite of the chemical reactions proceeding very rapidly due to a high stagnation temperature.

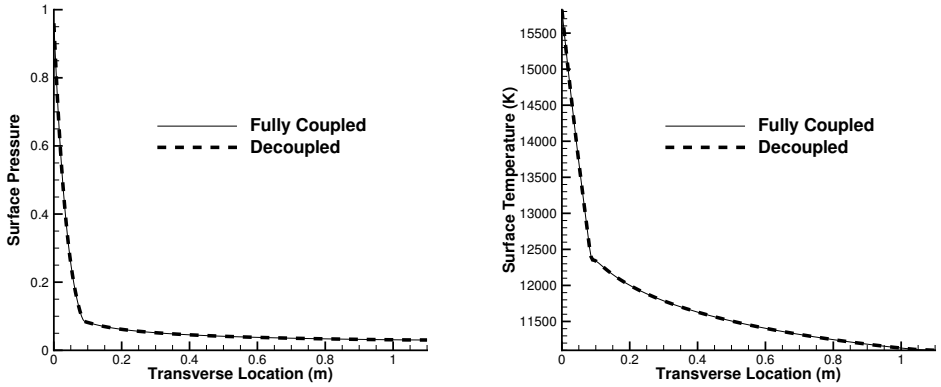

(a) Surface pressure

(b) Surface temperature

Figure 5: Sphere-cone predicted quantities.

2. Convergence Quality

The limits on the stability of the decoupled scheme derives from introducing explicitness in creating and destroying species. By scaling the magnitude of the chemical source term during the transient phase of the solve, this instability can be mitigated, and the convergence of decoupled scheme approaches that of the fully coupled scheme. Scaling of chemical source term was done identically between the decoupled and fully coupled scheme by ramping the factor ω from 0.001 to 1.0 over the first 500 timesteps. Figure 6 shows that the convergence of both schemes progresses nearly identically, with the decoupled scheme converging in significantly less computational time and, interestingly, fewer timesteps. This demonstrates that the decoupled scheme has significant potential to improve the effiency of high-velocity simulations, and that the stiffness of the source term can be overcome in the presence of large chemical reaction rates.

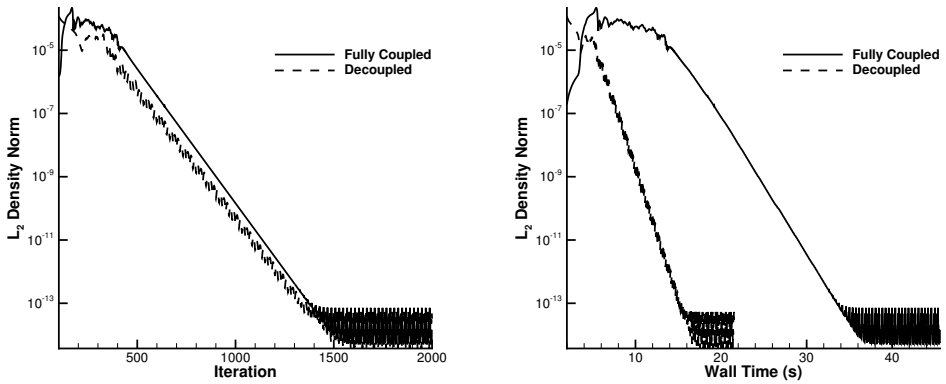

(a) Iterations to convergence (b) Computational time to convergence

Figure 6: Sphere-cone convergence details.

VI. Conclusion

The work presented here extends the derivation of a decoupled Steger-Warming scheme to that of the Roe FDS scheme. By decoupling the species mass fluxes, the computational cost and memory required in solving the implicit system is significantly less than that required by the traditionally, fully coupled implementation of the Roe FDS scheme. The motivation for this work was to aid in the development of an adjoint development for the reacting gas path of FUN3D. Because of the complexity of the high-energy physics and the large problem size associated with additional equations needed to preserve species mass, solving an adjoint in a fully coupled manner becomes prohibitively costly; therefore, improving the efficiency of the reacting gas flow solver path in FUN3D is extremely beneficial towards an adjoint formulation. More importantly, the decoupling of the species equations from the mixture equations presents significant opportunities to obtain accurate Jacobians for the Roe FDS scheme, which are extremely useful in solving the adjoint problem. The $\sim 1/3$ memory required and 200 percent speedup for the 11-species case, as well as the demostrated robustness at very high velocities, clearly illustrates significant advantages of the decoupled method for this inviscid implementation.

