To the Graduate Council:

I am submitting herewith a thesis written by Nicholas G. Currier entitled “A Hybrid Method for Flows in Local Chemical Equilibrium and Nonequilibrium.” I have examined the final electronic copy of this thesis for form and content and recommend that it be accepted in partial fulfillment of the requirements for the degree of Master of Science, with a major in Computational Engineering.

We have read this thesis and recommend its acceptance:

Daniel G. Hyams, Ph.D.
Associate Professor,
Computational Engineering
Chairperson/Principal Adviser

W. Roger Briley, Ph.D.
Professor,
Computational Engineering

Robert S. Webster, Ph.D.
Associate Research Professor,
Computational Engineering

Accepted for the Graduate Council:

Stephanie Bellar, Ph.D.
Interim Dean of The Graduate School

(Original signatures are on file with official student records)

A Hybrid Method for Flows in Local Chemical Equilibrium and Nonequilibrium

A Thesis Presented for
The Master of Science Degree
The University of Tennessee at Chattanooga

Nicholas G. Currier

August 2010

© by Nicholas G. Currier, 2010
All rights reserved.

Dedication

I would like to dedicate this thesis to my parents, Frederick and Janice Currier, who instilled in me a love of learning and have always encouraged me to pursue my dreams.

Acknowledgements

Foremost, I would like to thank my soon-to-be wife, Maria, for her unending support. Her encouragement has helped me immensely in the pursuit of this degree.

A very sincere thanks goes to my adviser, Dr. Daniel Hyams, for his guidance and for the many long conversations we have had regarding this work. Without his help, this thesis would not have been possible. I would like to thank my other committee members, Dr. Roger Briley and Dr. Robert Webster, for their guidance both regarding this work and in general. Dr. Kidambi Sreenivas also deserves my appreciation for his help and advice on several issues that came to light during this process.

Abstract

The primary objective of this work is to develop a more efficient chemically active compressible Euler equation solver. Currently, a choice between the physical accuracy of a finite-rate solver or the computational efficiency of an equilibrium flow solver must be made. The number of species modeled continues to increase with available computational resources. A method of further leveraging the increase in computational power is desired.

The hybrid chemistry scheme proposed here attempts to maintain the accuracy of finite-rate schemes while retaining some of the cost savings associated with equilibrium chemistry solvers. The method given uses a full finite-rate flux in regions where chemistry is slow compared to the advection rate and an equilibrium chemistry scheme in regions where the chemistry outpaces the fluid transport. Control volume switching is based on a locally defined Damköhler number. This method could be extremely useful for full reaction path modeling or the tracking of a very large number of species. The cost of symmetric Gauss-Seidel iterations grows like the number of species plus four, quantity squared. Thus, eliminating the increased cost of solving for a large number of unknowns in regions where it is unjustified can be very useful.

Tenasi, a University of Tennessee SimCenter research code, is used as a base for the new solver. The hybrid method is implemented and tested with an explicit solution technique in one dimension. In combination with a five species air chemistry model, a high-temperature shock tube is used as a verification test case. Results are compared with those from pure equilibrium, full finite-rate, perfect gas Euler, and exact perfect

gas Riemann solvers. Timings are also given, suggesting the cost savings that would be possible should the hybrid method be extended using implicit algorithms.

Contents

List of Tables	x
List of Figures	xi
Nomenclature	xii
1 Introduction	1
2 Computational Formulation	4
2.1 Governing Equations	4
2.1.1 Finite-rate Chemistry Regime	5
2.1.2 Equilibrium Chemistry Regime	8
2.1.3 Hybrid Chemistry Regime	15
2.2 HLLC Flux	20
2.3 Spatial Discretization	22
2.4 Temporal Discretization	24
2.5 Explicit Solution Technique	25
2.6 Chemistry	26
2.6.1 Thermodynamic Data	26
2.6.2 Reaction Rates	29
3 Implementation Issues	31

4	Numerical Verification	34
4.1	Problem Setup	34
4.2	Refinement Study	35
4.3	Verification	39
4.4	Results of Hybrid Regime	45
4.5	Performance	47
5	Conclusions	49
	Future Work	50
	Bibliography	51
A	Definition of Mapping Jacobians	57
A.1	For Finite-rate Regime	58
A.2	For Equilibrium Regime	60
	Vita	62

List of Tables

4.1	Performance comparison - Implicit solve (seconds)	47
4.2	Performance comparison - Explicit solve (seconds)	47

List of Figures

2.1	Illustration of fluxes used near hybrid interfaces	18
2.2	Illustration of residual contributions near hybrid interfaces	19
2.3	Illustration of equilibrium cell residual shifting and summation operation	20
4.1	Refinement study: perfect gas compressible Euler equations - density plot	36
4.2	Refinement study: perfect gas compressible Euler equations - velocity plot	37
4.3	Refinement study: perfect gas compressible Euler equations - pressure plot	38
4.4	Comparison: density comparison plot of four regimes	40
4.5	Comparison: velocity comparison plot of four regimes	41
4.6	Comparison: pressure comparison plot of four regimes	42
4.7	Species mass fractions N_2 and O_2 comparison	43
4.8	Species mass fractions N and O comparison	43
4.9	Species mass fractions NO comparison	44
4.10	Percentage of cells switched from equilibrium as solution progresses .	46

Nomenclature

$\alpha_{r,k}$	collisional efficiency for species k and reaction r
β	extrapolated state coefficient
Δq_i^n	nonconservative variable update for cell i and time level n
Δt	time step
δ	variable jump vector
$\dot{W}$	chemical source term vector
$\dot{w}_i$	mass production rate for species i
Γ	effective concentration
γ	specific heat ratio
$\hat{n}_x, \hat{n}_y, \hat{n}_z$	face area vector components in x , y , and z directions
$\hat{n}$	normalized face area vector
$\int_{\Omega}$	control volume integral
$\int_{\partial\Omega}$	control surface integral
κ, ϕ	extrapolation parameters
λ	eigenvalue

$\mathcal{M}$	mass constraint
$\mathcal{V}$	volume
$\nu''_{i,r}$	product stoichiometric coefficient for reaction r and species i
$\nu'_{i,r}$	reactant stoichiometric coefficient for reaction r and species i
$\mathfrak{R}$	residual
ρ	mixture density
ρ_i	density of species i
σ	equilibrium solver update scaling factor
θ	covariant velocity, $\hat{n}_x u + \hat{n}_y v + \hat{n}_z w + a_t$
$\tilde{R}$	mixture gas constant
$\vec{n}$	control volume face area vector
$\vec{U}$	flow field velocity vector, $u\hat{i} + v\hat{j} + w\hat{k}$
$\vec{V}$	control volume face velocity vector, $V_x\hat{i} + V_y\hat{j} + V_z\hat{k}$
ξ_1, ξ_2	temporal accuracy parameters
ζ	equilibrium solver maximum allowed relative change in species
$a_1, a_2, a_3, a_4, a_5, a_6$	thermochemistry curvefit coefficients
$a_{i,j}$	number of particles of element i in species j
a_t	control volume face velocity in direction of normal
c_L, c_R	speed of sound to the left and right of a face
c_p	specific heat at constant pressure

c_v	specific heat at constant volume
d	equilibrium solver equation of state
Da	local Damköhler number
e_t	specific total energy
Ea	activation energy
F	conservative flux vector
h	specific enthalpy
h_i^o	standard state enthalpy for species i
h_t	specific total enthalpy
$K_{b,r}$	backward reaction rate for reaction r
$K_{c,r}$	equilibrium rate constant for reaction r
$K_{f,r}$	forward reaction rate for reaction r
M_k	molecular weight of species k
NE	number of elemental species
NS	number of species
P	pressure
P_{std}	standard pressure = 1 atm
Q	conservative variable vector
q	nonconservative variable vector
R	specific mixture gas constant

R_i	specific gas constant
R_{univ}	universal gas constant
s	specific entropy
s_i^o	standard state entropy for species i
S_L, S_R, S_*	wave speed estimates
T	temperature
u, v, w	flow field velocity components in x , y , and z directions
V_x, V_y, V_z	grid velocities in x , y , and z directions
x, y, z	Cartesian coordinate directions
Y_i	mass fraction
z''	sum of product stoichiometric coefficients
z'	sum of reactant stoichiometric coefficients

Chapter 1

Introduction

The importance of high-fidelity Computational Fluid Dynamics (CFD) simulations has grown enormously in recent years. The search for new methods to increase modeling accuracy is of great interest, not only in academia, but also in industrial settings [1]. CFD is capable of modeling physics that are either too costly or too complicated to address with physical experiments. This ability helps drive the rapid dissemination of high technology solutions for public use.

One of the many facets included in this search is reacting chemistry. The addition of chemical reactions to fluid models can lead to new observations and more efficient design in areas such as combustion [2], hypersonic flow [3], and atmospheric propagation of chemically active components. The physics modeled by including chemical reactions in CFD are not trivial nor easily dismissed. In hypersonic flows, the addition of chemically reactive terms can significantly affect shock stand-off distances [4]. This change in shock characteristic can affect dynamic loading of a vehicle and therefore the required aerodynamic or structural design. The investigation of chemically active species as they are transported through Earth's atmosphere is a complex problem. Pollutant oxidation and breakdown is also difficult to observe and computer modeling stands as one of the more useful methods for investigating these environmental issues. The recent increase in concern for environmental emissions

makes these simulations very helpful in understanding our effect on the planet’s ecosystems [5, 6]. The utility of modeling chemistry in the context of combustion is obvious. Design of combustors is highly dependent on chemical kinetics, as well as flow field character [2]. Despite the usefulness of these methods, the increased fidelity of the modeling does not come without additional cost.

Equilibrium chemistry models are an attempt at decreasing the total cost of a simulation, with respect to finite-rate methods, while still including some relevant fluid property changes. Equilibrium methods are formulated such that only total mixture density, ρ , is solved as part of the system of equations. In doing this, the species densities, ρ_i ’s, are eliminated from the iterative solving step. They are instead updated only once per time step through a separate equilibrium chemistry solver. This process brings a nontrivial time savings via the reduction of the size of the system of equations. Several researchers have developed these methods with great success [7]. These models are useful for studying problems with very fast chemistry in comparison to the flow time scales.

Finite-rate chemistry includes individual species transport, diffusion and production terms, but can be very costly to implement for large numbers of species. In particular, an additional species continuity equation is required for each species in the flow. While this is necessary for tracking species transport, it makes the solution of the resulting linear system very expensive. For complex chemistry, with full reaction path modeling, the number of species and reactions modeled grows very quickly. Attempts have been made to eliminate noncritical reactions from simulations through the use of so-called global models [8]. Research has also been done on identifying constraining reactions in order to allow only those reactions that are necessary to be fully finite-rate [9, 10, 11]. However, most of this work relies on a priori knowledge of the chemistry involved and detailed analysis of the reaction rates within the context of the particular flow.

