

**A GENERALIZED MACCORMACK SCHEME FOR LOW MACH NUMBER,
CHEMICALLY-REACTING LARGE-EDDY SIMULATIONS**

A Dissertation
Presented to
The Academic Faculty

By

Timothy P. Gallagher

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Aerospace Engineering

Georgia Institute of Technology

August 2017

Copyright © Timothy P. Gallagher 2017

**A GENERALIZED MACCORMACK SCHEME FOR LOW MACH NUMBER,
CHEMICALLY-REACTING LARGE-EDDY SIMULATIONS**

Approved by:

Dr. Suresh Menon
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Stephen Ruffin
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Marilyn Smith
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Vaidyanathan Sankaran
Thermal and Fluid Sciences
United Technologies Research Center

Dr. Lakshmi Sankar
School of Aerospace Engineering
Georgia Institute of Technology

Dr. Yingjie Liu
School of Mathematics
Georgia Institute of Technology

Dr. Venkateswaran Sankaran
Aerospace Systems Directorate
Air Force Research Laboratory

Date Approved: June 2, 2017

ACKNOWLEDGEMENTS

No research is done in isolation and this work is certainly no exception. Every part of this thesis is the product of countless, often long, discussions and guidance and it builds on the decades of hard work done by my colleagues, predecessors, sponsors, advisor and instructors. I would like to start by thanking the Air Force Research Laboratory (AFRL) at Edwards Air Force Base and the United Technologies Research Center (UTRC) for providing the funding for this work. Dr. Venkateswaran Sankaran (AFRL) was instrumental in providing the history and the context of the techniques that provide the foundation for this work. Both he and Dr. Vaidyanathan Sankaran (UTRC) guided me through many of the technical challenges and provided an invaluable sounding board to develop solutions for the problems that arose. In addition, both provided personal and professional support and guidance over the past two years and I'm very happy to have them as mentors and collaborators.

This work is a small component of a much larger computational code, LESLIE. The ability to extend the operational range of the code to low Mach numbers with comparatively few lines of code is attributed to the efforts of the students, post-docs, and research engineers in the Computational Combustion Lab (CCL). The modern generation of LESLIE started in 2009. Over that time period, 17 unique authors have contributed over 500,000 lines of code and I'm proud to have been a part of that project. Within those years, there are only 2 hours during the 168 hours that make up an entire week where nobody has checked code into the repository. Everything that I have had the opportunity to work on is a product of the literally tireless efforts of those in the CCL. Together, we learned about software engineering, best practices, and techniques to ensure that everything the code can do is reproduced and accounted for. It was not always easy, we did not always agree, and we often got very worked up over things like tabs or spaces (and then how many of each), the classic vi vs. emacs editor wars, and which flavor of Linux is superior to the others. But

we persevered and the result is a robust computational framework that I hope enables the current and future researchers of the CCL to advance their own work with as little coding as possible.

Starting in the lab in 2006 as an undergraduate researcher means that I have had the opportunity to work with far too many people to remember them all. But I am fortunate to have made many great friends during that time and I am lucky enough to still have their support and friendship years after they have graduated and moved on. I got to share an office with Dr. Andrew Smith for most of my time in the CCL and I was able to bounce ideas and questions off of him the entire time. Dr. Matthieu Masquelet was another co-conspirator in our adventures to design the code and he was always ready to join us upstairs for conversations about what was working or not working. Together, the three of us took our adventures into the real world, which included some painfully hot and long bike rides, many pitchers of beer over the years, and Tour de France and Olympic curling watch-parties. More recently, much of the work for my masters degree, my doctorate, and the other projects I have been fortunate enough to be involved with has been in collaboration with Dr. Michel Akiki. Through it all, he has always been there to provide a technical or personal voice of reason when things would get frustrating, and I know that I can rely on him whenever I get stuck or need to vent.

I would also like to thank my thesis committee members for taking their time to serve, particularly Dr. Yingjie Liu from the School of Mathematics, who agreed to participate late in the process. My internal committee members, Dr. Stephen Ruffin, Dr. Marilyn Smith, and Dr. Lakshmi Sankar, have been among the best mentors for me in the department. Over the 13 years I have been at Tech, they have pushed and inspired me to learn and improve through their positions as instructors. All of my favorite classes were taught by one of these three professors, and I am fortunate to have been challenged by each of them.

Lastly, I have had incredible opportunities to learn and grow for the past 11 years because my advisor, Dr. Suresh Menon, took a chance on me when I was an undergrad. Students in

the CCL have access to a diverse group of colleagues and collaborators, a wide-ranging set of topics, and access to some of the best facilities in the world because of the work put in by Dr. Menon. He always gave me the opportunity to explore new ideas and techniques, which included letting me and my fellow students establish internal standards and practices on our own and to develop the computational framework that enabled this thesis. But, most importantly, he also was always there to pull us back to reality and keep us focused on the end goals, even though I didn't always come back easily. I've learned far more, about many more areas, than this thesis reflects and it is because Dr. Menon has given me opportunities and encouragement to branch out and take chances.

TABLE OF CONTENTS

Acknowledgments	iii
List of Tables	xi
List of Figures	xiii
Nomenclature	xix
Chapter 1: Thesis Overview	1
1.1 Motivation	2
1.2 Literature Review	7
1.2.1 Challenges at Low Mach Numbers	7
1.2.2 Numerical Methods for Low Mach Number Flows	8
1.2.3 Treatment of Stiff Chemical Kinetics	14
1.3 Objectives	18
1.4 Thesis outline	22
Chapter 2: Governing Equations	25
2.1 Non-reacting Navier-Stokes Equations	25
2.2 Reacting Navier-Stokes Equations	27
2.3 Equations of State	30

2.4	Transport Properties	32
2.5	Large-Eddy Simulation	34
2.5.1	Conservation of mass	37
2.5.2	Conservation of momentum	37
2.5.3	Conservation of energy	38
2.5.4	Conservation of subgrid kinetic energy	40
2.5.5	Conservation of partial densities	41
2.5.6	Chemistry-Turbulence Closure	42
2.6	Boundary Conditions	49
2.6.1	Characteristic Inflows and Outflows	50
2.6.2	Simple Inflows and Outflows	56
2.6.3	Walls	58
2.7	Preconditioning and Dual-Time	62
Chapter 3: Traditional MacCormack Scheme		73
3.1	1971 MacCormack Scheme	74
3.2	Artificial Dissipation	79
3.3	Low Mach Number Analysis	80
Chapter 4: Generalized MacCormack Scheme		85
4.1	Generalized MacCormack Scheme	86
4.1.1	Single-time integration	87
4.1.2	Low-Mach Number Formulation	88
4.1.3	Dual-time integration	90

4.2	Time Step Selection and Robustness	93
4.2.1	Local Time Stepping	94
4.2.2	Robust Preconditioning	95
4.2.3	Pressure Perturbation Tests	99
Chapter 5: Non-reacting Scheme Validation		107
5.1	Flow over Bump	108
5.2	Inviscid Vortex Convection	111
5.3	Lid-Driven Cavity	121
5.4	Decaying Isotropic Turbulence	132
5.5	Summary of results	138
Chapter 6: Extensions for Stiff Chemistry		141
6.1	Chemistry Integration in Single-Time Schemes	142
6.2	Fully-Coupled Chemistry in Dual-Time Schemes	147
6.3	Operator Splitting in Dual-Time Schemes	148
6.3.1	Physical-Time Splitting	148
6.3.2	Pseudo-Time Splitting	150
6.4	Consistent Splitting in Dual-Time Schemes	151
6.5	Summary	152
Chapter 7: Reacting Scheme Validation		155
7.1	Zero-Dimensional Ignition	156
7.2	One-Dimensional Ignition	160

7.3	Three-Dimensional Ignition	169
7.4	Summary of results	176
Chapter 8: Applications		179
8.1	Volvo Flygmotor	179
8.2	Cambridge Slot Burner	186
Chapter 9: Conclusions and Future Work		203
9.1	Summary	203
9.2	Future Work	212
Appendix A: Nonlinear Instabilities in the MacCormack Scheme		217
A.1	Model Equations	217
A.2	Vortex Convection	219
Appendix B: Computational Settings and Performance		223
References		246

LIST OF TABLES

2.1	Definitions for the common limiting values of the reference velocity used in Equation (2.113) to determine the preconditioning parameter.	67
3.1	Alternating difference direction pattern for MacCormack scheme [117]. . .	76
4.1	Definitions for the common limiting values of the reference velocity used in Equation (2.113) to determine the preconditioning parameter.	100
4.2	Pulse strength order of magnitude for the maximum stable pressure pulse measured in atmospheres for the various combinations of preconditioning definitions from Table 2.1	104
4.3	Comparison in maximum stable pulse strength between traditional preconditioning using definitions (A) and (D) from Table 4.1 and the new preconditioning definitions designed to enforce the same limits as the non-linear time step control from Lian et al. [60].	105
5.1	Convergence information for the lid-driven cavity for each combination of flow conditions, meshes and schemes.	127
7.1	Errors measured by the L_2 norm of the volume-integrated temperature in Case 1A as a function of time for each scheme and CFL number, relative to the results for Scheme C at $CFL_c = 1$. The error for the explicit scheme relative to this case is 4.11×10^{-6}	162
8.1	Species and Lewis numbers used with the Smooke-Giovangigli mechanism.	188
8.2	Sectors of the $(\hat{P}, \hat{Q}, \hat{R})$ space divided by the discriminant surfaces along with their topological characteristics [156] and physical interpretations.	192