A. Decoupled Flux Derivation

For the Roe flux difference splitting scheme, the species mass fluxes are given by

$$F_{\rho_s} = \frac{\rho_s^L \overline{U}^L + \rho_s^R \overline{U}^R}{2} - \frac{\tilde{c}_s(\lambda_1 dv_1 + \lambda_2 dv_2) + \lambda_3 dv_{3_s}}{2} \tag{25}$$

$$dv_1 = \frac{p^R - p^L + \tilde{\rho}\tilde{a}(\overline{U}^R - \overline{U}^L)}{\tilde{a}^2} \tag{26}$$

$$dv_2 = \frac{p^R - p^L - \tilde{\rho}\tilde{a}(\overline{U}^R - \overline{U}^L)}{\tilde{a}^2} \tag{27}$$

$$dv_{3_s} = \frac{\tilde{a}^2(\rho_s^R - \rho_s^L) - \tilde{c}_s(p^R - p^L)}{\tilde{a}^2} \tag{28}$$

$$\lambda_1 = \mid \overline{\mathbf{U}} + \tilde{a} \mid, \quad \lambda_2 = \mid \overline{\mathbf{U}} - \tilde{a} \mid, \quad \lambda_3 = \mid \overline{\mathbf{U}} \mid \tag{29}$$

where the ˜ notation signifies a Roe-averaged quantity, given by:

$$\tilde{\mathbf{U}} = w\tilde{\mathbf{U}}^L + (1-w)\tilde{\mathbf{U}}^R \tag{30}$$

$$w = \frac{\tilde{\rho}}{\tilde{\rho} + \rho^R} \tag{31}$$

$$\tilde{\rho} = \sqrt{\rho^R \rho^L} \tag{32}$$

The species mass fluxes must sum to the total mass flux; thus, the total mixture mass flux is given as

$$F_\rho = \sum_s F_{\rho_s} = \frac{\rho^L \overline{U}^L + \rho^R \overline{U}^R}{2} - \frac{\tilde{c}_s(\lambda_1 dv_1 + \lambda_2 dv_2) + \lambda_3 dv_3}{2} \tag{33}$$

$$dv_3 = \frac{\tilde{a}^2(\rho^R - \rho^L) - (p^R - p^L)}{\tilde{a}^2} \tag{34}$$

Multiplying Eq. (33) by the Roe-averaged mass fraction and substituting it into Eq. (25) results in:

$$F_{\rho_s} = \tilde{c}_s F_\rho + \frac{(c_s^L - \tilde{c}_s)\rho^L(\overline{U}^L + |\,\tilde{U}\,|)}{2} + \frac{(c_s^R - \tilde{c}_s)\rho^R(\overline{U}^R - |\,\tilde{U}\,|)}{2} \tag{35}$$

It should be noted here that the Roe-averaged normal velocity, $\tilde{U}$, requires an entropy correction in the presence of strong shocks.[14] This correction has no dependence on the species mass fractions; therefore, it does not change the form of the Jacobian for this decoupled scheme. The notation can be further simplified by defining the normal velocities as follows:

$$\lambda^+ = \frac{\overline{U}^L + |\,\tilde{U}\,|}{2}, \quad \lambda^- = \frac{\overline{U}^R - |\,\tilde{U}\,|}{2} \tag{36}$$

Finally, substituting Eq. (36) into Eq. (35) yields the final result for calculating the species flux in the decoupled system:

$$F_{\rho_s} = \tilde{c}_s F_\rho + (c_s^L - \tilde{c}_s)\rho^L \lambda^+ + (c_s^R - \tilde{c}_s)\rho^R \lambda^- \tag{37}$$