In this work, a method is presented that attempts to retain finite-rate accuracy in regions where the chemical and flow time scales are similar while reducing cost

in regions where the chemistry is reacting faster than the individual species are transported out of a localized area. A local Damköhler number [12] is defined and used as a trigger to switch chemistry modeled in each control volume from equilibrium to finite-rate as appropriate. A method of flux restriction is also given so that the solution method remains consistent along equilibrium and finite-rate control volume interfaces. Tenasi, a University of Tennessee SimCenter research code, is used as a framework for implementing the changes brought about by the new method. Results and timings from a one dimensional transient shock tube are included as verification of the method’s applicability to a commonly encountered CFD test problem. Results are also included from both equilibrium and full finite-rate versions of the solver as a method of comparison.

Chapter 2

Computational Formulation

This chapter outlines the mathematics involved in implementing both a finite-rate and equilibrium Euler equation solver. The hybrid regime is developed and serves as the basis for the remainder of this work. The necessary numerical fluxes, discretization methods and discrete chemistry models are also examined.

The primary motivation driving the development of a hybrid chemistry regime is the reduction of computational cost associated with the use of very high degree of freedom chemistry models, i.e. large number of species and reactions. Since this method switches control volumes in the domain during runtime from equilibrium to finite-rate and back again, both methods of computing chemistry must be available to the solver.

2.1 Governing Equations

The Euler equations were used during this work in order to limit the number of influences on convergence, solution and cost. These methods should also be fully applicable to viscous flows with the usual considerations implied therein.

In this work, the hybrid method has been implemented in conjunction with an explicit time marching algorithm in one dimension. However, it is readily extensible to higher dimensions without much additional consideration and the full 3D form of

the equations is given. The issue of running an implicit solver is of much greater difficulty since the block sizes in the left-hand side matrix may be of different sizes.

The Euler equations can be written in integral (conservation) form as

$$\frac{\partial}{\partial t} \int_{\Omega} Q \, d\mathcal{V} + \int_{\partial\Omega} F \cdot \hat{n} \, dA = \dot{W} \quad (2.1)$$

where the solution vector, Q , and the flux vector, F , will be defined uniquely in each of the regimes.

The variables Q given in Equation 2.1 are of the conservative type. While the fluxes must be conservative for continuity reasons, the stored variables need not be. In fact, with regards to chemically reactive flows, it is often more computationally prudent to store them in their nonconservative form in order to make pressure information readily available without computation. The nonconservative variables are referred to as q . The transformation between these and the form given above is a function of the variables themselves and the thermodynamic state defined by the chemistry information available in that particular control volume.

2.1.1 Finite-rate Chemistry Regime

Finite-rate chemistry is developed under the assumption that each discrete chemical state must be tracked. This is the most general chemistry method and is completely applicable to cases where chemical action happens on a time scale that is slower, at least in some reactions, than the transport of fluid via advection or diffusion. The requirement that each individual species in the flow field be tracked explicitly can make this chemistry regime quite expensive, as the number of equations in the Euler set is now $NS + 4$. If the regime is extended to run implicitly, as has been done in Tenasi, the linear sparse matrix solver begins to dominate the solution time. The cost of a symmetric Gauss-Seidel (SGS) solve on this system scales as $\sim (NS + 4)^2$, which becomes quite expensive very quickly.

The variable set in the finite-rate case is defined by

$$Q = \begin{bmatrix} \rho_1 \\ \rho_2 \\ \vdots \\ \rho_i \\ \vdots \\ \rho_{NS} \\ \rho u \\ \rho v \\ \rho w \\ \rho e_t \end{bmatrix}, q = \begin{bmatrix} \rho_1 \\ \rho_2 \\ \vdots \\ \rho_i \\ \vdots \\ \rho_{NS} \\ u \\ v \\ w \\ P \end{bmatrix}, \dot{W} = \begin{bmatrix} \dot{w}_1 \\ \dot{w}_2 \\ \vdots \\ \dot{w}_i \\ \vdots \\ \dot{w}_{NS} \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad (2.2)$$

where Q is the vector of conservative variables, q is the stored-vector or nonconservative variables, and $\dot{W}$ is the chemical source term associated with the mass production for each tracked species.

The conservative flux associated with this regime is

$$F \cdot \hat{n} = \begin{bmatrix} \rho_1 \theta \\ \rho_2 \theta \\ \vdots \\ \rho_i \theta \\ \vdots \\ \rho_{NS} \theta \\ \rho u \theta + P \hat{n}_x \\ \rho v \theta + P \hat{n}_y \\ \rho w \theta + P \hat{n}_z \\ \rho h_t \theta - a_t P \end{bmatrix} \quad (2.3)$$

where

$$\theta = \hat{n}_x u + \hat{n}_y v + \hat{n}_z w + a_t \quad (2.4)$$

and the grid velocity term a_t is defined as

$$a_t = -[\hat{n}_x V_x + \hat{n}_y V_y + \hat{n}_z V_z] \quad (2.5)$$

where $\vec{V}$ is the control volume face velocity vector.

The chemical source term for each species in the equation set is defined as a combination of the rate contributions from each reaction in which that species participates. It is written as

$$\dot{w}_i = M_i \sum_{r=1}^{NR} (\nu''_{i,r} - \nu'_{i,r}) \Gamma \left[K_{f,r} \prod_{k=1}^{NS} \left(\frac{\rho_k}{M_k} \right)^{\nu'_{k,r}} - K_{b,r} \prod_{k=1}^{NS} \left(\frac{\rho_k}{M_k} \right)^{\nu''_{k,r}} \right] \quad (2.6)$$

where $M_{i/k}$ is the molecular weight of each species indexed either i or k . The terms $\nu'_{i/k,r}$ and $\nu''_{i/k,r}$ are the left and right-hand chemical stoichiometric coefficients respectively for a species indexed either i or k . The coefficient $K_{f,r}$ is the forward rate of a particular reaction and $K_{b,r}$ is the backward rate. It is worth noting that $K_{f,r}$ and $K_{b,r}$ have units that are dependent on the stoichiometric coefficients. That is, $K_{f,r}$ has units of $(\frac{mol}{m^3})^{1-z'} \frac{1}{s}$ and $K_{b,r}$ has units of $(\frac{mol}{m^3})^{1-z''} \frac{1}{s}$, where z' is the sum of the left-hand side coefficients and z'' is the sum of the right-hand side coefficients for a particular reaction. The term Γ is used as a modifier for the rate of progress of a reaction if a third body is present. It is considered an effective concentration of a species for cases where there is enhanced collisional efficiency $\alpha_{r,k}$ [13]. This effective concentration can be higher than the actual concentration. If all species contribute equally to a reaction, $\alpha_{r,k} = 1$. This term can be computed via the expression

$$\Gamma = \sum_{k=1}^{NS} \alpha_{r,k} \frac{\rho_k}{M_k} \quad (2.7)$$

The backward rate of a reaction is typically determined from the combination of a forward rate ($K_{f,r}$), which is computed as a function of temperature, and the equilibrium constant ($K_{c,r}$), which is computed via standard thermodynamic properties of the involved reactants and products. A full description can be found in Section 2.6.2.

2.1.2 Equilibrium Chemistry Regime

For flow fields in which the assumption of “fast” chemistry can be made, equilibrium chemistry offers a possible cost reduction over full finite-rate analysis. This method is developed on the assumption that chemistry changes happen much faster than advection or diffusion and can thus be modeled as being instantaneous. In short, it assumes that all cells are in local equilibrium. This idea is by no means perfect but it does increase physical phenomenon modeling capabilities for certain classes of flow problems.

For the equilibrium regime, the conservative and nonconservative (stored) variables are

$$Q = \begin{bmatrix} \rho \\ \rho u \\ \rho v \\ \rho w \\ \rho e_t \end{bmatrix}, q = \begin{bmatrix} \rho \\ u \\ v \\ w \\ P \end{bmatrix}, \dot{W} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad (2.8)$$

where $\dot{W}$ is identically zero since chemical composition is adjusted instantaneously through the equilibrium solver. The interest is in local bulk thermodynamic properties only, which is the primary contribution of the chemistry to the model.

The numerical conservative flux is given as

$$F \cdot \hat{n} = \begin{bmatrix} \rho\theta \\ \rho u\theta + P\hat{n}_x \\ \rho v\theta + P\hat{n}_y \\ \rho w\theta + P\hat{n}_z \\ \rho h_t\theta - a_t P \end{bmatrix} \quad (2.9)$$

where θ and a_t have been defined in Equations 2.4 and 2.5, respectively.

Cox developed a method for solving for equilibrium composition given a particular thermodynamic state in his work [7]. This is implemented and used here for all equilibrium solving duties. The nonlinear system that must be solved can be written in the standard form, $Ax = b$.

$$A = \begin{bmatrix} \frac{\partial \dot{w}_1}{\partial \rho_1} & \cdots & \frac{\partial \dot{w}_1}{\partial \rho_{NS}} & \frac{\partial \dot{w}_1}{\partial T} \\ \vdots & \ddots & \vdots & \vdots \\ \frac{\partial \dot{w}_{NS-NE}}{\partial \rho_1} & \cdots & \frac{\partial \dot{w}_{NS-NE}}{\partial \rho_{NS}} & \frac{\partial \dot{w}_{NS-NE}}{\partial T} \\ \frac{\partial \mathcal{M}_1}{\partial \rho_1} & \cdots & \frac{\partial \mathcal{M}_1}{\partial \rho_{NS}} & \frac{\partial \mathcal{M}_1}{\partial T} \\ \vdots & \ddots & \vdots & \vdots \\ \frac{\partial \mathcal{M}_{NE}}{\partial \rho_1} & \cdots & \frac{\partial \mathcal{M}_{NE}}{\partial \rho_{NS}} & \frac{\partial \mathcal{M}_{NE}}{\partial T} \\ \frac{\partial d}{\partial \rho_1} & \cdots & \frac{\partial d}{\partial \rho_{NS}} & \frac{\partial d}{\partial T} \end{bmatrix} \quad (2.10)$$

where NS is the total number of species present in the flow field and NE is the number of elemental species represented. The system consists of $NS + 1$ equations. There are $NS - NE$ mass action law equations, $\dot{w}_i$, and NE elemental species mass constraint equations, $\mathcal{M}$. The number of elemental species need not be true elemental species but must contain the element that they represent. For example, HO may serve as a representative for either elemental hydrogen H or elemental oxygen O . The final

equation is an energy equation for updating mixture temperature. The rest of the system is

$$x^T = [\Delta\rho_1, \dots, \Delta\rho_{NS}, \Delta T] \quad (2.11)$$

$$b^T = -[\dot{w}_1, \dots, \dot{w}_{NS-NE}, \mathcal{M}_1, \dots, \mathcal{M}_{NE}, d] \quad (2.12)$$