B.1	Preconditioning settings for the flow over a bump case.	223
B.2	Preconditioning settings for the inviscid vortex case.	224
B.3	Preconditioning settings for the lid-driven cavity case.	225
B.4	Relative wall times for the lid-driven cavity case. All timing metrics are relative to the GMC-ST simulation on the same grid.	225
B.5	Preconditioning settings for the non-reacting isotropic turbulence case.	226
B.6	Preconditioning settings for the zero-dimensional ignition cases.	227
B.7	Relative wall times for the zero-dimensional ignition case. All timing metrics are relative to the GMC-ST simulation.	227
B.8	Preconditioning settings for the one-dimensional ignition cases.	229
B.9	Preconditioning settings for the three-dimensional ignition case.	230
B.10	Preconditioning settings for the Volvo case.	230
B.11	Preconditioning settings for the Cambridge slot burner cases.	231

LIST OF FIGURES

1.1	Illustration of the resolved and modeled regions of each computational technique in an example spectrum of turbulent kinetic energy.	3
1.2	The solution of a low Mach number flow over a circular bump using a compressible scheme exhibits undesirable oscillations and asymmetries in the velocity field, demonstrating the common quality issues that arise for many schemes at low Mach numbers.	5
1.3	Illustrations of several typical combustion configurations that contain low speed flow regions that often dictate the dynamics of the system.	6
2.1	Schematic of a laminar, premixed flame structure where δ_F is thermal thickness, and δ_r is the thickness of the reaction zone; thermal conductivity and species diffusion transport energy from the primary reaction zone into the unburned gases, creating the preheat zone.	43
2.2	Peters' modified Borghi diagram showing the regimes of premixed, turbulent combustion in log-log scale [45].	45
3.1	Pressure oscillations in a smooth flow using various forms of the non-linearity switch. Oscillations with the original switch are large enough to degrade solution quality at low Mach numbers. The details of the new switch are in Chapter 4 and Appendix A.	78
4.1	Convergence history of the GMC-PC for the lid-driven cavity at Reynolds number of 1000 and Mach number of 0.3 with/without alternating the difference directions.	93
4.2	Reference Mach number and the velocity perturbation response for increasing pressure perturbation strengths when using definitions (A), (D), and (G) in Table 2.1 for flows with different mean Mach numbers, M	97

4.3	Reference Mach number and the velocity perturbation response for increasing pressure perturbation strengths, comparing original preconditioner (black) and the preconditioner augmented with limits based on the non-linear time step control from Lian et al. [60] (blue) with different mean Mach numbers, M	103
4.4	Residual for the $M = 0.001$ case at the maximum stable pressure pulse with and without the $U_r = \alpha_{u2c}$ preconditioner definition.	105
5.1	Schematic of the domain for the inviscid flow through a channel with a bump.	109
5.2	Mesh used for the inviscid flow over a bump in a channel case containing 512×256 cells.	109
5.3	Steady-state solutions for the inviscid flow over a bump at $M_\infty = 0.001$, showing the degradation that occurs when compressible schemes (GMC-ST) are applied at low Mach numbers. Isolines of Mach number and velocity contours highlight the smoothness and symmetry of the solution using the preconditioned scheme, GMC-PC.	110
5.4	Computational domain for the inviscid vortex convection case.	112
5.5	Velocity profiles across the vortex convecting at $M = 0.5$, using the GMC-ST and the GMC-PC schemes. Reprinted with permission from Gallagher et al.[117].	113
5.6	Pressure profiles across the vortex convecting at $M = 0.5$, using the GMC-ST and the GMC-PC schemes.	114
5.7	Velocity profiles across the vortex convecting at $M = 0.001$ illustrating the limitations of applying preconditioning to the original MacCormack scheme (MC-PC) when run at a $CFL_c = 1$. The improved scheme (GMC-PC) shows little numerical dissipation under these conditions. Reprinted with permission from Gallagher et al.[117].	115
5.8	Velocity profiles across the vortex convecting at $M = 0.001$ for the GMC-PC scheme at two different physical CFL numbers compared to the GMC-ST and theoretical profiles. Both limiting cases of the GMC-PC scheme show improved dissipative properties relative to the GMC-ST scheme at these conditions.	116
5.9	Velocity profiles across the vortex convecting at $M = 0.001$, using the GMC-ST and GMC-PC schemes with various convergence criteria. Reprinted with permission from Gallagher et al.[117].	117

5.10	Convergence history for several physical time iterations of the vortex convection case at two physical CFL numbers. Discontinuities in the residual indicate the transition to the next physical-time iteration. Reprinted with permission from Gallagher et al.[117].	119
5.11	Comparison of the solutions using the GMC-ST and GMC-PC schemes, where the GMC-ST solution requires three times finer mesh in both directions for comparable accuracy to the GMC-PC solution on the original mesh. Reprinted with permission from Gallagher et al.[117].	120
5.12	Convergence histories for the GMC-ST and GMC-PC schemes on the finest grids for each configuration. Reprinted with permission from Gallagher et al.[117].	128
5.13	U- and V-velocity profiles for both Reynolds numbers at $M = 0.3$ indicating both the GMC-PC and GMC-ST schemes are capable of capturing the correct solution.	129
5.14	U- and V-velocity profiles for both Reynolds numbers at $M = 0.001$ demonstrating the inability of the GMC-ST scheme to capture the correct results. Reprinted with permission from Gallagher et al.[117].	130
5.15	Steady-state velocity contours in the lid-driven cavity for the $M = 0.001$ at $Re = 1000$ using the GMC-PC scheme (filled contour) and the GMC-ST scheme (lines). Contours are at 17 levels from $ u = 0$ (light) to $ u = u_{\text{wall}}$ (dark). Reprinted with permission from Gallagher et al.[117].	131
5.16	Error in velocity relative to reference data [123] for each mesh tested. The blue dashed line represents second-order accurate convergence. Reprinted and adapted with permission from Gallagher et al.[117].	131
5.17	Time variation of the vorticity fluctuations and Kolmogorov length scale in decaying isotropic turbulence using the GMC-ST and GMC-PC schemes. Reprinted with permission from Gallagher et al.[117].	133
5.18	Turbulent kinetic energy spectra at $\tau_t = 3.5$ for the GMC-ST and GMC-PC schemes. Reprinted with permission from Gallagher et al.[117].	134
5.19	Z-vorticity (normal to the plane) in an isotropic decaying turbulence field along the center plane for the GMC-PC and GMC-ST schemes at an eddy-turnover time $\tau_t = 3.5$. Reprinted with permission from Gallagher et al. [117].	136
5.20	Velocity correlations for the GMC-PC and GMC-ST schemes at an eddy-turnover time $\tau_t = 3.5$	137

7.1	Evolution of key species with time from the Cantera constant volume simulation	156
7.2	Eigenvalue and temperature evolution with time from the Cantera constant volume simulation. The magnitude of the real component of the eigenvalues is on the right axis where open circles represent negative real components and filled squares represent positive real components. Temperature is the dashed line on the right axis.	157
7.3	Number of pseudo-iterations per physical time step in the zero-dimensional reactor for each scheme, averaged over bins $10\ \mu\text{s}$ wide. Scheme B does not require any pseudo-iterations for this configuration and is omitted. Scheme C failed to converge for the two largest physical time steps.	159
7.4	Initial fields for the one-dimensional test cases with a temperature inhomogeneity (Case 1A) and with both a temperature and compositional inhomogeneity through the introduction of atomic H (Case 1B).	161
7.5	Volume-averaged temperature profiles for each scheme and CFL number for Case 1A. Scheme B at $\text{CFL}_c = 1000$ is unstable due to the instantaneous volumetric expansion at the large, physical time step while Scheme C at the same CFL number is unstable at the point of maximum heat release. Both are omitted.	166
7.6	Instantaneous OH mass fraction at various instants in time through the ignition process for Case 1A. Symbols are the same as in Figure 7.5	167
7.7	Temperature and OH profiles for each scheme and CFL number for Case 1B. Scheme C is unstable for all physical CFL numbers tested and is omitted. 168	
7.8	Front surface area (the isosurface of $\tilde{Y}_{\text{H}_2} = 8.5 \times 10^{-4}$) and pressure throughout the ignition process for the two schemes at each CFL number, compared to the 3D DNS data from Yu and Bai [142].	171
7.9	Contours of HO_2 for both schemes and each CFL number along the center plane at $t = 0.9\tau_0$ with the ignition front defined as the surface of $\tilde{Y}_{\text{H}_2} = 8.5 \times 10^{-4}$ indicated as black lines.	175
7.10	Isosurfaces of the ignition front defined as the surface of $\tilde{Y}_{\text{H}_2} = 8.5 \times 10^{-4}$ at $t = 0.9\tau$ for both schemes at $\text{CFL}_c = 1$, where the blue regions indicate the slow, deflagration front propagation ($s_d^* < 1.1S_L$) and the yellow regions indicate the rapid, spontaneous ignition front propagation ($s_d^* > 1.1S_L$). . .	176