Forming the convective contributions to the Jacobians is straightforward. Because the $\mathbf{U}'$ level variables are constant, only the left, right, and Roe-averaged state mass fractions vary. Differentiating Eq. (37) with respect to the mass fraction, c_s, the left and right state contributions are

$$\frac{\partial F_{\rho_s}}{\partial c_s^L} = wF_\rho + (1-w)\rho^L\lambda^+ - w\rho^R\lambda^- \tag{38}$$

$$\frac{\partial F_{\rho_s}}{\partial c_s^R} = (1-w)F_\rho + (w-1)\rho^L\lambda^+ + w\rho^R\lambda^- \tag{39}$$

Because there is no dependence between species in decoupled convective formulation, the Jacobian block elements are purely diagonal for the convective contributions, of the form

$$\begin{pmatrix} \frac{\partial F_{\rho_1}}{\partial c_1} & & 0 \\ & \ddots & \\ 0 & & \frac{\partial F_{\rho_{ns}}}{\partial c_{ns}} \end{pmatrix} \tag{40}$$

Acknowledgments

The authors would like to thank the Entry Systems Modeling project for their funding and support of this research. Also, the authors would like to recognize the FUN3D team at NASA Langley Research Center, for their support in integrating aspects of the compressible gas path into the reacting gas path of FUN3D.

References

[1]Fischer, P. F.: *Scaling Limits for PDE-Based Simulation (Invited)*, American Institute of Aeronautics and Astronautics. 2015/10/21 2015.

[2]Candler, G. V.; Subbareddy, P. K.; and Nompelis, I.: Decoupled Implicit Method for Aerothermodynamics and Reacting Flows. *AIAA Journal*, Vol. 51, no. 5, 2015/04/23 2013, pp. 1245–1254.

[3]MacCormack, R. W.; and Candler, G. V.: The solution of the Navier-Stokes equations using Gauss-Seidel line relaxation. *Computers and Fluids*, Vol. 17, no. 1, 1989, pp. 135–150.

[4]Roe, P.: Approximate Riemann solvers, parameter vectors, and difference schemes. *Journal of Computational Physics*, Vol. 43, no. 2, 1981, pp. 357 – 372.

[5]Nishikawa, H.: Implementing a Real-Gas Roe Solver into Implementing a Real-Gas Roe Solver into CFL3D. April 2004. URL http://www.cfdnotes.com.

[6]Shuen, J.-S.; Liou, M.-S.; and Leer, B. V.: Inviscid flux-splitting algorithms for real gases with non-equilibrium chemistry. *Journal of Computational Physics*, Vol. 90, no. 2, 1990, pp. 371–395.

[7]Saad, Y.: *Iterative Methods for Sparse Linear Systems*, Society for Industrial and Applied Mathematics. 2003, pp. 391–392.

[8]Park, C.: *Assessment of two-temperature kinetic model for ionizing air*, American Institute of Aeronautics and Astronautics. 2015/10/23 1987.

[9]Gnoffo, P. A.; Gupta, R. N.; and Shinn, J. L.: Conservation Equations and Physical Models for Hypersonic Air Flows in Thermal and Chemical Nonequilibrium. Tehnical Paper 2867, NASA, 1989.

[10]George, A.; and Liu, J. W.: *Computer Solution of Large Sparse Positive Definite*. Prentice Hall Professional Technical Reference, 1981.

[11]Anderson, W. K.; and Bonhaus, D. L.: "An Implicit Upwind Algorithm for Computing Turbulent Flows on Unstructured Grids". *Computers and Fluids*, Vol. 23, no. 1, 1994, pp. 1–22.

[12]Bartels, R.; Vatsa, V.; Carlson, J.-R.; and Mineck, R.: FUN3D Grid Refinement and Adaptation Studies for the Ares Launch Vehicle. *28th AIAA Applied Aerodynamics Conference*, AIAA, 2015/10/06 2010.

[13]ACM SIGPLAN 2007 Conference on Programing Language Design and Implementation: *How to Shadow Every Byte of Memory Used by a Program*, June 2007.

[14]Harten, A.: High resolution schemes for hyperbolic conservation laws. *Journal of Computational Physics*, Vol. 49, no. 3, 1983, pp. 357–393.