Here, d denotes the caloric equation of state written in terms of the total energy per unit volume

$$d = \rho e_t - \sum_{j=1}^{NS} \rho_j \left[\int_{T_{ref}}^T c_{v_j}(T) dT + h_{f_j} \right] - \rho \frac{\|\vec{U}\|^2}{2} = 0 \quad (2.13)$$

where specific heats at constant volume, c_{v_j} , are a function of temperature and computed via curve fit data. The total energy per unit volume ρe_t and velocity $\vec{U}$ are available from flow field computations. The term h_{f_j} is the heat of formation at T_{ref} . The term $\dot{w}$ denotes the law of mass action which has been modified for equilibrium chemistry.

$$\dot{w}_i(\rho_j, T) = M_i \sum_{r=1}^{NR} (\nu''_{i,r} - \nu'_{i,r}) \Gamma \left[K_{c,r} \prod_{k=1}^{NS} \left(\frac{\rho_k}{M_k} \right)^{\nu'_{k,r}} - \prod_{k=1}^{NS} \left(\frac{\rho_k}{M_k} \right)^{\nu''_{k,r}} \right] = 0, \quad (2.14)$$

$i = 1, \dots, NS - NE$

where the index i only covers nonelemental species. This is derived from Equation 2.6 by setting $K_{b,r} = 1$ and substituting the definition from Equation 2.86 for $K_{f,r}$. This modification models the limiting case of chemical equilibrium, which is an infinite backwards reaction rate. This avoids the difficulties involved in attempting to determine those limit reaction rates. At equilibrium, the bracketed terms are identically zero. The elemental mass conservation equations ensure that mass is conserved and are written as

$$\mathcal{M}_i = \sum_{j=1}^{NS} \frac{a_{ij} \rho_j}{M_j \rho} - \left[\sum_{j=1}^{NS} \frac{a_{ij} \rho_j}{M_j \rho} \right]_{t=0} = 0, \quad (2.15)$$

$i = 1, \dots, NE$

where a_{ij} is the total number of particles of element i in species j . The index i covers only species designated as elemental. These are strictly referenced back to $t = 0$ total mass fractions in a pure equilibrium regime as indicated. For the hybrid chemistry regime, they are referenced to the beginning of each time step on a control volume basis. There is no concept of globally defined mass fractions in a flow field where mass injection and extraction can be tracked on boundaries.

The partial derivatives of the laws of mass action that appear in the Jacobian matrix A are

$$\frac{\partial \dot{w}_i}{\partial \rho_j} = \frac{M_i}{M_j} \sum_{r=1}^{NR} (\nu''_{i,r} - \nu'_{i,r}) \Gamma K_{b,r} \left[K_{c,r} \prod_{k=1}^{NS} \left(\frac{\rho_k}{M_k} \right)^{\bar{\nu}'_{k,r}} \nu'_{j,r} - \prod_{k=1}^{NS} \left(\frac{\rho_k}{M_k} \right)^{\bar{\nu}''_{k,r}} \nu''_{j,r} \right] \quad (2.16)$$

$i = 1, \dots, NS - NE$

$$\frac{\partial \dot{w}_i}{\partial T} = M_i \sum_{r=1}^{NR} (\nu''_{i,r} - \nu'_{i,r}) \Gamma K_{b,r} \left[\frac{dK_{c,r}}{dT} \prod_{k=1}^{NS} \left(\frac{\rho_k}{M_k} \right)^{\nu'_{k,r}} \right] \quad (2.17)$$

$i = 1, \dots, NS - NE$

where in Equation 2.16

$$\bar{\nu}'_{k,r}, \bar{\nu}''_{k,r} = \begin{cases} \nu'_{k,r}, \nu''_{k,r} & \text{if } k \neq j, \\ \nu'_{k,r} - 1, \nu''_{k,r} - 1 & \text{if } k = j \text{ and } \nu'_{k,r}, \nu''_{k,r} \neq 0, \\ 0 & \text{if } \nu'_{k,r}, \nu''_{k,r} = 0 \end{cases} \quad (2.18)$$

The mass constraint partial derivatives are expressed as

$$\frac{\partial \mathcal{M}_i}{\partial \rho_j} = \frac{a_{ij}}{\rho M_j} \quad (2.19)$$

$i = 1, \dots, NE$

$$\frac{\partial \mathcal{M}_i}{\partial T} = 0 \quad (2.20)$$

$i = 1, \dots, NE$

The partial derivatives of the equation of state d are

$$\frac{\partial d}{\partial \rho_j} = e_j = \int_{T_{ref}}^T c_{v_j} dT + h_{f_j} \quad (2.21)$$

$$\frac{\partial d}{\partial T} = \rho c_v \quad (2.22)$$

where c_v is the mixture specific heat given in Section 2.6.

The above system is for the situation where the complete thermodynamic system (CTS) is known in terms of density and internal energy, $CTS = f(\rho, e)$. There are situations that arise in which the system is known in terms of pressure and density, $CTS = f(P, \rho)$, or in terms of pressure and temperature, $CTS = f(P, T)$. For the case where pressure and density are given, the linearized system remains very similar. The only difference is that closure is now provided by the thermal equation of state. It is defined as

$$d = \sum_{j=1}^{NS} \rho_j R_j T - P_0 = 0 \quad (2.23)$$

where P_0 is the initial pressure. The new partial derivatives for the last equation in the system are now

$$\frac{\partial d}{\partial \rho_j} = R_j T \quad (2.24)$$

$$\frac{\partial d}{\partial T} = \rho R \quad (2.25)$$

where R is the specific mixture gas constant.

The case of known pressure and temperature, $CTS = f(P, T)$ requires more modifications. The dependent variables in this case are now partial densities and total density, with the thermal equation of state being replaced by the caloric equation of state. Here, the total density equation must be left in place to couple the state equation to the system, despite the fact it seems to duplicate information contained

in the partial densities. The new linearized system for $CTS = f(P, T)$ is

$$A = \begin{bmatrix} \frac{\partial \dot{w}_1}{\partial \rho_1} & \cdots & \frac{\partial \dot{w}_1}{\partial \rho_{NS}} & \frac{\partial \dot{w}_1}{\partial \rho} \\ \vdots & \ddots & \vdots & \vdots \\ \frac{\partial \dot{w}_{NS-NE}}{\partial \rho_1} & \cdots & \frac{\partial \dot{w}_{NS-NE}}{\partial \rho_{NS}} & \frac{\partial \dot{w}_{NS-NE}}{\partial \rho} \\ \frac{\partial \mathcal{M}_1}{\partial \rho_1} & \cdots & \frac{\partial \mathcal{M}_1}{\partial \rho_{NS}} & \frac{\partial \mathcal{M}_1}{\partial \rho} \\ \vdots & \ddots & \vdots & \vdots \\ \frac{\partial \mathcal{M}_{NE}}{\partial \rho_1} & \cdots & \frac{\partial \mathcal{M}_{NE}}{\partial \rho_{NS}} & \frac{\partial \mathcal{M}_{NE}}{\partial \rho} \\ \frac{\partial d}{\partial \rho_1} & \cdots & \frac{\partial d}{\partial \rho_{NS}} & \frac{\partial d}{\partial \rho} \end{bmatrix} \quad (2.26)$$

and

$$x^T = [\Delta \rho_1, \dots, \Delta \rho_{NS}, \Delta \rho] \quad (2.27)$$

with the right-hand side vector remaining the same. The partial derivatives of the law of mass action with respect to the species densities remain the same as in Equation 2.16. The partial derivatives of the law of mass action with respect to total density are zero. The partial derivatives of the mass constraints are

$$\begin{aligned} \frac{\partial \mathcal{M}_i}{\partial \rho_j} &= \frac{a_{ij}}{\rho M_j} \\ i &= 1, \dots, NE \end{aligned} \quad (2.28)$$

$$\begin{aligned} \frac{\partial \mathcal{M}_i}{\partial \rho} &= -\frac{1}{\rho} \sum_{j=1}^{NS} \frac{a_{ij}}{M_j} Y_j \\ i &= 1, \dots, NE \end{aligned} \quad (2.29)$$

The thermal equation of state for constant temperature and pressure can be written as

$$d = \sum_{j=1}^{NS} \rho_j R_j T_0 - P_0 = 0 \quad (2.30)$$

where the terms P_0 and T_0 are the initial conditions. The partial derivatives for the equation of state are now

$$\frac{\partial d}{\partial \rho_j} = R_j T_0 \quad (2.31)$$

$$\frac{\partial d}{\partial \rho} = 0 \quad (2.32)$$

The two limiters described by Cox as “catastrophic” and “relative” are also implemented. It can be shown that as long as the scaling is applied to all species updates, the mass constraints will not be violated. However, the linearized system will not be satisfied unless the temperature is also modified. This can be remedied by simply resolving the last row of the governing system for the modified temperature update, using the modified species updates. The scaling factor is given as

$$\sigma = \max \left\{ \sigma_{min}, \left[\min_i \left(1, \frac{\zeta \rho_i^n}{|\Delta \rho_i^{old}|} \right), \rho_i > 0, i = 1, \dots, NS \right] \right\} \quad (2.33)$$

where σ_{min} is a minimum value used to prevent the limiter from dramatically slowing down the convergence of the solver. The value ζ is the maximum allowable relative change in the species and is defined as

$$\zeta \geq \left| \frac{\Delta \rho_i^{new}}{\rho_i^n} \right| \quad (2.34)$$

$$i = 1, \dots, NS$$

The primary difference between the two limiters is that the relative limiter is applied at every iteration and the catastrophic limiter is applied only when an update would cause a negative density. The limiter is applied as

$$\Delta \rho_i^{new} = \sigma \Delta \rho_i^{old} \quad (2.35)$$

$$i = 1, \dots, NS$$

It is important to note that negative species densities cannot be prevented should a density be identically zero or an update be negative. In this case, the negative species density is excluded from the scaling factor evaluation, σ . This allows the negative density to again become positive where it would otherwise be inhibited.

2.1.3 Hybrid Chemistry Regime

After observing the characteristics of both finite-rate and equilibrium chemistry models, it is simple to formulate a class of problems for which neither is very well suited. Finite-rate chemistry is detailed enough to model most classes of chemistry problems with high accuracy. It performs very well, given good information on reaction rates and a reliable method of determining thermodynamic properties for each species involved over the temperature range of interest. However, as stated before, this method can get quite expensive in terms of computational cost. Equilibrium chemistry models, on the other hand, work well for “fast” chemistry but lack accuracy in regions where this assumption is not valid.

The hybrid scheme given here attempts to maintain the accuracy of finite-rate methods while leveraging the cost savings associated with equilibrium chemistry solvers. This method uses a full finite-rate flux in regions where chemistry is slow compared to the advection rate and an equilibrium chemistry scheme in regions where the chemistry outpaces the fluid transport. This maintains the full resolution of finite-rate chemistry where needed and then falls back to equilibrium in the other regions to reduce the overall cost of the simulation. This could be extremely useful for full reaction path modeling or the tracking of a very large number of species. The cost of symmetric Gauss-Seidel iterations grows like $\sim (NS + 4)^2$. Thus, eliminating this cost in regions where it is unjustified can be very powerful.