8.1	Schematic of the computational domain used for the premixed, bluff-body stabilized flame simulations using the specifications of the Model Validation for Propulsion workshop [79].	180
8.2	Grid clustering around the bluff body and shear layers.	181
8.3	Temperature contours along the centerplane of the Volvo flygmotor using the GMC-PC scheme and the explicit scheme. Images are not at the same time instant.	182
8.4	Spanwise vorticity contours along the centerplane of the Volvo flygmotor using the GMC-PC scheme and the explicit scheme. Images are not at the same time instant.	183
8.5	Profiles along the centerline behind the bluff-body for the GMC-PC scheme compared to the explicit scheme. Results for the explicit scheme on several consecutively finer meshes provide context for the GMC-PC scheme's results.	184
8.6	Transverse profiles at the downstream locations $x/H \in \{0.375, 0.95, 1.53, 3.75, 9.4\}$ for the GMC-PC scheme, the explicit scheme, and experimental measurements. Results for the explicit scheme are on several consecutively finer meshes.	185
8.7	Schematic of the Cambridge slot burner.	194
8.8	Grid within the important regions of the domain, including the near-wall location where the flame will anchor.	195
8.9	Mean temperature profiles at two streamwise locations downstream from the burner exit for both the premixed and stratified conditions. Experimental and RANS data digitized from Darbyshire et al. [83]. The vertical dashed line marks the centerline of the rod.	196
8.10	Instantaneous spanwise vorticity along the center plane of the burner for both the premixed and stratified configurations.	197
8.11	Instantaneous contours of reactant-side and product-side radical species indicative of the leading and trailing edges of the flame brush for both configurations along the center plane.	198
8.12	Sectors representing different critical points in the topological analysis of the invariants of the velocity gradient tensor. Sectors are defined in Table 8.2. Yellow dashed lines represent the surfaces where the discriminant is zero.	199

8.13	Topological analysis of the velocity gradient tensor invariants for the pre-mixed case conditioned on $\tilde{T} = 800 \text{ K}, 1300 \text{ K}, 1800 \text{ K}$ from left to right. . .	200
8.14	Topological analysis of the velocity gradient tensor invariants for the stratified case conditioned on $\tilde{T} = 800 \text{ K}, 1300 \text{ K}, 1800 \text{ K}$ from left to right. . . .	201
A.1	Comparison of the schemes without the nonlinearity switch, with the original 1971 nonlinearity switch, and with the new nonlinearity switch given by Equation (4.1) on the linear and nonlinear model equations. Reprinted with permission from Gallagher et al.[117].	219
A.2	Velocity and pressure profiles across the vortex using the original, 1971 nonlinearity switch, the new switch proposed in this work (Equation (4.1)), and no nonlinearity switch for one turn around the domain. Reprinted with permission from Gallagher et al. [117].	221

NOMENCLATURE

A_i	Amplitude of the i -th mode of a vibrating wall
A	Linearized flux Jacobian for the Navier-Stokes system in the X-direction
A	Velocity gradient tensor
A'	Linearized flux Jacobian for the Navier-Stokes system in the direction normal to the boundary
A	Cell face area vector
B	Linearized flux Jacobian for the Navier-Stokes system in the Y-direction
B'	Linearized flux Jacobian for the Navier-Stokes system in the first tangential direction to the boundary
C_k	Constant used in the evaluation of the dynamically thickened flame model
C	Linearized flux Jacobian for the Navier-Stokes system in the Z-direction
C'	Linearized flux Jacobian for the Navier-Stokes system in the second tangential direction to the boundary
D_k	Diffusion coefficient for species k into the mixture
$D_{k^{sgs}}$	Dissipation of subgrid kinetic energy
\hat{D}	Non-dimensionalized discriminant of a tensor characteristic equation
D	Jacobian representing the derivative of the source term vector with respect to the primitive variable vector
D	Viscous and source terms for the Navier-Stokes system used when applying characteristic inflow and outflow boundary conditions
D''	Viscous, source, and transverse terms for the Navier-Stokes system used when applying characteristic inflow and outflow boundary conditions in the rotated system
E	Efficiency factor used in the dynamically thickened flame model
E	Total energy

F	Thickening factor used in the dynamically thickened flame model
F_R	Reacting inviscid flux tensor
F_{NR}	Non-reacting inviscid flux tensor
$\mathcal{G}(t, \mathbf{x})$	Spatial filter used to derive the LES equations
\mathbf{G}	Amplification matrix
\mathbf{G}_R	Reacting viscous flux tensor
\mathbf{G}_{NR}	Non-reacting viscous flux tensor
H	Bluff-body height
H	Total enthalpy
\mathbf{H}^{sgs}	Inviscid subgrid transport of enthalpy vector
\mathbf{J}	Species mass diffusion vector
K	Relaxation coefficient used for partially-reflective characteristic inflows and outflows
L_{11}	Integral length scale
\mathbf{L}_p	Selector matrix to extract the pressure field in the pressure-only artificial dissipation
\mathbf{L}	Vector of wave amplitudes used when applying characteristic inflow and outflow boundary conditions
\mathbf{L}_{PC}	Vector of wave amplitudes used when applying characteristic inflow and outflow boundary conditions in the preconditioned system
M	Mach number
M_t	Turbulent Mach number
N_k	Total number of species
N_m	Number of modes for a vibrating wall
N_s	Number of species used to define the flame sensor in the dynamically thickened flame model
P	First tensor invariant
$\Delta \mathbf{P}^{AD4}$	Fourth-order artificial dissipation term for the pressure field only
\hat{P}	Non-dimensionalized first tensor invariant

P	Variable vector corresponding to acoustic flux component
Q	Second tensor invariant
ΔQ	Incremental change in the primitive variable vector
\hat{Q}	Non-dimensionalized second tensor invariant
Q	Primitive variable set used in the pseudo-time integration of the preconditioned system of equations
Q_s	Vector of symmetrizing primitive variables
R	Gas constant of the mixture
R	Third tensor invariant
R_u	Universal gas constant
\hat{R}	Non-dimensionalized third tensor invariant
R	Residual vector containing the spatial fluxes and non-stiff source terms
S	Constant used in Sutherland's law
S	Sensor used to detect the flame location in the dynamically thickened flame
S^ϕ	Sensor based on variable ϕ used to determine when the second-order artificial dissipation is needed
S_F	Total cell face area
S_L	Laminar flame speed in a premixed flame
S_n	Cell face area normal to a wall
S_{ij}	Component of the rate-of-strain tensor
S	Eigenvector matrix
S_{PC}	Eigenvector matrix for the preconditioned system
S	Cell face area vector
S_I	Cell face area vector projected on the I computational coordinate
S_J	Cell face area vector projected on the J computational coordinate
S_K	Cell face area vector projected on the K computational coordinate
S_R	Reacting source term vector
S_{NR}	Non-reacting source term vector

\mathbf{S}_{NS}	Vector of the non-stiff components of the source term vector
\mathbf{S}_S	Vector of the stiff components of the source term vector
T	Fluid temperature
T_f	Final temperature of a reacting system
T_i	Initial system temperature
T_x	Cell face area tangential to the X-direction
T_y	Cell face area tangential to the Y-direction
T_z	Cell face area tangential to the Z-direction
T_{ref}	Reference temperature used in Sutherland's law
$\mathcal{T}_{k^{sgs}}$	Diffusion of subgrid kinetic energy
U	Contravariant velocity at the cell face
U_r	Reference velocity used to define the preconditioning parameter
$U_{r,A}$	Reference velocity multiplied by cell face area
\mathbf{U}'_{io}	Primitive vector set used to apply the inflow–outflow boundary conditions, rotated to align with the boundary
\mathbf{U}_w	Primitive variable set used to impose wall boundary conditions
\mathbf{U}_{io}	Primitive variable set used to impose inflow–outflow boundary conditions
V	Cell volume
V_∞	Reference inflow velocity
V_n	Velocity component normal to a wall
V'_n	Velocity component normal to a wall caused by vibrations
$V'_{effusion}$	Velocity of effusion prescribed through a wall
\mathbf{V}_k	Diffusion velocity vector of the k -th species
W	Mixture molecular weight
W_k	Molecular weight of species k
$\Delta \mathbf{W}$	Incremental change in the conservative variable vector
\mathbf{W}'	Conservative variable set rotated to be normal to the boundary