The conservative Q and nonconservative q variables needed in each control volume are dictated by examining each local state and making the decision of which regime is more appropriate given the current flow conditions. The hybrid method chosen here

is based on a locally defined Damköhler number (Da) [12]. If Da is above a certain threshold, then the advection time is assumed to dominate the local flow character and the control volume is switched to equilibrium chemistry. If below the threshold, finite-rate chemistry is used in an attempt to capture any partial reaction product transport. This local Damköhler number for species i is defined here as

$$Da_i = \frac{t_{res}}{t_{chem_i}} = \frac{|\dot{w}_i|L_c}{\rho_i\|\vec{U}\|} \quad (2.36)$$

where $|\dot{w}_i|$ is the magnitude of the mass production rate, L_c is a characteristic length associated with the current control volume, ρ_i is the density of species i and $\|\vec{U}\|$ is the local velocity magnitude. This is computed for each species and the minimum for each cell is chosen to perform the switching. The minimum is chosen with a preference for accuracy over cost savings. Simply stated, if any species in a cell dictates the resolution of finite-rate chemistry, the entire cell uses finite-rate chemistry. This switch is currently computed and checked at every time step, but could be frozen for several time steps to eliminate the expense of the Da number calculation should the flow physics be evolving “gently” enough.

Obviously, this hybrid scheme leads to a mixture of finite-rate cells and equilibrium cells throughout the domain. This brings into focus some problems related to the mixing of two distinct methods that must be resolved.

Firstly, due to memory consistency issues, a single way of storing the independent variable vector must be selected and maintained. It is possible to change the memory storage format as required, but this is likely to segment heap-allocated continuous memory arrays and cause unnecessary cache thrashing. Unstructured meshes tend to have less than ideal data reuse patterns, so a poor implementation in variable storage could cause serious efficiency issues. Unfortunately, there is not much of a choice here. Tracking of individual species is required when using finite-rate; therefore, this is the data structure that must be used should the control volume be switched to the finite-rate method. Each vector of unknowns is stored in memory as in the finite-rate

method shown in Equation 2.2. Since the total mixture density is needed instead of partial densities when a cell happens to switch to equilibrium, a way to derive this from stored variables is needed. This is done via a simple summation

$$\rho = \sum_i^{NS} \rho_i \quad (2.37)$$

This implementation provides a method of storing independent variables consistently while requiring no additional computation to switch equation sets.

The next issue that must be addressed is how to deal with fluxes on interfaces between finite-rate and equilibrium control volumes. Fluxes are accumulated to the residual by looping over each edge and making contributions directly to each control volume’s specified array location in the right-hand side of the block linear system. Care must be taken to ensure that contributions to each residual location are of the appropriate “type”. If this were done in a naive way without considering interfaces between equilibrium and finite-rate cells, the residual would be garbled. Take for example a particular edge that happens to lie between an equilibrium cell and a finite-rate cell. The first flux contribution from the equilibrium side would be placed in the residual location for ρ_1 and the contribution from the finite-rate side would be placed in the ρ residual location for the equilibrium side. Placing a contribution from ρ_1 into the residual location for ρ is very close to what should be achieved but the first case of putting a bulk density contribution into a partial density residual location is unacceptable.

The solution turns out to be a bit more complicated in implementation but is very general and works quite well. Since the problem encountered is that the flux from the equilibrium cell is inconsistent with a finite-rate structured residual vector, the decision is made to default to full finite-rate fluxes for any cell (finite-rate or equilibrium) which lies on one of these interfaces. This is illustrated in Figure 2.1.

Figure 2.1: Illustration of fluxes used near hybrid interfaces

Figure 2.2: Illustration of residual contributions near hybrid interfaces

Though hybrid interface fluxes are now consistent, the same problem now exists between equilibrium cells on hybrid interfaces and those surrounded by other equilibrium cells. Consider an equilibrium cell on a hybrid interface and one of its equilibrium neighbors. The flux computed on the edge between them will be of the equilibrium type but since the cell on the hybrid interface must be able to accept finite-rate fluxes, the equilibrium flux will overwrite useful information in its residual accumulation vector. The solution here turns out to be quite intuitive. By shifting the flux to accumulate to the ρ_{NS} location and below, coherency is maintained. In particular, all mass flux contributions from equilibrium cells, $(\rho\theta)_{eq}$, to the residual vector will be in location NS only. The flux contributions from any finite-rate neighbors will be scattered in any of the first NS locations. The locations to which the scattering occurs are shown in Figure 2.2. However, since only total mass flux and not species specific mass flux is needed for the equilibrium interface cell, a simple summation and reshift before solving puts all the variables in the expected locations. The reshift is not explicitly required; however, the need to modify stride distances (jump required to reach the memory location of the next vector) in the core solver of the code is eliminated and it is likely that this gives a more predictable data access pattern for compiler optimization. This is illustrated in Figure 2.3.

$$\begin{array}{ccc}
\left[\begin{array}{c} \text{Equilibrium} \\ \text{before shifting} \\ \rho_1\theta \\ \rho_2\theta \\ \vdots \\ \rho_i\theta \\ \vdots \\ \rho_{NS} + (\rho\theta)_{eq} \\ \rho u\theta + P\hat{n}_x \\ \rho v\theta + P\hat{n}_y \\ \rho w\theta + P\hat{n}_z \\ \rho h_t\theta - a_tP \end{array} \right] & \Uparrow & \left[\begin{array}{c} \text{Equilibrium} \\ \text{after shifting} \\ \rho\theta \\ \rho u\theta + P\hat{n}_x \\ \rho v\theta + P\hat{n}_y \\ \rho w\theta + P\hat{n}_z \\ \rho h_t\theta - a_tP \\ 0 \\ \vdots \\ 0 \\ 0 \\ 0 \end{array} \right]
\end{array}$$

Figure 2.3: Illustration of equilibrium cell residual shifting and summation operation

The higher-order finite-rate flux components can be restricted like

$$\rho\theta = \sum_{i=1}^{NS} \rho_i\theta + (\rho\theta)_{eq} \quad (2.38)$$

with the flux component in location NS containing both total mixture density contributions from equilibrium cells, $(\rho\theta)_{eq}$, as well as the partial density contributions of species NS from any adjacent finite-rate cells. Components $NS+1$ through $NS+4$ - the momentum and energy variables - are simply shifted to the correct residual location for solving the block linear system.

2.2 HLLC Flux

High resolution near discontinuities is of the utmost importance for compressible CFD solvers where shocks are a common occurrence. These cases can be reduced to a Riemann problem. Several approximate Riemann solvers are available and in widespread use in CFD. One such method is given by Roe [14], however, it requires the full decomposition of the eigensystem. Instead, the approximate Riemann solver of Harten, Lax, and van Leer (HLL) [15] for two distinct contact waves (HLLC)

is implemented. This avoids the full eigensystem development while still providing accurate results. It was found by Buvaneswari that the differences between a Roe flux and HLLC are almost nonexistent as applied to both equilibrium and nonequilibrium chemically active flows [16]. The development below follows the discussion of Toro [17].

First, the maximum and minimum eigenvalues of the system must be found by

$$\lambda_{4R} = \theta_R + c_R \quad (2.39)$$

$$\lambda_{5L} = \theta_L - c_L \quad (2.40)$$

$$\lambda_4 = \bar{\theta} + \bar{c} \quad (2.41)$$

$$\lambda_5 = \bar{\theta} - \bar{c} \quad (2.42)$$

where c_L, c_R is the speed of sound based on either the left or right state at the control volume face and $\bar{c}$ is some averaged speed of sound between the two states. This can either be a simple average or something more sophisticated such as a Roe average. The term θ has been defined previously, only here the averaged version $\bar{\theta}$ is used. Next, wave speed estimates S_L , S_R , and S_* are defined as

$$S_L = \min(\lambda_{5L}, \lambda_5) \quad (2.43)$$

$$S_R = \max(\lambda_{4R}, \lambda_4) \quad (2.44)$$

$$S_* = \frac{P_R - P_L + \rho_L \theta_L (S_L - \theta_L) - \rho_R \theta_R (S_R - \theta_R)}{\rho_L (S_L - \theta_L) - \rho_R (S_R - \theta_R)} \quad (2.45)$$

and the HLLC state is determined from

$$Q^{HLLC} = \begin{cases} Q_L & \text{if } 0 \leq S_L \\ Q_{*L} & \text{if } S_L \leq 0 \leq S_* \\ Q_{*R} & \text{if } S_* \leq 0 \leq S_R \\ Q_R & \text{if } 0 \geq S_R \end{cases} \quad (2.46)$$

where Q_L and Q_R indicate that the entire state is derived from either the left or the right side of the interface. This is the condition that arises from a supersonic flow, in which the speed of sound is not sufficient to overcome the flow speed, where all values are upwinded. For the cases in which the HLLC state is defined as Q_{*K} , where K is either L or R , the following relations are appropriate.

$$\rho_K^{HLLC} = \frac{S_K - \theta_K}{S_K - S_*} \rho \quad (2.47)$$

and is applicable to bulk density as well as species density. The estimated pressure in the star region is

$$P_* = P_K + \rho_K(\theta_K - S_K)(\theta_K - S_*) \quad (2.48)$$

Since the fluxes are required to be conservative, the variables must be transformed to their equivalent conservative form before passing to the flux computation routine.

This gives us

$$\rho u_K^{HLLC} = \frac{S_K - \theta_K}{S_K - S_*}(\rho_K u_K) + \frac{P_* - P_K}{S_K - S_*} \hat{n}_x \quad (2.49)$$

$$\rho v_K^{HLLC} = \frac{S_K - \theta_K}{S_K - S_*}(\rho_K v_K) + \frac{P_* - P_K}{S_K - S_*} \hat{n}_y \quad (2.50)$$

$$\rho w_K^{HLLC} = \frac{S_K - \theta_K}{S_K - S_*}(\rho_K w_K) + \frac{P_* - P_K}{S_K - S_*} \hat{n}_z \quad (2.51)$$

$$\rho e_t^{HLLC} = \frac{S_K - \theta_K}{S_K - S_*} \rho e_{tK} + \frac{P_* S_* - P_K \theta_K}{S_K - S_*} \quad (2.52)$$

The HLLC method provides a means of obtaining the appropriate upwinding for supersonic flow conditions as well as a characteristic driven blending for subsonic regions in the field.

2.3 Spatial Discretization

The code described uses a finite volume method for the integration of the applicable fluxes. Therefore, a control volume must be defined over which these fluxes will be

conserved. This particular work uses unstructured meshes for discretizing a particular spatial domain. Within these meshes, a control volume is defined as the median dual surrounding any particular node. That is, element centroids are connected to edge midpoints and face centroids to form a closed cell. This is referred to as being node-centered and is fully extendable to 3D. Several good references are available to describe the implementation details [18, 19].