\mathbf{W}_R	Reacting conservative solution vector
\mathbf{W}_{NR}	Non-reacting conservative solution vector
X	Species mole fraction
Y	Species mass fraction
\mathbf{Y}_k^{sgs}	Subgrid mass flux of the species k
$b_{1..6,k}$	Polynomial curve fit coefficients used to evaluate species specific thermal conductivity for species k
c	Speed of sound
c'	Preconditioned acoustic velocity
c_ν	Coefficient used in the definition of the turbulent viscosity
c_p	Specific heat capacity at constant pressure for the mixture
c_v	Specific heat capacity at constant volume
$c_{1..6,k}$	Polynomial curve fit coefficients used to evaluate species specific viscosity for species k
c_ϵ	Coefficient in the definition of the subgrid kinetic energy dissipation model
$c_{p,k}$	Specific heat capacity at constant pressure for species k
$c_{v,k}$	Specific heat capacity at constant volume for species k
e	Internal energy
$e_{k,0}$	Sensible energy of species k at a reference temperature
f	Factor used to toggle between slip and no-slip wall boundary conditions
f_δ	Auxiliary function used to compute the efficiency factor in the dynamically thickened flame model
f_i	Frequency of the i -th mode of a vibrating wall
f_u	Auxiliary function used to compute the efficiency factor in the dynamically thickened flame model
f_{Re}	Auxiliary function used to compute the efficiency factor in the dynamically thickened flame model
f_{cut}	Cutoff frequency induced by the LES filter
g_i	One-dimensional spatial filter

h	Internal enthalpy of the mixture
h_k	Enthalpy of the k -th species
$h_{k,0}$	Sensible enthalpy of species k at a reference temperature
k	Wave number
k^{sgs}	Resolved subgrid turbulent kinetic energy
k_0	Wave number with peak energy content
k_{cut}	Cutoff wave number induced by the LES filter
l	Length scale used to determine the unsteady preconditioning reference velocity
l_t	Turbulent integral length scale
\hat{n}	Unit vector normal to the cell face
Δp	Pressure gradient across cell faces used to determine the preconditioning reference velocity
p	Pressure
p_∞	Reference pressure
p_g	Gauge pressure
$pd_{k^{sgs}}$	Pressure-dilation correlation for the subgrid kinetic energy
\mathbf{q}	Heat flux vector
\mathbf{q}_k^{sgs}	Subgrid heat transfer vector due to unresolved turbulent diffusion of species k
q	Heat flux at the wall boundary
s	Fluid entropy
s_d	Reaction front displacement speed
s_d^*	Density-weighted reaction front displacement speed
Δt	Physical time step increment
Δt_{cut}	Cutoff time step representing the smallest resolvable features in time
t	Physical time
\mathbf{u}	Velocity vector

u	Convective velocity
u	Velocity in the X-coordinate direction
u'	Preconditioned convective velocity
u'	Turbulent velocity fluctuation
u_1	Flow velocity normal to the boundary
u_A	Velocity projected on the cell face, multiplied by the cell face area
u_t	Flow velocity in the first tangential direction to the boundary
u'_{effusion}	Velocity component in the X-direction for effusion through a wall
u'_{vib}	Velocity component in the X-direction for a vibrating wall
u_{tt}	Flow velocity in the second tangential direction to the boundary
v	Velocity in the Y-coordinate direction
v'_{effusion}	Velocity component in the Y-direction for effusion through a wall
v'_{vib}	Velocity component in the Y-direction for a vibrating wall
\boldsymbol{w}	Variable vector for the diagonalized Euler equations
w	Velocity in the Z-coordinate direction
w'_{effusion}	Velocity component in the Z-direction for effusion through a wall
w'_{vib}	Velocity component in the Z-direction for a vibrating wall
Δx	Incremental change in the X-coordinate direction between cell faces
x_n	Coordinate normal to the boundary
x_t	Coordinate in the first tangential direction to the boundary
x_{tt}	Coordinate in the second tangential direction to the boundary
Δ_i	Grid spacing in the i -th direction
$\bar{\Delta}$	Size of the LES filter
$\Gamma(K, D)$	Function used to compute the efficiency factor in the dynamically thickened flame model
$\mathbf{\Gamma}$	Preconditioning matrix for the primitive variable set
$\mathbf{\Gamma}_C$	Preconditioning matrix for the conservative variable set

Γ'_{io}	Preconditioning matrix for the primitive variable set used for inflows and outflows, rotated to align with the boundary
Λ	Eigenvalue matrix
Λ_{PC}	Eigenvalue matrix for the preconditioned system
Ω	Domain used for integration
$\partial\Omega$	Surface of the domain used for integration
Π_{kj}	Function used in the definition of the mixture viscosity when calculated using Wilke's formula
Θ	Preconditioning parameter
α	Function used in the evaluation of the dynamically thickened flame model
α_v	Coefficient of thermal expansion
$\alpha_{1\dots7,k}$	Polynomial curve fit coefficients used to evaluate species heat capacity and sensible enthalpy
α_{p1}	Coefficient used to limit the pressure response to velocity perturbations
α_{u1}	Coefficient used to limit the velocity response to pressure perturbations
α_{u2}	Coefficient used to limit the velocity response to pressure perturbations
β	Constant used in the evaluation of the dynamically thickened flame model
δ_F	Thickness of the flame
δ_F^0	Unstrained thermal thickness of a premixed flame
δ_r	Thickness of the reaction zone in a flame
δ_{ij}	Kronecker delta
ϵ	Reference Mach number in the preconditioned system
$\epsilon^{(2)}$	Coefficient used for the Jameson–Schmidt–Turkel second-order artificial dissipation term
$\epsilon^{(4)}$	Coefficient used for the Jameson–Schmidt–Turkel fourth-order artificial dissipation term
ϵ_4	Artificial dissipation coefficient used in the pressure-only dissipation term
ϵ_{TF}	User-defined cutoff value used in the dynamically thickened flame sensor
η	Kolmogorov length scale

ϕ_{target}	Target value for a generic variable imposed at the partially-reflective and non-reflective characteristic inflows and outflows
γ	Ratio of specific heats
κ	Condition number of a system
κ	Small cutoff coefficient for ensuring the preconditioning reference velocity remains defined at stagnation points
$\kappa(T)$	Thermal diffusivity of the fluid
$\kappa^{(2)}$	User-defined dissipation coefficient for the second-order Jameson–Schmidt–Turbel artificial dissipation
$\kappa^{(4)}$	User-defined dissipation coefficient for the fourth-order Jameson–Schmidt–Turbel artificial dissipation
κ_k	Species specific thermal conductivity for species k
λ	Eigenvalue of a system
$\lambda(T)$	Bulk viscosity of the fluid
$\lambda_{PC,1,\dots,6+k}$	Eigenvalues of the preconditioned system
$\boldsymbol{\lambda}$	Vector of eigenvalues
$\mu(T)$	Viscosity of the fluid
μ_k	Species specific viscosity for species k
μ_{ref}	Reference viscosity used in Sutherland’s law
ν_t	Turbulent viscosity
$\dot{\omega}$	Chemical reaction rate for a species
ω	Frequency of the eigenmodes for the simple boundary conditions
ω	User-specified weight used to control the local time step smoothing
ω_{RMS}	Root-mean-square of vorticity fluctuations
ϕ	Equivalence ratio
ρ	Fluid density
σ_e	Spectral radius of the non-preconditioned system
σ_p	Spectral radius of the preconditioned system
σ^{sgs}	Subgrid heating due to viscous dissipation

$\Delta\tau$	Pseudo-time step increment
τ	Pseudo-time
τ	Viscous stress tensor
τ_0	Ignition delay time
τ_t	Eddy turnover time
τ_{xi}	Components of the viscous stress tensor in the X-coordinate direction
τ_{yi}	Components of the viscous stress tensor in the Y-coordinate direction
τ_{zi}	Components of the viscous stress tensor in the Z-coordinate direction
θ	Normalized temperature
θ_{effusion}	Angle of effusion prescribed through a wall
τ^{sgs}	Subgrid viscous stress tensor
θ_k^{sgs}	Subgrid species diffusion flux of the species k
CFL_{cut}	Cutoff CFL number for an ideal LES time step
CFL_c	CFL number based on the acoustic eigenvalue
CFL_{u+c}	CFL number based on the largest acoustic eigenvalue
$\text{CFL}_{u'}$	CFL number based on the initial turbulent velocity fluctuations
CFL_u	CFL number based on the convective eigenvalue
Re	Reynolds number
CFL	Courant–Friedrichs–Lewy stability criterion
Da	Damköhler number
Ka	Karlovitz number
Le_k	Species specific Lewis number
Pr	Prandtl number
Pr_t	Turbulent Prandtl number
Re_t	Turbulent Reynolds number
Re_Δ	Reynolds number at the effective filter size used in the dynamically thickened flame model

Sc_t	Turbulent Schmidt number
$(\cdot)_g$	Subscript indicating the value of a variable in the ghost cell used to impose wall boundary conditions
$(\cdot)_i$	Subscript indicating the value of a variable inside the flow, adjacent to a wall
$(\cdot)_w$	Subscript indicating the value of a variable on the wall
$\langle \cdot \rangle$	Mean value of a variable
$(\cdot)''$	Fluctuating component of a spatially filtered variable
(\cdot)	Perturbation of a variable in Fourier space
$(\bar{\cdot})$	Mean component of a spatially filtered variable
$(\tilde{\cdot})$	Farv�re averaged variable

SUMMARY

Chemically reacting flows contain a wide range of regimes with many velocity and time scales. The increasing access to computational resources enables higher-fidelity simulations of these flows. In order to take advantage of these capabilities, numerical schemes must be robust, efficient and accurate in all of the regimes present in the flow. Pressure-based schemes are suitable for many low Mach number flows, but are limited to low velocities and relatively small temperature variations. Density-based schemes struggle to converge in low-speed flows due to the time-step restrictions imposed by the acoustic velocity, which may be orders of magnitude larger than the convective velocity. Furthermore, such codes may exhibit excessive numerical dissipation due to improper scaling of the dissipative properties of the scheme. Chemical reactions introduce another set of temporal scales associated with the kinetics mechanism used to model the system. These scales are often much smaller than the convective or acoustic scales and impose additional restrictions on the time-step. This disparity requires numerical schemes designed to handle the challenges that occur in low Mach number, chemically reacting flows. Analysis of density-based schemes at the low Mach number limit suggests that the development of improved, robust preconditioning with suitable operator splitting techniques leads to improved solution fidelity.