The discretization of spatial terms over each control volume can be expressed as

$$\mathcal{V}_i \frac{\partial Q_i}{\partial t} + \mathfrak{R}_i = 0 \quad (2.53)$$

where $\mathfrak{R}$ is the spatial residual and contains all contributions from numerical fluxes and source terms where applicable [18]. This quantity is defined for every control volume i and $\frac{\partial Q_i}{\partial t}$ here represents the change of the conservative variables with respect to a particular step in time. This expression must be generalized to solve directly for the nonconservative variables via

$$\mathcal{V}_i \frac{\partial Q_i}{\partial q_i} \frac{\partial q_i}{\partial t} + \mathfrak{R}_i = 0 \quad (2.54)$$

where $\frac{\partial Q_i}{\partial q_i}$ can be thought of as a mapping of the conservative solution into nonconservative variable space. These terms arise on the diagonal of the linear system and are shown completely expanded in appendix A.

In this work, a MUSCL extrapolation method was used to extend spatial resolution to second-order accuracy. A full description and development of the method is available in Van Leer [20] and described as applied to the 2D compressible Euler equations in Berthon [21]. Here, this extrapolation can be expressed as

$$\delta_R = Q_{R+1} - Q_R \quad (2.55)$$

$$\delta_L = Q_L - Q_{L-1} \quad (2.56)$$

$$\delta_C = Q_R - Q_L \quad (2.57)$$

which are the jumps used in computing the extrapolated state.

$$\beta_L = \frac{2 \delta_C \delta_L + \epsilon}{\delta_C^2 + \delta_L^2 + \epsilon} \quad (2.58)$$

$$\beta_R = \frac{2 \delta_C \delta_R + \epsilon}{\delta_C^2 + \delta_R^2 + \epsilon} \quad (2.59)$$

where ϵ is a small perturbation to avoid dividing by zero, selected here to be 1×10^{-6} for double precision arithmetic. The states at the control volume interfaces can then be calculated by

$$Q_{L+\frac{1}{2}} = Q_L + \psi_L \phi \beta_L ((1 - \kappa)(\delta_L) + (1 + \kappa)(\delta_C)) \quad (2.60)$$

$$Q_{R-\frac{1}{2}} = Q_R - \psi_R \phi \beta_R ((1 - \kappa)(\delta_R) + (1 + \kappa)(\delta_C)) \quad (2.61)$$

where $Q_{L/R \pm \frac{1}{2}}$ are the cell face extrapolated states and $\psi_{L/R}$ are computed limiters of which many are available [22]. The parameters ϕ and κ control the order of the spatial accuracy. For a first-order scheme $\phi = 0$, second-order and third-order accuracy can be obtained with $\phi = \frac{1}{4}$ with $\kappa = -1$ and $\kappa = \frac{1}{3}$ respectively.

2.4 Temporal Discretization

Next, the time derivative must be discretized with suitable accuracy. During model development it was found that the test problem stability was driven entirely by the chemistry terms, whether finite-rate or equilibrium. Since an explicit solve was used, this dictated extremely small time steps. Due to the time step size being so small, it was found that attempting to run a higher-order time stepping scheme simply added additional cost without the addition of any noticeable increase in temporal resolution. For this reason, only a first-order time stepping scheme was enabled.

For some of the timings given in the performance section, an implicit method was used. Beam and Warming [23] give a general expression for implicit schemes as

$$\Delta q_i^n = \frac{\xi_1 \Delta t}{1 + \xi_2} \frac{\partial}{\partial t} (\Delta q_i^n) + \frac{\Delta t}{1 + \xi_2} \frac{\partial}{\partial t} (\Delta q_i^n) + \frac{\xi_2}{1 + \xi_2} \Delta q_i^{n-1} \quad (2.62)$$

where a first-order time accurate Euler implicit scheme is given by $\xi_1 = 1$, $\xi_2 = 0$ and second-order time Euler implicit method by $\xi_1 = 1$, $\xi_2 = 1/2$. The term $\Delta q_i^n = q_i^{n+1} - q_i^n$. Only a first-order time scheme is used for this work but the second-order scheme is implemented as well.

The above statement can be simplified for the specific case where $\xi_1 = 1$ [18]. This can be written as

$$\Delta q_i^n = \frac{\Delta t}{1 + \xi_2} \frac{\partial}{\partial t} (q_i^{n+1}) + \frac{\xi_2}{1 + \xi_2} \Delta q_i^{n-1} \quad (2.63)$$

Substituting the time derivative from Equation 2.54 gives

$$\frac{\mathcal{V}_i}{\Delta t} \frac{\partial Q_i}{\partial q_i} \left(\Delta q_i^n - \frac{\xi_2}{1 + \xi_2} \Delta q_i^{n-1} \right) = -\frac{\mathfrak{R}_i^{n+1}}{1 + \xi_2} \quad (2.64)$$

2.5 Explicit Solution Technique

For this particular study, an explicit solution technique was used, though not necessarily in the standard sense. Though the fluxes are conservative, the nonconservative variables are actually solved for in order to eliminate repeat calculations of pressure, etc. Therefore, the mapping in Equation 2.54 must be used to arrive at a suitable method for updating the stored variables q_i . Rearranging the spatially and temporally discretized terms into the form $Ax = b$ gives

$$\left(\frac{\mathcal{V}_i}{\Delta t} \right) \left(\frac{\partial Q_i}{\partial q_i} \right) (\Delta q_i^n) = -\mathfrak{R}_i^{n+1} \quad (2.65)$$

which is Equation 2.64 with $\xi_2 = 0$. Note that this form is a first-order temporal accurate method that contains a mapping term $\frac{\partial Q_i}{\partial q_i}$, which can be evaluated either

discretely or analytically, depending on what might be more appropriate. The structure resulting from a mesh discretized in this fashion is a block diagonal matrix which is solved at each node by LU decomposition with partial pivoting and back substitution.

2.6 Chemistry

The chemistry components used in the solver are largely modularized from the compressible Euler equations solver in order to allow quick changes of chemistry models and thermodynamic databases for individual species. This section describes the methods utilized within the chemistry library for this research. The particular databases and reaction sets used for verification are detailed in the problem setup section in Chapter 4.

2.6.1 Thermodynamic Data

Determining thermodynamic properties of individual chemical species is not a trivial task. One method of determining these properties is through the use of NASA 7/9 coefficient polynomials. Here, the 7 coefficient variety is used exclusively, though for most species the 9 coefficient versions are also available. These are significantly more accurate, with the error at maximum temperature being reduced one to two orders of magnitude from the older curve fits. The accuracy is not necessarily an issue in either case, with the 7 coefficient version of the polynomials typically exhibiting one-tenth of one percent to one percent error at peak temperatures [24].

The thermodynamic curve fit coefficients are valid over two semi-standard ranges, 200-1000K and 1000-6000K, with some researchers [25] adding additional coefficients for ranges beyond these tables. The thermodynamic properties of a given species can be derived from these curve fits as a function of temperature. The relation used for

specific heats is

$$\frac{c_{p_i}}{R_i} = a_1 + a_2T + a_3T^2 + a_4T^3 + a_5T^4 \quad (2.66)$$

which is the specific heat at constant pressure. The constant volume specific heat can be calculated as

$$c_{v_i} = c_{p_i} - R_i \quad (2.67)$$

where $R_i = \frac{R_{univ}}{M_i}$ is the specific gas constant for the species. Specific enthalpy is calculated via

$$\frac{h_i}{R_iT} = a_1 + \frac{a_2T}{2} + \frac{a_3T^2}{3} + \frac{a_4T^3}{4} + \frac{a_5T^4}{5} + \frac{a_6}{T} \quad (2.68)$$

Specific entropy can be calculated with

$$\frac{s_i}{R_i} = a_1 \ln T + a_2T + \frac{a_3T^2}{2} + \frac{a_4T^3}{3} + \frac{a_5T^4}{4} + a_7 \quad (2.69)$$

Gibbs free energy, which can be used as a fundamental indication of an equilibrium state by minimizing the total free energy in a mixture, can be found by

$$\frac{g_i}{R_iT} = \frac{h_i}{R_iT} - \frac{s_i}{R_i} = a_1(1 - \ln T) - \frac{a_2}{2}T - \frac{a_3}{6}T^2 - \frac{a_4}{12}T^3 - \frac{a_5}{20}T^4 + \frac{a_6}{T} - a_7 \quad (2.70)$$

All of the above properties are calculated on a per species basis and are not very useful with regards to a CFD solver. What is desired are the bulk fluid properties in a particular control volume. First, the mass fraction for any species i is defined as

$$Y_i = \frac{\rho_i}{\rho} \quad (2.71)$$

Using the following relations, several useful properties can be computed for a localized flow field.

$$c_v = \sum_{i=1}^{NS} Y_i c_{v_i} \quad (2.72)$$

$$c_p = \sum_{i=1}^{NS} Y_i c_{p_i} \quad (2.73)$$

$$\gamma = \frac{c_p}{c_v} = 1 + \frac{R}{c_v} \quad (2.74)$$

$$R_i = \frac{R_{univ}}{M_i} \quad (2.75)$$

$$R = \sum_{i=1}^{NS} Y_i R_i \quad (2.76)$$

$$e_i = h_i - \frac{P}{\rho} = h_i - R_i T \quad (2.77)$$

Partial pressures are calculated via

$$P_i = \rho_i R_i T \quad (2.78)$$

and total mixture pressure can be computed from

$$P = \sum_{i=1}^{NS} \rho_i R_i T \quad (2.79)$$

Specific total enthalpy for the mixture can be calculated as

$$h_t = \sum_{i=1}^{NS} Y_i h_i + \frac{1}{2} \|U\|^2 \quad (2.80)$$

and likewise, specific total energy can be calculated via

$$e_t = \sum_{i=1}^{NS} Y_i e_i + \frac{1}{2} \|U\|^2 \quad (2.81)$$

The square of the speed of sound is defined as the partial derivative of pressure with respect to density and can be written as

$$c^2 = \left(\frac{\partial P}{\partial \rho} \right)_s \quad (2.82)$$

This can be evaluated analytically as was done by Cox [7]. However, the full evaluation can be quite expensive to compute. Here, an approximation was used and can be written as

$$c^2 = \gamma RT \quad (2.83)$$

where γ and R are both defined for the total mixture of gases.

2.6.2 Reaction Rates

The reaction rates for a particular reaction can be computed via several methods. In this research, the Arrhenius and modified Arrhenius forms are used. The Arrhenius equation is

$$K = Ae^{-E_a/R_{univ}T} \quad (2.84)$$

where A is a prefactor and E_a is the activation energy of the reaction in question. The modified Arrhenius form is

$$K = A \left(\frac{T}{T_0} \right)^n e^{-E_a/R_{univ}T} \quad (2.85)$$

where T_0 is some reference temperature and n is a unitless power. Other forms do exist but are not used here.

Reaction rates are typically given in terms of forward reaction rate, $K_{f,r}$, only. It is then left to the researcher to derive the backward reaction rate, $K_{b,r}$. This can be done with

$$K_{b,r} = \frac{K_{f,r}}{K_{c,r}} \quad (2.86)$$

where the equilibrium coefficient $K_{c,r}$ must be determined first. It can be computed from thermodynamic properties alone [13].

$$K_{p,r} = e^{\left(\frac{\Delta s_{r,r}^o}{R_{univ}} - \frac{\Delta h_{r,r}^o}{R_{univ}T} \right)} \quad (2.87)$$

where

$$\Delta s_{r,r}^o = \sum_{i=1}^{NS} (\nu''_{i,r} - \nu'_{i,r}) s_i^o \quad (2.88)$$

and

$$\Delta h_{r,r}^o = \sum_{i=1}^{NS} (\nu''_{i,r} - \nu'_{i,r}) h_i^o \quad (2.89)$$

in which s_i^o and h_i^o are the standard state entropy and enthalpy, respectively. From this the reaction rate based on concentration can be derived, which is

$$K_{c,r} = K_{p,r} \left(\frac{P_{std}}{R_{univ}T} \right)^{\sum_{i=1}^{NS} (\nu''_{i,r} - \nu'_{i,r})} \quad (2.90)$$

where P_{std} is 1 atmosphere.

Chapter 3

Implementation Issues

When implementing a given algorithm or combination of algorithms within a particular programming environment - here C++ - there is never a shortage of hurdles that must be overcome. Whether for efficiency, numerical stability, or generalization of methods for the researcher's end use, these sometimes intractable problems consume a large portion of a project's total completion time. This particular research was no exception. The following discussion is an attempt to classify and document the issues encountered with the hope that future work can benefit from this experience.

During the linear system solve, $Dx = b$, where each node is updated explicitly via an LU decomposition with partial pivoting and a back substitution, there is an opportunity for a significant savings in the hybrid chemistry solver. Each node has a matrix D that consists of the partial derivative terms $\frac{\partial Q}{\partial q}$ and the temporal term. Depending on whether a cell is switched to equilibrium or finite-rate, this matrix will be of size 5x5 or $(NS + 4) \times (NS + 4)$, respectively. With a combined LU factorization and back substitution costing $\sim \frac{n^3}{3} + n^2$ multiplications and a similar number of additions, it is easy to see how not taking advantage of the zero structure in equilibrium cells could be an expensive mistake. Memory is allocated on a global basis with an assumption that each cell could require the full finite-rate space if the

switching happens in that way. The memory array is indexed assuming each cell uses the full amount but the LU decomposition call is made with the correct matrix size for each cell. While this may not seem like such a large computational savings, remember that the goal of developing this method is to allow a very large number of species to be modeled at once. It is worth noting that this method does use more memory than is explicitly required but the time savings was deemed more important.

A similar issue occurs with the right hand side of the linear system. One must make the decision to constantly make memory allocations and deallocations for the residual vector; or, index as if each cell uses the maximum amount of memory. Due to the shifting operation outlined in Section 2.1.3, the full rank vector must be available to all control volumes regardless of the current switched state it is in. This makes the decision a simple one; however, it is worth noting that this is suboptimal with regard to memory usage. This method does, however, make the flux scatter operation extremely efficient and eliminates any interpretation calls when placing flux components within the residual vector.

The equilibrium solver “black box” was the source of many issues during the development of the hybrid method. One of the most troubling issues encountered was the inability to converge a particular set of reactions quickly to within a required tolerance on an equilibrium state. It was discovered through some considerable amount of analysis and observation that certain reactions, due to their disparate reaction rates, inhibit solution progress in the early iterations of an equilibrium solve. It seems that by excluding these reactions from the first few iterations and later including them, convergence rates can be drastically enhanced. The solutions obtained were checked with the baseline equilibrium states without the “freezing” and it was confirmed that the solutions were identical to within machine tolerance.

Another method of enhancing convergence of the “black box” solver was to provide good initial guesses. Here, the current state of a control volume was provided as an initial guess. Cox also outlines a projection method based on the update vector Δq that he found effective but it was not investigated here [7]. It was observed that the

limiters given in his dissertation are quite aggressive and turning them off after a set number of iterations greatly reduced the cost associated with achieving an equilibrium state. This should not be done without care, as an unlimited update can very easily diverge should the system be poorly conditioned.

Chemical Jacobian freezing was also investigated such that during an equilibrium solve, the Jacobian is only updated after a given number of iterations. This is common practice in CFD when dealing with implicit methods and was proven to indeed reduce the large cost associated with the Jacobian computation. However, the results differed very slightly and when coupled with the fluid solver, the stability of the entire method suffered. The source of this instability was not determined but it is thought to be related to the numerical “noise” introduced by the freezing operation.

The expense of the equilibrium solver, Equations 2.10 - 2.35, was a driving factor in this research. A large number of iterations in solving for an equilibrium composition can completely destroy any speedup achieved by switching a cell from finite-rate. Therefore, it was of the utmost importance to the viability of this research that the solver converged quickly and reliably for a large variety of initial states. This is no small task considering the large amount of nonlinearity inherent to chemical reaction systems. The equilibrium solver is called once per control volume per time step and convergence for the cell was typically reached in under 10 iterations.

Chapter 4

Numerical Verification

For verification of the hybrid method's suitability, a one-dimensional shock tube problem was chosen. The dramatic changes at the shock front made the problem challenging enough to test the stability of the methods implemented, yet the problem was small enough to be run with a varying set of flow solvers.

4.1 Problem Setup

A 5 species air model with the species N_2 , O_2 , NO , N , and O was used. It includes the 6 reactions in which they are involved [26]. These reactions are

where M is a variety of species that can act as catalysts. This is different for every reaction and it is possible for reacting species to act as catalysts in reactions in which they may or may not participate.

The particular shock tube initial conditions that were used in this study are as follows

$$q = \begin{bmatrix} \rho \\ u \\ v \\ w \\ P \end{bmatrix} = \begin{bmatrix} 2.122E - 01 \frac{kg}{m^3} \\ 0.000E + 00 \frac{m}{s} \\ 0.000E + 00 \frac{m}{s} \\ 0.000E + 00 \frac{m}{s} \\ 1.000E + 06 Pa \end{bmatrix}_L \longleftrightarrow \begin{bmatrix} 1.153E + 00 \frac{kg}{m^3} \\ 0.000E + 00 \frac{m}{s} \\ 0.000E + 00 \frac{m}{s} \\ 0.000E + 00 \frac{m}{s} \\ 1.000E + 05 Pa \end{bmatrix}_R$$

which corresponds to an equilibrium state calculation based on the above pressure and a temperature of 9000K on the left and 300K on the right. The Burcat tables [24] were used in initial investigations but the lack of curve fit data above 6000K necessitated the transition to Gupta multi-temperature curve fits in the NASA 7 coefficient form [25]. The initial mass fractions were set to standard atmosphere with 79% N_2 and 21% O_2 by mass. The initial species densities for the finite-rate and hybrid solvers were also obtained from the equilibrium solver.

These initial conditions were utilized consistently across all regimes, and the flow variables were nondimensionalized by the left-hand state in the CFD solver. Boundary conditions were Dirichlet and set to be consistent with the initial states. This was possible because the solution was not allowed to evolve to the point in time where the shock wave would begin to interfere with the mesh extents.

4.2 Refinement Study

A refinement study was performed in order to aid in determining the appropriate resolution on which to run this particular shock tube problem. The perfect gas compressible Euler solver was used for this test with second-order spatial flux

discretization and first-order time accuracy. A variation between 200 and 800 points was used with a run performed at every 200 point interval. These and all subsequent runs were performed on 1D grids with uniform spacing.

Figure 4.1: Refinement study: perfect gas compressible Euler equations - density plot

It can be seen from Figure 4.1 that 200 point case lacks the resolution to adequately capture the density magnitude. The 600 and 800 point resolutions are almost indistinguishable.

Figure 4.2: Refinement study: perfect gas compressible Euler equations - velocity plot

Figure 4.3: Refinement study: perfect gas compressible Euler equations - pressure plot

In Figure 4.2 significant oscillations are seen in both the 200 and 400 grid point resolutions. All solutions exhibit a non-physical blip at the shock front. This is most likely due to the higher-order spatial resolution not adequately damping numerical perturbances.

After examining the plots for density, velocity, and pressure it was decided that a 600 point grid provided appropriate shock capturing while still allowing for a reasonable amount of computational expense. It is worth noting that the time steps required to run the chemistry cases explicitly were on the order of 1×10^{-6} . The total run of 300,000 time steps allowed for the simulation of approximately 0.3 time units.

4.3 Verification

Several runs were made given the above initial conditions with each of the available solvers. This set includes the pure finite-rate solver, the pure equilibrium solver, and the perfect gas compressible Euler solver. Considering that each of the solvers operate with separate native variables, a great amount of care was taken to ensure that the runs were equivalent, except for the included type of chemistry. The cases run the spectrum from no chemistry (perfect gas Euler) to full detailed chemistry (finite-rate Euler). Again, they are plotted in comparison to a perfect gas exact Riemann solver solution.

It is worth noting that given the chosen resolution of the grid, the second-order spatial resolution, the high gradients near the shock front, and the immediate chemistry changes, the pure equilibrium solver was not run at the full resolution of 600 points. This solver was stable at lower grid resolutions, but the abrupt changes that occur at the shock front, in combination with the small control volumes, led to spurious oscillations that were of high frequency and magnitude. Thus, the equilibrium solver data shown is run on a 300 point grid or half the resolution of the other cases. This provided necessary numerical dissipation which reduced the magnitude of the flow features present at the shock front. The hybrid solver indicates that the Damköhler numbers in the region of the shock front were extremely small, implying that equilibrium is in fact a very poor assumption for those regions. The equilibrium solver results are shown simply for comparison with the other solvers as a demonstration of the drastically different character which can be exhibited simply by modifying the chemistry methods used.

In Figure 4.4 it is obvious that the addition of chemistry drastically changes the magnitude of the density spike. The hybrid and finite-rate solvers show a nearly identical prediction of this density increase, which indicates that the hybrid solver is accurately tracking the shock front. In fact, for all the plots shown, it is quite difficult to discern any difference between the two. The equilibrium chemistry solver shows a

Figure 4.4: Comparison: density comparison plot of four regimes

considerably higher density increase at the shock front. This is due to the assumption of infinitely fast chemistry that must be made with solvers of that type. This density spike is most likely an effect of the pressure spike caused at the shock front, by allowing more diatomic gas reassociation than would be physically possible, and is felt via the ideal gas equation of state. Here, it can be seen that equilibrium chemistry is likely a poor decision for modeling a high-temperature shock tube problem.