In this work, a dual-time framework with low-Mach preconditioning is developed for complex, chemically reacting large-eddy simulations. A new version of the well-known MacCormack scheme is proposed and the resulting scheme improves the solution quality significantly at low Mach numbers. An established ordinary differential equation solver for stiff systems treats the stiffness associated with the chemical source terms. Methods to couple the PDE and ODE solvers in both pseudo-time and in physical time are proposed and analyzed. Validation of the non-reacting scheme and the coupled reacting scheme using canonical test cases demonstrates the improved solution fidelity and simulations of representative industrial applications demonstrate the combined scheme.

CHAPTER 1

THESIS OVERVIEW

A number of schemes and approaches exist for simulating turbulent, chemically reacting flows across the Mach number spectrum. Many of these schemes are suited for a particular regime and ill-suited for others. For example, schemes based on the incompressible form of the governing equations are not designed to capture the features of moderate or high Mach number flows. On the other hand, schemes designed around high Mach number flows are required to capture sharp gradients that may occur but these schemes are overly dissipative for smooth flows. For many chemically reacting, turbulent systems, a full range of Mach numbers may appear throughout the domain and numerical schemes must be designed to handle all of the possible regimes.

Section 1.2 discusses many of the approaches for extending solvers to low Mach number flows. The schemes are typically central schemes with explicitly added artificial viscosity for stability (analysis of upwind schemes often relies on writing the scheme as a central scheme with artificial viscosity arising from the upwinding). The Lax-Wendroff family of schemes is a unique class of schemes that are central in nature and inherently stable, requiring no additional artificial dissipation. To the author's knowledge, this class of schemes has not been analyzed previously in the low Mach number limit. The focus of this work is to develop an improved preconditioning approach for chemically reacting flows using a member of the Lax-Wendroff family of numerical schemes commonly used for Large Eddy Simulation (LES), the classic MacCormack predictor-corrector scheme.

The remainder of this chapter provides a foundation for the analysis that follows. It starts by motivating the interest in low Mach number scheme behavior in chemically reacting, turbulent flows and detailing the challenges that arise in these conditions. A number of schemes are surveyed to identify the important features required for low Mach numbers as

well as the techniques used to address the integration of chemical source terms. It concludes with an overview of the accomplishments of this work and an outline of this thesis.

1.1 Motivation

Computational resources continue to expand in both size and power, enabling simulations of large aerospace systems to use methods with increasing fidelity as time moves forward. In the realm of computational fluid dynamics, there are several classifications of methods that may be used for turbulent, chemically reacting flows. At the highest fidelity levels, direct numerical simulation (DNS) attempts to resolve every aspect of the flow in question. This demands large amounts of resources and, barring revolutionary breakthroughs, simulating practical aerospace applications remains impossible for the foreseeable future. On the other end, the Reynolds-Averaged Navier-Stokes (RANS) approach reduces the majority of the resolution requirements imposed by DNS through temporally filtering the governing equations. The filtering operation implies the solution is for the mean, or time-averaged, flow field. Consequently, all turbulent features and their effects on the mean flow must be modeled. This permits rapid solutions and enables parametric design studies that are useful for engineering applications. However, by filtering in time, the complex dynamics and unsteady features that may be important in combustion are omitted. In between the two methods lies a technique called Large Eddy Simulation (LES). Computational resources now permit their use on highly complex, turbulent, chemically reacting flow fields. LES uses a spatial filtering technique to remove only the smallest scales of the turbulent flow field from the solution. This solves the problem of stringent resolution requirements that arise in DNS. At the same time, resolving the temporal features of the flow field and requiring models for only the smallest, scales solves the limitations of RANS approaches. Figure 1.1 illustrates the different classifications with respect to an example turbulent kinetic energy spectrum.

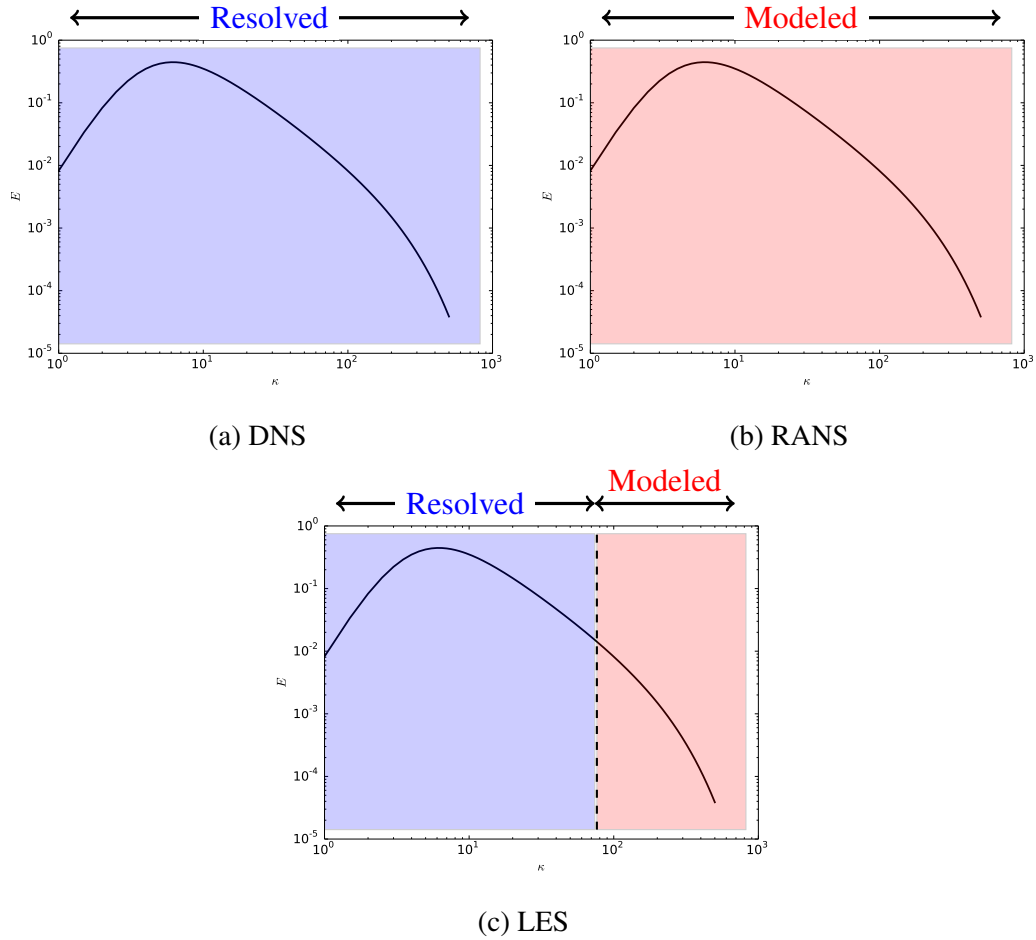


Figure 1.1: Illustration of the resolved and modeled regions of each computational technique in an example spectrum of turbulent kinetic energy.

Proper LES requires accurate models to represent the turbulent scales that are not directly resolved. These turbulent scales represent a portion of the inertial range and the entirety of the dissipation range of the turbulent spectrum, and as such, the models replacing those scales serve to mimic their dissipative effects. This generally takes the form of a turbulent viscosity that is added to the physical viscosity in the viscous stress tensor. In many instances, the form of these models is derived from or based on theoretical work and formalized independent of the numerical scheme to which they are coupled for simulations.

When formulating the model independently from the numerical schemes, several assumptions are generally made. It is often assumed that the filter used to separate scales is

a perfect, low-pass filter in spectral space. However, in implicitly filtered LES, the filter itself is never formally defined; instead, the mesh and the numerical scheme provide the filter. Although it is often assumed that these filters are box-filters in physical space, the actual behavior of the filter depends on the numerical scheme and is unlikely to be an ideal (ie. sharp) filter [1]. This means that when a model is used in a simulation, its behavior is intimately tied to the numerical scheme used. Therefore, it is essential to understand the characteristics of the numerical scheme in the regimes of interest and to have numerical schemes whose behavior can be controlled.

The dissipative behavior of the numerical scheme is important to understand and control because it is easy for the influence of the turbulence model to be overwhelmed by the numerical dissipation in the scheme. The ideal numerical scheme for LES is one that has the minimal amount of numerical dissipation required for stability to ensure the model performs optimally. This also suggests the ideal scheme should have dissipative properties that are independent of flow conditions because turbulent, reacting flows contain a wide range of conditions throughout the domain. It is well-known that without proper treatment, many numerical schemes used for compressible flows are excessively dissipative in regions of the flow with low Mach numbers [2, 3].

The overly dissipative nature of schemes used for compressible flows at low Mach numbers is not the only difficulty encountered. Schemes that employ explicit temporal integration methods show minimal dissipation from the temporal discretization, but stability constraints impose small time steps based on the acoustic wave speeds. When the disparity between the acoustic wave speed and the convective wave speed is large, these schemes take an excessive number of time steps to resolve the features of interest. Additionally, these schemes are prone to numerical oscillations that result from the decoupling of the velocity and pressure fields at low Mach numbers, known as checker-boarding. Figure 1.2 exemplifies the oscillations in the velocity field when an explicit, compressible scheme is used at low Mach numbers (details of this case are provided in Chapter 5). The oscillations




Figure 1.2: The solution of a low Mach number flow over a circular bump using a compressible scheme exhibits undesirable oscillations and asymmetries in the velocity field, demonstrating the common quality issues that arise for many schemes at low Mach numbers.

and the asymmetry are emblematic of the solution degradation that occurs for many schemes at low Mach numbers.