In Figures 4.5 and 4.6, the difference made by adding chemical terms to a perfect gas compressible flow solver in high-temperature air flows is again evident. As shown in Figures 4.8 and 4.9, the pure equilibrium solver predicts significantly more of the species *NO* and less monatomic *N*. The presence of additional *NO* suggests that *N* and *O* are combining in preference to remaining monatomic. This is an effect of allowing the chemical composition to settle to the lowest free energy state. The slight

Figure 4.5: Comparison: velocity comparison plot of four regimes

dip in velocity immediately at the shock front shown in the hybrid and finite-rate data is thought to be caused by the reaction front lagging the shock front slightly. This observation corresponds well to what is seen in the mass fraction plots in Figures 4.7, 4.8, and 4.9. The velocity spike shown in the equilibrium velocity data is an artifact of the “fast” chemistry assumption. When compared with the finite-rate data, the prediction of the spike corresponds precisely with the region in which the finite-rate regime predicts a sag. The lack of magnitude in the rest of the equilibrium shock region data is thought to be an artifact of the same assumptions.

Figure 4.6: Comparison: pressure comparison plot of four regimes

The pressure relaxation wave for finite-rate chemistry in Figure 4.6 is slightly slower than the perfect gas compressible solver. In contrast, the equilibrium solver shows a faster relaxation wave than predicted by either the perfect gas exact Riemann or perfect gas compressible Euler solvers. This is to be expected as chemistry is allowed to progress based on available energy in a control volume with no concept of rate of progression taken into account.

Figure 4.7: Species mass fractions N_2 and O_2 comparison

Figure 4.8: Species mass fractions N and O comparison

Figure 4.9: Species mass fractions *NO* comparison

4.4 Results of Hybrid Regime

The hybrid regime performed almost exactly as expected. The region that was switched to finite-rate begins in the few cells directly adjacent to the shock front, and as the solution evolves, the finite-rate solver is switched on in progressively more cells, tracking the shock front. The results using this method were indistinguishable from the results using a full finite-rate flow solver. This indicates that the hybrid method does an adequate job of detecting regions where the expense of a higher resolution flux might be warranted.

Several values for the switching cutoff were used in an attempt to determine if a poor choice could cause the solution to show incorrect characteristics. Eppard and Godfrey showed that the limiting of the Damköhler number to accelerate convergence by modifying reaction rate constants in equilibrium regions can be implemented with great success [27]. They found that for Damköhler number limiting between 3.0 and 10.0, acceptable accuracy is achieved. Their work suggests that at Damköhler numbers below 3.0, finite-rate chemistry effects begin to dominate. Though the definition of Damköhler number they used is not exactly the same as that used here, they are very similar. The cutoff region around 3.0 was thus investigated. However, despite great effort, a small enough whole number was not found that affected the solution character. The Damköhler numbers in the shock region were found to be on the order $1.0e^{-10}$ and smaller. It is expected that this behavior is specific only to this 1D shock tube case and that other flows will indeed show differences based on the cutoff point chosen.

Figure 4.10 shows the percentage of cells that are switched to equilibrium as the shock wave progresses. The area under this curve represents the fraction of time savings that would be realized for this particular case. The maximum time savings would correspond to all cells being in an equilibrium state, while there would be no cost benefit to the hybrid scheme if all cells in the flow field were running as finite-rate. As would be expected, the solver starts with 100% of cells being set as equilibrium

and as the solution progresses, more and more cells are switched to finite-rate as determined by the local Damköhler number. For Figure 4.10 and all the hybrid cases shown here, the Damköhler cutoff was set to one.

Figure 4.10: Percentage of cells switched from equilibrium as solution progresses

4.5 Performance

The performance of each of the chemically active solvers was measured to provide an accurate representation of the cost associated with each technique. A benchmark of 100 time steps on a 600 point grid was used. The same 5 species air model was run for each of the three solvers. When run implicitly, the number of SGS cycles was varied to give some sense of the cost associated with each additional relaxation sweep. The fastest is marked and the percent difference from the leader is indicated for each additional solver. The hybrid solver was not extended to being implicit and was thus not run for those timing cases.

Table 4.1: Performance comparison - Implicit solve (seconds)

Regime	5 SGS	%diff	8 SGS	%diff	11 SGS	%diff
Equil	3.225	fastest	3.276	fastest	3.355	fastest
F. Rate	3.719	15.4	3.883	18.5	4.071	21.3
Hybrid	N/A	N/A	N/A	N/A	N/A	N/A

It can be seen from Table 4.1 that the fastest solver is equilibrium, as expected. This solver was faster by approximately 15% for five SGS sweeps. It was observed that for each additional SGS sweep added, the equilibrium solver increases its lead by approximately one percentage point. It must be remembered that these results are for a five species air formulation and the cost associated with the SGS solve is expected to grow as $\sim (NS + 4)^2$. Hence, the computational savings are amplified by the complexity of the chemistry system.

Table 4.2: Performance comparison - Explicit solve (seconds)

Regime	Explicit	%diff
Equil	2.573	290.0
F. Rate	0.659	fastest
Hybrid	2.547	287.0

In Table 4.2, the finite-rate solver was faster by a very large margin. This is due to the additional expense of constructing and solving the linear system associated

with equilibrium chemistry. In this test, the additional cost of SGS iterations were not incurred and this made the equilibrium solver much more expensive in terms of computation time. Since the case here had all cells set to equilibrium to start with, and the shock propagation had not yet forced finite-rate solving in very many cells, the hybrid solver was approximately as expensive as the equilibrium solver.

Looking at Tables 4.1 and 4.2, some conclusions can be made about the cost benefits of using a hybrid solver. If running in an explicit manner, it is not prudent to incur the cost of equilibrium cells. Since the primary motivation for equilibrium chemistry is reducing the cost of a linear solve by handling chemistry outside of SGS iterations, this is to be expected. However, all indicators suggest that an implicit hybrid solver would fall somewhere between equilibrium and finite-rate solvers in terms of expense, while maintaining some higher degree of physical accuracy.

Chapter 5

Conclusions

The goal of this research was to develop a method that incorporated the cost savings of equilibrium chemistry and the full resolution of finite-rate chemistry. A method was developed that allows runtime switching of a control volume between equilibrium and finite-rate based on a local Damköhler number.

This hybrid solver was used to model a standard 1D shock tube problem and the results were compared with a full range of other compressible solvers. It was found that the accuracy of the new method rivaled that of the full finite-rate solver. In fact, for this problem, the solutions were indistinguishable. This indicates that the hybrid method does a relatively good job of identifying areas in which the solution would benefit from the added expense of a finite-rate regime.

According to the test timings of both the equilibrium and finite-rate solvers when run implicitly, it is expected that the hybrid method would indeed exhibit cost savings if it were extended to being fully implicit. This method is readily implemented in existing explicit flow solvers with chemistry capability but significant hurdles related to differently sized flux Jacobians must be crossed to make the method compatible with existing implicit solution methods.

Future Work

The hybrid chemistry solver was intended as a proof of concept for future extension using implicit methods. In this work, only the explicit time marching scheme has been implemented. Since explicit methods are limited by the CFL condition, the time steps that must be used are quite small. In order for the full benefit of this method to be felt, it must be extended by utilizing an implicit scheme. The major issues involved relate to the calculation of the flow Jacobians and the solution of a blocked linear system in which the block sizes vary. It is the opinion of this researcher that these limitations can be overcome using an operation similar to multigrid's prolongation and restriction. However, the structure of the system will dictate that Jacobians either be modified as necessary during an SGS sweep, which is computationally expensive, or that Jacobians of both ranks be stored for every node lying on a hybrid interface. It remains to be seen which of these would be more desirable.

Bibliography

Bibliography

- [1] A. Ern, V. Giovangigli, and M. D. Smooke, “Numerical study of a three-dimensional chemical vapor deposition reactor with detailed chemistry,” *Journal of Computational Physics*, vol. 126, pp. 21–39, 1996. 1
- [2] M. Cazalens, M. Rullaud, and J. P. Frenillot, “Computational methodology for carbon monoxide emission for aeroengine combustor design,” *Journal of Propulsion and Power*, vol. 24, pp. 779–787, 2008. 1, 2
- [3] D. Knight, “Rto wg 10: Test cases for cfd validation of hypersonic flight,” in *40th AIAA Aerospace Sciences Meeting and Exhibit*, 2002. 1
- [4] S. J. Henderson and J. A. Menart, “Adding chemical equilibrium to a parallel, unstructured euler/navier-stokes flow solver,” in *36th AIAA Fluid Dynamics Conference and Exhibit*, 2006. 1
- [5] L. M. Frohn, J. H. Christensen, and J. Brandt, “Development of a high-resolution nested air pollution model,” *Journal of Computational Physics*, vol. 179, pp. 68–94, 2002. 2
- [6] B. D’Anna, S. Langer, E. Ljungstrom, C. J. Nielsen, and M. Ullerstam, “Rate coefficients and arrhenius parameters for the reaction of the no3 radical with acetaldehyde and acetaldehyde-1d,” *Physical Chemistry Chemical Physics*, vol. 3, pp. 1631–1637, 2001. 2

- [7] C. F. Cox, *An Efficient Solver for Flows in Local Chemical Equilibrium*. PhD thesis, Mississippi State University, 1992. 2, 9, 29, 32
- [8] R. Rogers and W. Chinitz, “Using a global hydrogen-air combustion model in turbulent reacting flow calculations,” *AIAA Journal*, vol. 21, pp. 586–592, 1983. 2
- [9] J. C. Keck, “Rate-controlled constrained-equilibrium theory of chemical reactions in complex systems,” *Progress in Energy and Combustion Science*, vol. 16, pp. 125–154, 1990. 2
- [10] V. Yousefian, “Rate-controlled constrained-equilibrium thermochemistry algorithm for complex reacting systems,” *Combustion and Flame*, vol. 115, pp. 66–80, 1998. 2
- [11] S. Rao, C. J. Rutland, and S. B. Fiveland, “Reduced chemistry modeling - a new method for identifying reactions in partial equilibrium,” in *16th AIAA Computational Fluid Dynamics Conference*, 2003. 2
- [12] G. R. Inger, “Scaling nonequilibrium-reacting flows: The legacy of gerhard damkohler,” in *34th AIAA Thermophysics Conference*, 2000. 3, 16
- [13] R. J. Kee, M. E. Coltrin, and P. Glarborg, *Chemically Reacting Flow : Theory and Practice*. Wiley-Interscience, 2003. 7, 29
- [14] P. L. Roe, “Characteristic-based schemes for the euler equations,” *Annual Review of Fluid Mechanics*, vol. 18, pp. 337–365, 1986. 20
- [15] A. Harten, P. D. Lax, and B. van Leer, “On upstream differencing and godunov-type schemes for hyperbolic conservation laws,” *SIAM Review*, vol. 25, pp. 35–61, 1983. 20
- [16] J. Buvaeswari, *Accurate Upwind Procedures for Real-Gas Flows in Equilibrium and Non-Equilibrium*. PhD thesis, Indian Institute of Science, 1998. 21

- [17] E. F. Toro, *Riemann Solvers and Numerical Methods for Fluid Dynamics - A Practical Introduction*. Springer - Verlag, 1999. 21
- [18] D. G. Hyams, *An Investigation of Parallel Implicit Solution Algorithms for Incompressible Flows on Unstructured Topologies*. PhD thesis, Mississippi State University, May 2000. 23, 25
- [19] J. Blazek, *Computational Fluid Dynamics: Principles and Applications*. Elsevier Science Ltd., 2001. 23
- [20] B. van Leer, “Towards the ultimate conservative difference scheme v. a second-order sequel to godunov’s method,” *Journal of Computational Physics*, vol. 135, pp. 229–248, 1997. 23
- [21] C. Berthon, “Robustness of muscl schemes for 2d unstructured meshes,” *Journal of Computational Physics*, vol. 218, pp. 495–509, 2006. 23
- [22] M. Berger and M. J. Aftosmis, “Analysis of slope limiters on irregular grids,” in *43rd AIAA Aerospace Sciences Meeting and Exhibit*, 2005. 24
- [23] R. M. Beam and R. F. Warming, “An implicit factored scheme for the compressible navier-stokes equations,” *AIAA Journal*, vol. 16, pp. 393–402, 1978. 25
- [24] A. Burcat and B. Ruscic, “Third millenium ideal gas and condensed phase thermochemical database for combustion with updates from active thermochemical tables,” tech. rep., Argonne National Laboratory, 2005. 26, 35
- [25] R. N. Gupta, K.-P. Lee, R. A. Thompson, and J. M. Yos, “Calculations and curve fits of themodynamic and transport properties for equilibrium air to 30 000 k,” tech. rep., NASA, 1991. 26, 35

- [26] R. N. Gupta, J. M. Yos, R. A. Thompson, and K.-P. Lee, “A review of reaction rates and thermodynamic and transport properties for an 11-species air model for chemical and thermal nonequilibrium calculations to 30 000 k,” tech. rep., NASA, 1990. 34
- [27] W. M. Eppard and A. G. Godfrey, “Damkohler limiting of chemical time scales,” in *43rd AIAA Aerospace Sciences Meeting and Exhibit*, 2005. 45

Appendix

Appendix A

Definition of Mapping Jacobians

This appendix describes the process by which the mapping Jacobians are derived. These terms appear due to the choice of storing nonconservative variables in lieu of maintaining the conservative variables that are required for a fully conservative flux formulation. It should be noted that the derivatives shown here are strictly for use with curvefit thermo-chemistry but alternatives for use with vibrational chemistry are derived in a similar manner.

A.1 For Finite-rate Regime

The mapping Jacobian for the finite-rate regime is written as

$$\frac{\partial Q}{\partial q} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 1 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & \ddots & 0 & 0 & \dots & 0 \\ u & u & \dots & \rho & 0 & 0 & 0 \\ v & v & \dots & 0 & \rho & 0 & 0 \\ w & w & \dots & 0 & 0 & \rho & 0 \\ \frac{\partial \rho e_t}{\partial \rho_1} & \frac{\partial \rho e_t}{\partial \rho_2} & \dots & \rho u & \rho v & \rho w & \frac{\partial \rho e_t}{\partial P} \end{bmatrix} \quad (\text{A.1})$$

The total energy term can be expressed as

$$\rho e_t = \sum_{j=1}^{NS} \rho_j e_j + \frac{1}{2} \sum_{j=1}^{NS} \rho_j \|\vec{U}\|^2 \quad (\text{A.2})$$

where $\|\vec{U}\|$ is the magnitude of the velocity vector. The partial derivative of total energy with respect to a species partial density is

$$\frac{\partial \rho e_t}{\partial \rho_i} = e_j + \sum_{j=1}^{NS} \rho_j \frac{\partial e_j}{\partial \rho_i} + \frac{1}{2} \|\vec{U}\|^2 \quad (\text{A.3})$$

where the specific energy of a given species is

$$e_j = h_j - \frac{P_j}{\rho_j} = h_j - R_j T \quad (\text{A.4})$$

Substituting the specific enthalpy relation from Equation 2.68 into the above equation we get

$$e_j = R_j T \left[(a_1 - 1) + \frac{a_2 T}{2} + \frac{a_3 T^2}{3} + \frac{a_4 T^3}{4} + \frac{a_5 T^4}{5} + \frac{a_6}{T} \right] \quad (\text{A.5})$$

Remembering the ideal gas law

$$T = \frac{P}{\sum_{i=1}^{NS} \rho_i R_i} \quad (\text{A.6})$$

and defining

$$\tilde{R} = \sum_{i=1}^{NS} \rho_i R_i = \frac{P}{T} \quad (\text{A.7})$$

we can substitute to arrive at

$$e_j = R_j \left[(a_1 - 1) \frac{P}{\tilde{R}} + \frac{a_2 P^2}{2 \tilde{R}^2} + \frac{a_3 P^3}{3 \tilde{R}^3} + \frac{a_4 P^4}{4 \tilde{R}^4} + \frac{a_5 P^5}{5 \tilde{R}^5} + a_6 \right] \quad (\text{A.8})$$

Then, taking the partial derivative with respect to a species i 's partial density, we can derive

$$\begin{aligned} \frac{\partial e_j}{\partial \rho_i} &= R_j \left\{ \frac{\partial}{\partial \rho_i} \left[(a_1 - 1) \frac{P}{\tilde{R}} + \frac{a_2 P^2}{2 \tilde{R}^2} + \frac{a_3 P^3}{3 \tilde{R}^3} + \frac{a_4 P^4}{4 \tilde{R}^4} + \frac{a_5 P^5}{5 \tilde{R}^5} + a_6 \right] \right\} \quad (\text{A.9}) \\ &= R_j \left[\frac{\partial \tilde{R}}{\partial \rho_i} \right] \left[-(a_1 - 1) \frac{P}{\tilde{R}^2} - 2 \frac{a_2 P^2}{2 \tilde{R}^3} - 3 \frac{a_3 P^3}{3 \tilde{R}^4} - 4 \frac{a_4 P^4}{4 \tilde{R}^5} - 5 \frac{a_5 P^5}{5 \tilde{R}^6} \right] \\ &= -R_j R_i \left[(a_1 - 1) \frac{P}{\tilde{R}^2} + \frac{a_2 P^2}{\tilde{R}^3} + \frac{a_3 P^3}{\tilde{R}^4} + \frac{a_4 P^4}{\tilde{R}^5} + \frac{a_5 P^5}{\tilde{R}^6} \right] \\ &= -\frac{R_j R_i}{\tilde{R}} [(a_1 - 1)T + a_2 T^2 + a_3 T^3 + a_4 T^4 + a_5 T^5] \end{aligned}$$

Starting from Equation A.2, we can derive the partial derivative of total density with respect to pressure.

$$\frac{\partial \rho e_t}{\partial P} = \sum_{j=1}^{NS} \rho_j \frac{\partial e_j}{\partial P} \quad (\text{A.10})$$

Taking the partial derivative of Equation A.8 with respect to pressure we have

$$\frac{\partial e_j}{\partial P} = R_j \left\{ \frac{\partial}{\partial P} \left[(a_1 - 1) \frac{P}{\tilde{R}} + \frac{a_2 P^2}{2\tilde{R}^2} + \frac{a_3 P^3}{3\tilde{R}^3} + \frac{a_4 P^4}{4\tilde{R}^4} + \frac{a_5 P^5}{5\tilde{R}^5} + a_6 \right] \right\} \quad (\text{A.11})$$

$$\begin{aligned} &= R_j \left[\frac{(a_1 - 1)}{\tilde{R}} + 2 \frac{a_2 P}{2\tilde{R}^2} + 3 \frac{a_3 P^2}{3\tilde{R}^3} + 4 \frac{a_4 P^3}{4\tilde{R}^4} + 5 \frac{a_5 P^4}{5\tilde{R}^5} \right] \\ &= \frac{R_j}{\tilde{R}} [(a_1 - 1) + a_2 T + a_3 T^2 + a_4 T^3 + a_5 T^4] \end{aligned} \quad (\text{A.12})$$

A.2 For Equilibrium Regime

The mapping Jacobian for the equilibrium regime is written as

$$\frac{\partial Q}{\partial q} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ u & \rho & 0 & 0 & 0 \\ v & 0 & \rho & 0 & 0 \\ w & 0 & 0 & \rho & 0 \\ \frac{\partial \rho e_t}{\partial \rho} & \rho u & \rho v & \rho w & \frac{\partial \rho e_t}{\partial P} \end{bmatrix} \quad (\text{A.13})$$

The total energy term expressed in terms of total density is

$$\rho e_t = \sum_{i=1}^{NS} [\rho Y_i (h_i - R_i T)] + \frac{1}{2} \rho \|\vec{U}\|^2 \quad (\text{A.14})$$

Again, substituting the enthalpy relation from Equation 2.68, we can derive the partial derivative of total energy with respect to total density

$$\begin{aligned}
\frac{\partial \rho e_t}{\partial \rho} &= \frac{\partial}{\partial \rho} \sum_{i=1}^{NS} \left[\rho Y_i R_i \left\{ (a_1 - 1) \frac{P}{\rho R} + \frac{a_2 \left(\frac{P}{\rho R} \right)^2}{2} + \frac{a_3 \left(\frac{P}{\rho R} \right)^3}{3} + \frac{a_4 \left(\frac{P}{\rho R} \right)^4}{4} \right. \right. \\
&\quad \left. \left. + \frac{a_5 \left(\frac{P}{\rho R} \right)^5}{5} + a_6 \right\} \right] + \frac{1}{2} \|\vec{U}\|^2 \quad (\text{A.15}) \\
&= \sum_{i=1}^{NS} \left[Y_i R_i \left\{ -\frac{1}{2} a_2 \left(\frac{P}{\rho R} \right)^2 - \frac{2}{3} a_3 \left(\frac{P}{\rho R} \right)^3 - \frac{3}{4} a_4 \left(\frac{P}{\rho R} \right)^4 \right. \right. \\
&\quad \left. \left. - \frac{4}{5} a_5 \left(\frac{P}{\rho R} \right)^5 \right\} \right] + \frac{1}{2} \|\vec{U}\|^2 \\
&= \sum_{i=1}^{NS} \left[Y_i R_i \left\{ -\frac{1}{2} a_2 T^2 - \frac{2}{3} a_3 T^3 - \frac{3}{4} a_4 T^4 - \frac{4}{5} a_5 T^5 \right\} \right] + \frac{1}{2} \|\vec{U}\|^2
\end{aligned}$$

It should be noted that here Y_i is not taken to be a function of ρ since it is assumed that the mass fractions will not be affected by a density change. Strictly speaking, this may not be the case as the chemical composition may indeed vary depending on what terms are specified as fixed.

Vita

Nicholas Graham Currier was born in Merced, California, on July 27, 1985. He graduated Van Buren County High School in 2004 after beginning full-time studies at Tennessee Technological University during his senior year in 2003. During his time at TTU, he was a member of Pi Tau Sigma (mechanical engineering honor society), Mortar Board, ASME, and the university's Formula SAE team. He graduated Cum Laude with a Bachelor of Science degree in Mechanical Engineering in December 2007 and joined the University of Tennessee SimCenter as a graduate research assistant shortly thereafter. Nicholas plans to continue his research at the SimCenter and pursue a Ph.D.