In the context of turbulent combustion, low Mach number flows are common and occur in many systems that also contain high-speed or compressible flows. For example, many freely-propagating flames and the deflagration of solid propellants occur at velocities ranging from millimeters per second to a few meters per second. Gas turbines may have velocities on the order of 100 meters per second, but also contain recirculation regions or stagnation regions where velocities approach zero. Low-speed regions exist in many swirl combustors around the annulus after the dump-plane or in the form of recirculation bubbles due to complex vortex breakdown processes; reverse flow regions behind bluff-body stabilized flames are similarly low speed. Even combustion problems classically envisioned as high-speed, such as rocket motors with converging nozzles, exhibit regions of low-speed or stagnation flow within the combustion chamber. Figure 1.3 illustrates these types of flows. In many of these cases, the low speed regions control the behavior of the system through flame stabilization.

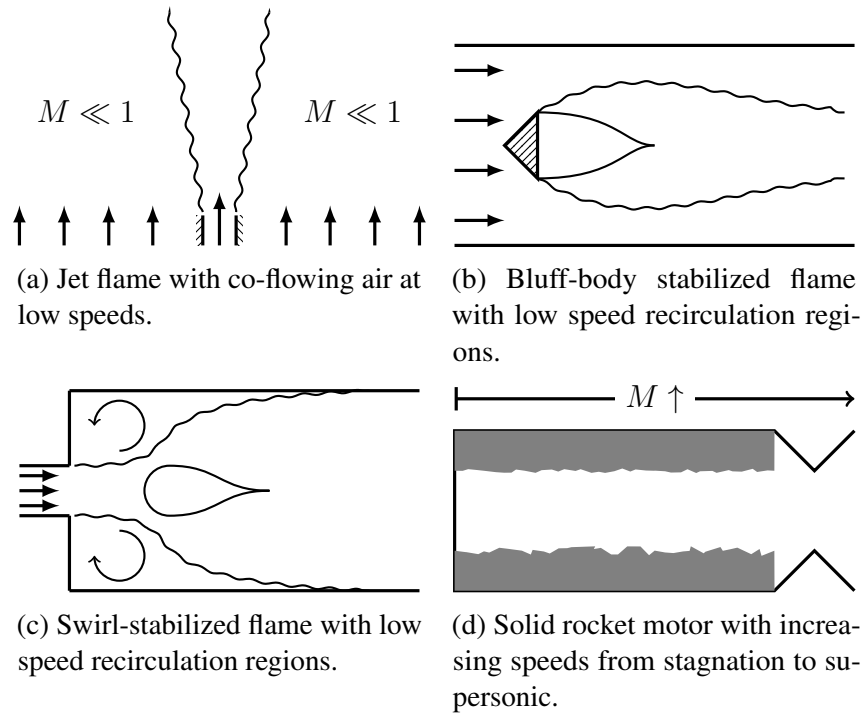


Figure 1.3: Illustrations of several typical combustion configurations that contain low speed flow regions that often dictate the dynamics of the system.

Therefore, accurate simulation of these systems requires a numerical scheme that enables proper control of its dissipative properties, including Mach number independence.

The current challenges in engineering drive the need for accurate simulations. Specifically, many of the current challenges are centered on prediction and control of combustion instabilities and/or improving fuel efficiencies and reducing pollutant emissions. The former requires accurate resolution of the flow field, particularly the coupling between the pressure field, the velocity, and the reaction rates, in all regions of the flow. The latter requires the use of detailed chemical kinetics that are capable of predicting both low- and high-temperature fuel oxidation and production of pollutants such as NO_x. These chemical mechanisms often have time scales that are much smaller than the acoustic or convective time scales, resulting in even more stringent temporal resolution requirements than typically associated with explicit schemes. The overarching goal of this work is to devise a numerical scheme for

LES that is capable of accurately simulating all regions of a chemically reacting, turbulent flow field using detailed chemical kinetics.

1.2 Literature Review

With the prevalence of low Mach number conditions alongside compressible and high-Mach numbers in many interesting systems, schemes that have properties independent of Mach number are desirable. One must understand the origin and nature of the difficulties at these conditions before schemes to address the challenges may be developed. This section outlines the challenges that occur at low Mach numbers. Following this, a number of schemes for these conditions are presented and assessed relative to the desired goal of turbulent, chemically reacting LES. Finally, the section concludes with a discussion of approaches to treat stiff chemical source terms, particularly in the context of the low Mach number conditions of interest.

1.2.1 Challenges at Low Mach Numbers

The numerical difficulties in low Mach number flows emerge from the disparity in the eigenvalues, λ_i , of the system [2, 3]. Mathematically, if the condition number, or stiffness, of a matrix is $\kappa = \max|\lambda_i|/|\lambda_j|$, the system using the matrix becomes more difficult to solve (or stiffer) as $\kappa \rightarrow \infty$. For the Euler equations, the matrix of concern is the flux Jacobian and the resulting eigenvalues in three dimensions are $\lambda \in \{u + c, u - c, u, u, u\}$ where u is the convective velocity and c is the acoustic velocity such that the eigenvalues $u + c$ and $u - c$ correspond to the right- and left-running acoustic waves. This gives a condition number $\kappa = 1 + 1/M$ where M is the Mach number [3]. In low velocity regions of the flow, and most evidently at stagnation points, the condition number of the Euler equations tends towards infinity and the equations are ill-conditioned. For Navier-Stokes calculations, the eigenvalues and condition number remain the same for flows that are convection-dominated (ie. the Reynolds number is large, $Re \gg 1$). As the velocity approaches zero, two branches

occur in the eigenvalues depending on the relationship between the Mach number and the Reynolds number [3]. In the event that $Re \ll 1$ and $Re/M \gg 1$, the acoustic eigenvalues remain the same but the three remaining eigenvalues are imaginary with $|\lambda| = u/Re$, giving a condition number $\kappa = (1 + 1/M)Re$. This is an ill-conditioned case because $Re/M \gg 1$. For the situation where $Re \ll 1$ and $Re/M \ll 1$, the first two eigenvalues are no longer related to the acoustic velocity while the remaining three are unchanged from the previous branch. The first eigenvalue becomes u and the second is imaginary with $|\lambda| = u\gamma/Re$ where γ is the ratio of specific heats. The resulting condition number is $\kappa = \gamma/Re$, which is again ill-conditioned because $Re \ll 1$. This stiffness and the resulting ill-conditioning presents numerical challenges that need to be addressed for accurate and efficient simulation of low Mach number regions in a flow.

1.2.2 Numerical Methods for Low Mach Number Flows

Multiple methods exist to reduce the difficulties due to stiffness at low Mach numbers. These methods are generally classified as either pressure-based methods or density-based methods, depending on the original form of the governing equations [4]. Pressure-based methods originate from the incompressible form of the governing equations and use corrections to recouple the velocity and pressure terms. In the incompressible limit, the governing equations are no longer hyperbolic and the pressure becomes independent of density, requiring the pressure to be updated using a Poisson equation rather than the equation of state. Time marching schemes like those used for compressible flows are not suited for this regime [5] and specifying the boundary conditions becomes difficult [6].

A class of iterative schemes called the semi-implicit, pressure-linked equations (SIMPLE) uses staggered grids by storing the pressure and velocity at the cell-centers and cell-vertices respectively [7, 8]. Acharya et al. [8] provide a review of the SIMPLE scheme as well as several of its variants that improve certain aspects of the original scheme such as the SIMPLE-Consistent (SIMPLEC) and SIMPLE-Revised (SIMPLER) schemes. Karki

and Pantankar [9] extended the SIMPLER algorithm by using two steps for the pressure correction. The first step updates the thermodynamic variables and the second updates the velocities. These extensions allowed the SIMPLER scheme to treat flows with temperature and density variations. Asymptotic analysis based on Mach number enabled Munz *et al* [10] to develop multiple pressure corrections for the SIMPLE schemes, extending the range to higher Mach numbers.

Non-iterative approaches based on pressure corrections also exist. The Pressure-Implicit with Splitting of Operators (PISO) scheme removes the iterative process from the SIMPLE schemes [11]. This family of schemes, when properly extended, is capable of solving reacting flows [12] and may be applied to moderate subsonic through supersonic Mach numbers [13]. Although the schemes may be used for a range of Mach numbers, compressible subsonic, transonic, and supersonic Mach number flows present numerical difficulties [13]. Both the SIMPLE and PISO family of schemes were designed for use on staggered grids to avoid cell-to-cell oscillations – so-called checker-boarding – due to the decoupling of pressure and velocity [4]. Extensions of the schemes to collocated grids are possible through proper interpolation techniques to introduce the required coupling. The well-known Rhie and Chow scheme eliminates the checker-boarding that appears on collocated grids [14] and may be viewed as a fourth-order artificial dissipation term applied to the pressure field only [4].

It is possible to apply the time-stepping techniques used in the compressible solvers to the incompressible equations. This is done by introducing an artificial compressibility factor δ into the governing equations such that an artificial equation of state relates an artificial pressure and artificial density, $p = \rho/\delta$ [6, 15]. With the addition of this term, the governing equations become hyperbolic, just as if they represented a compressible flow, and time marching schemes are again effective [5]. The value of δ is chosen to optimize convergence and is constrained by the independence of the final solution on its value. It can be shown

that artificial compressibility is equivalent to density-based preconditioning techniques with the proper definitions of δ [2].

The artificial compressibility approach takes the incompressible equations and creates a system that may be solved using a compressible solver. However, it is also possible to approach the problem from the other side and devise a method that prevents the density-based methods from becoming ill-conditioned in the low Mach number limit. This approach is desirable for flows with a wide-range of Mach numbers where the compressible equations may be required in some regions of the flow, while others are effectively incompressible; it is also important for flows with significant thermal effects where the density is variable but the flow velocities are low. These methods are also usually based on collocated grids, which require significantly less book-keeping, are easier to implement, and require less memory compared to the staggered schemes [4, 16]. The low Mach number extensions for density-based methods fall into two categories: asymptotic methods and preconditioning methods.

Asymptotic methods are based on perturbation analysis of the governing equations. This is done by expanding the equations using Taylor-series approximations based on Mach number and this has been done in the context of reacting flows [17–19]. Components of the perturbed equations may be neglected, thereby reducing the differences in the wave speeds. However, neglecting these components eliminates the ability to resolve those components when they may be important. This prevents the application of asymptotic methods to acoustic-combustion interactions and instabilities, as well as flows with a wide range of velocities. Because of this limitation, these methods are not considered further.

Preconditioning techniques modify the eigenvalues of the coupled system of equations to reduce the stiffness [2, 3, 5, 20]. In contrast with the asymptotic methods, preconditioning methods may be designed to modify the eigenvalues based on local flow conditions. This enables the retention of acoustic waves when physically relevant and allows optimal convergence rates and accuracy regardless of Mach number [3]. Choi and Merkle [21]

introduced a matrix preconditioning scheme to the compressible equations and compared the convergence across Mach numbers to the convergence of the artificial compressibility approach. The authors found that with this matrix preconditioning approach, the rate of convergence is independent of Mach number and convergence is rapid for Mach numbers as low as $M = 0.05$ [21]. Shortly thereafter, Turkel [22] sought to combine several specialized preconditioning approaches into a unified approach. A generalized preconditioner that uses two parameters, α and β , is demonstrated for a number of different dependent variable choices. It is shown that many of the existing preconditioners, including that of Choi and Merkle [21], are a subset of this generalized preconditioner and the Choi–Merkle preconditioner can be recovered by setting $\alpha = 0$. The modifications to the eigenvalues guarantee the system is well-conditioned, ensuring accuracy and efficiency across a wide range of Mach numbers [3].

This Mach number independent behavior is intimately tied to the dissipative properties of the discretization approach. The finite volume approach on collocated grids stores all of the solution variables at the cell-centers and interpolates this solution to quadrature points on the faces between cells. The interpolation techniques may be of any order of accuracy and may be symmetric (central) about the face or biased (upwind). The number of quadrature points on a face, combined with the interpolation order of accuracy, determines the formal, global order of accuracy of the scheme. Central interpolation contains no dissipative errors (although dispersive errors may exist) and is therefore ideal for smooth flows [4, 23]. The non-dissipative behavior of the interpolation minimizes errors, but also leads to numerical instabilities. Artificial viscosity techniques [24] or high-order, low-pass filters [25–29] may be used to add sufficient dissipation to provide numerical stability. Unfortunately, these techniques use the eigenvalues of the system. For low Mach number flows without preconditioning, the dissipation introduced is excessive. Proper scaling of the dissipative terms through preconditioning provides proper bounds and restores accuracy in a Mach number independent fashion [3]. Interpolation schemes with a stencil biased in the upwind

direction are typically more robust than central interpolation schemes owing to the inherent dissipation in the interpolation. These schemes likewise require proper scaling at low Mach numbers to ensure accuracy.

Schemes like the Convective Upwind and Split Pressure (CUSP) [30, 31] and the Advection Upwind Split Method (AUSM) [32] seek to control the dissipation through decomposing the total flux between cells into convective and pressure fluxes. This enables the interpolation procedure for each type of flux to be chosen independently. The CUSP scheme, for example, chooses an interpolation based on the local Mach number. If the flow is supersonic, both fluxes are upwinded; otherwise, the convective term uses an upwind interpolation while the pressure term uses a downwind interpolation. This matches the domain of dependence based on the eigenvalues of the system in subsonic flow [30]. The AUSM scheme is similar, but the implementation of the dissipation terms is different [33]. Both CUSP and AUSM have been extended to low Mach numbers with preconditioning [3, 5, 34–37].

In addition to the schemes just presented, there exists an additional class of schemes using central interpolation but possessing inherent numerical dissipation. This eliminates the need for additional numerical dissipation in smooth flows. These schemes use Taylor-series expansions in the independent variables, resulting in the well-known Lax-Wendroff scheme and several variants [4]. The MacCormack scheme [38, 39] is one such variant and is a popular choice for chemically reacting large-eddy simulations (LES). The original version of the scheme does not use any artificial dissipation or filtering for smooth flows. It is inherently stabilized by applying alternating sequences of first-order derivatives, resulting in a second-order accurate scheme. Using alternating first-order derivatives leads to a small numerical stencil, and this means the MacCormack scheme is easy and efficient to implement in parallel computational environments. As an explicit scheme, there are no global operations required to advance the solution in time. The two-step formulation is well-suited for non-linear equations, including those with source terms, and eliminates much

of the computational expense of the original Lax-Wendroff scheme [4]. Although frequently used for compressible flows [4], the MacCormack scheme has not been previously analyzed or applied to low Mach number flows. It is this scheme that forms the foundation of this work.

The early studies using preconditioning focused primarily on the steady-state solution and therefore the preservation of time-accuracy was not required. For unsteady calculations where time-accuracy is required, a dual-time procedure is created where a pseudo-time derivative is introduced to the governing equations and the solution is marched in pseudo-time until convergence is reached. The resulting field is the unsteady solution at the resulting physical time level [40]. Withington, Shuen and Yang [41] first demonstrated the dual-time approach combined with preconditioning for chemically reacting flows. The authors used a first-order implicit solver in pseudo-time to drive convergence to the solution of a second-order accurate implicit scheme in physical-time for each physical timestep. The chemical source terms were computed using a point-wise, implicit scheme; no special treatment was given to equalize the chemical, convective, and acoustic eigenvalues. This approach was generalized to include a viscous preconditioning parameter [20] and demonstrated in multi-dimensional flows [42]. Again, no special treatment was given to the chemical eigenvalues and the chemical source terms were computed using an implicit scheme. Weiss and Smith [5] proposed a preconditioner that was not restricted to ideal gases and retained the derivatives of temperature with density for solving flows where density is only a function of temperature. This preconditioner was extended to account for chemical species and used to solve reacting flows; as with the other authors, the chemical source terms were treated implicitly and no special treatment was given to the chemical eigenvalues [43]. In a later review, Turkel [2] indicates that many of the successful preconditioning applications have used the Weiss–Smith form of the preconditioner and it is this form that is used in the current work.

The introduction of preconditioning removes the time-accuracy of the equations by modifying the wave speeds [3]. This time-accuracy is restored through the introduction of a pseudo-time derivative that contains the preconditioning instead of applying the preconditioning to the physical-time derivative. The pseudo-time derivative is iterated until the converges to a “steady-state”, at which point the governing equations return to their unaltered form [3]. This two-level time scheme is known as a dual-time scheme, first introduced by Jameson [44] in the context of accelerating the convergence of an implicit scheme. In that instance, the objective was to use a multi-grid acceleration technique alongside local time stepping for the convergence acceleration. Both of these techniques are only valid for steady-state solutions, but the introduction of the dual-time procedure extended them to unsteady solutions. The inner iterator may be either explicit or implicit, which permits considerable flexibility in implementation. In particular, choosing an explicit scheme as the pseudo-time iterator maintains all of the simplicity and efficiency of implementation on parallel computers to yield an implicit scheme in physical time. Combined with techniques to accelerate convergence such as local-time stepping or multi-grid, and/or techniques to increase the local time step values such as low-Mach number preconditioning, dual-time schemes relax many of the difficulties with implementing implicit schemes.

1.2.3 Treatment of Stiff Chemical Kinetics

Even for convection dominated flows at moderate speeds – like those occurring in aerospace and industrial combustion systems – multi-component flow and chemical reactions introduce a new set of time scales and eigenvalues. The chemical reactions themselves may introduce eigenvalues that are the inverse of the characteristic chemical time scales [45]. These eigenvalues can be well over $1 \times 10^{20} \text{ s}^{-1}$, which introduces considerable stiffness into the system in the flame fronts for even moderate flow velocities [3]. Additionally, the sign of the chemical eigenvalues affects the stability of the numerical methods used. Positive eigenvalues that arise during pre-heat and ignition phases of combustion are unconditionally

stable for explicit schemes but only conditionally stable for implicit schemes [46]. The reverse is true for negative eigenvalues that occur throughout the flame. This disparity demands the use of sophisticated solvers such as DVODE [47], dynamic/adaptive methods to reduce stiffness [48, 49], or the use of time-step restriction techniques in the scheme used to integrate the full system of equations.

Reduction of the stiffness in the chemical source terms may be achieved using multiple techniques. The dynamic adaptive hybrid integration (AHI) approach integrates the fast species and reactions using an implicit scheme while the slow reactions are integrated using an explicit scheme [50]. The speed of the species is determined during the simulation based on their respective analytical timescales. By removing the slow species from the system requiring implicit integration, the size of the source term Jacobian is significantly reduced and in turn the cost of the implicit solver is also reduced.

In 2007, Sankaran and Oefelein [51] presented a preconditioning method for integrating stiff chemical mechanisms in a zero-dimensional reactor. The authors attempt to derive a generic scaling parameter based on the local combustion process to solve the problem of *ad hoc* scaling parameters used by others [3, 52]. This scaling parameter is found using a simple, global mechanism and the parameter modifies the eigenvalue of the reacting system, turning the large, positive eigenvalue into a negative eigenvalue, stabilizing the system. However, the authors were not able to generalize the process for a generic, multi-step chemical mechanism. More recent work by Hansen and Sutherland [53] extend the work of Sankaran and Oefelein [51] for use with larger, more complex chemical mechanisms. These efforts for preconditioning the source term integration are still confined to zero-dimensional reactors and have not yet been applied to multi-dimensional flows with reactions.

The second approach is to treat the chemical source terms as a local, pointwise problem decoupled from the spatial fluxes. This approach is known as operator splitting and permits the use of dedicated solvers for solving the stiff chemical source term differential equations that arise within each computational point [54–59]. Such solvers are optimized for treating

the disparate scales in detailed kinetics. The most popular of these methods are based on implicit methods for sparse systems. Numerous solvers exist and combine information contained in the eigenvalues with careful selection of time step size to control the error and minimize the stiffness of the problems. The DVODE solver [47] is a popular choice.

Lian et al. [46] discuss how source terms arising from chemical reactions, as well as those from turbulence models, can negatively impact the stability and convergence of the solution. Positive source terms are best handled with an explicit integration scheme while negative source terms (sinks) are best handled implicitly. These positive source terms cause the solution to grow and care must be taken to choose appropriate time steps lest the solution grow too much and therefore deviate from the physical and accurate values causing instability. Sinks, on the other hand, work to stabilize the solution by driving it to zero. Explicit treatment of sink terms may lead to unphysical sign changes if the integration step is too large, causing instability or unphysical values. In the context of the dual-time framework, care must be taken in choosing the physical time step to avoid instabilities due to source terms [46, 60].

Source terms are not the only difficulty in unsteady flows; strong non-linearities preclude large time step advances [60]. Lian et al. [60] introduce a method to limit the time step by estimating the change in the solution over the chosen time step using the residual. An upper bound on the time step is user-specified and represents the ideal integration size based on the users' requirements. If this estimate exceeds an acceptable tolerance, a smaller time step is chosen so that the limit on the change in the solution is enforced. In combination with the local time step acceleration technique employed with the dual-time stepping procedure, it is shown that this limiting procedure is typically confined to small regions of flows where non-linearities or source terms are strong [60]. By only using the limiting approach where needed, the stability and accuracy of the scheme is not compromised while the increase in cost is relatively minimal. Unfortunately, applying these non-linear time step controls to

multi-step, explicit schemes is not straightforward. Techniques to enforce these limits using the preconditioning parameter directly are proposed in this work.

Venkateswaran et al. [52] discuss the specific issues that arise in chemically reacting flows. The authors identified two primary regimes in the chemical reaction process: the first, a non-linear phase, is characterized by both positive and negative eigenvalues (sources and sinks, respectively); the second phase is called the linear phase and is characterized by predominately negative eigenvalues. Stability analysis shows that implicit schemes are unconditionally stable for chemical mechanisms only when all the eigenvalues are negative — in the presence of large, positive eigenvalues (strong sources), implicit schemes may be unstable and explicit schemes should be used [3, 52]. These large, positive eigenvalues are located in regions of ignition and may dominate the other terms in the species equations while the resulting heat release dominates the other terms in the energy equation. To address these issues, it is possible to gradually introduce the heat release and chemical source terms through an *ad hoc* scaling function that increases through the iterations until the full heat release and source terms are reached [52]. Although effective, this destroys the time accuracy of the simulation and may be used only when the steady state is desired. This is not a problem when using the dual-time formulation, where the convergence to “steady state” in pseudo-time is the objective. For time-accurate solutions, it is possible to adapt the time step, reducing it when the chemical reactions introduce large, positive eigenvalues and increasing it when the eigenvalues are predominately negative; this approach is also *ad hoc* and a generic way to scale the time step based on the local combustion process is unknown [52]. At transonic, supersonic and hypersonic conditions, the flow time scales and the chemical time scales may be close to one another and the system will be well-conditioned. However, at lower speeds like those occurring in many combustion applications, the disparity between the flow time scales and the chemical time scales increases and the conditioning of the system becomes worse. This is exacerbated when preconditioning is used because the flow

time scale is now close to the convective time scale instead of the much-smaller acoustic time scale [61].

1.3 Objectives

The motivating interest is to develop a methodology that extends the existing compressible capabilities for LES in the LESLIE code [62, 63] into the low Mach number regime with stiff chemistry, creating an all-Mach number, reacting flow solver. In order to accomplish this, the following objectives have been met:

1. **Analyze the MacCormack scheme at low Mach numbers and develop a generalized, dual-time form of the scheme using preconditioning to address deficiencies and improve performance.**

Analysis of the 1971 form of the MacCormack predictor-corrector scheme [39] shows that the inherent dissipation is ill-posed at low Mach numbers, consistent with the behavior of unmodified artificial dissipation schemes [2, 3]. The straight-forward application of preconditioning to the original scheme ensures proper scaling of the inherent dissipation for steady preconditioning, but unsteady preconditioning reverts to the overly-dissipative form. To address these limitations, a new, generalized form of the scheme is proposed where the original scheme is augmented with a pressure-only artificial dissipation similar to the Rhie and Chow scheme [14]. The results indicate the scheme exhibits less dissipation for low Mach number flows relative to the original 1971, explicit scheme and the 1971 scheme with preconditioning. The reduction in dissipation means the scheme is more efficient at low Mach numbers as it takes significantly refined grids with the non-preconditioned scheme to match the solution quality. Moving towards LES, it is also shown that the temporal evolution of isotropic decaying turbulence at low Mach numbers exhibits less numerical dissipation than the original scheme and the resulting flow statistics are insensitive to the physical time step chosen for the time steps tested. This is an important step to ensuring that

the numerical dissipation is minimal and that the physical and modeled turbulent dissipation are the dominant contributors to the dissipation in the flow.

Enhancements to the robustness and efficiency of the preconditioned scheme are made through the introduction of new preconditioning definitions. These definitions restrict the amount of preconditioning applied to the equations when the local conditions show strong non-linearities. The restrictions are derived using a Fourier analysis technique to impose limits used in other, non-linear time step control schemes [60]. By limiting the preconditioner instead of changing the time step, the new preconditioner definitions may be used with any preconditioned scheme. Additionally, the Fourier technique used permits the design of preconditioners to impose new limits based on the local flow conditions.

2. Develop schemes to couple external, stiff chemical source term solvers to the dual-time framework to enable robust, accurate simulations with detailed chemical kinetics.

The introduction of preconditioning permits larger time step sizes for each cell in the pseudo-time iteration relative to the explicit integration scheme without preconditioning. As a result, the integration of the chemical source terms must be treated carefully to ensure accuracy despite the large difference between the time step size and the characteristic time scales of the reactions. This problem is exacerbated with stiff chemical mechanisms. Common approaches to mitigate stiffness issues in single-time schemes based on operator splitting techniques introduce errors that, when applied to dual-time schemes, can compromise the accuracy of the resulting scheme. The traditionally, fully coupled approach to treating chemistry in dual-time schemes contains *ad hoc* modifications that can lead to a lack of robustness and/or efficiency. Two new schemes are proposed to address these limitations and are based on approaches used to integrate chemical source terms in single time schemes. Care has been taken

in the newly proposed schemes to understand and mitigate the errors that arise from operator splitting.

The first scheme is an extension of the consistent splitting scheme, first demonstrated by Schwer et al. for single-time schemes, to dual-time schemes [64]. This scheme uses a stiff solver to integrate the chemical source terms and spatial fluxes together, the former treated implicitly and the latter treated explicitly using the generalized MacCormack scheme. Integrating these together eliminates the steady-state errors that otherwise degrade the accuracy of operator split schemes [65] and ensures that all error terms arising from the mixed implicit/explicit treatment of the terms are contained on the left-hand side of the system [64]. The importance of the consistent splitting approach is demonstrated by using an alternate scheme that employs a traditional operator splitting scheme in pseudo-time. Because operator split schemes are known to converge to the incorrect steady-state without special modifications [65], operator splitting in pseudo-time leads to unacceptable errors in the physical time accuracy of the scheme as shown in Chapter 7

The second scheme is based on the well-known Strang splitting [66] in physical time, where the preconditioned scheme is used to solve the stages associated with the flux operators and an external, stiff solver is used to solve the stages associated with the reactions. This is a straight-forward and logical extension of the operator splitting done in single-time schemes, but it requires subtle changes in a dual-time scheme to ensure consistency with the original governing equations when a dual-time scheme is used to integrate the fluxes. Additionally, the scheme has splitting errors in physical time proportional to the time step and the eigenvalues of the Jacobian from the chemical mechanism. These errors may be unacceptably large for large physical time step values, but acceptable at smaller choices. Despite these drawbacks, the scheme requires very few chemical integration steps relative to the other approaches and this can result in substantial cost reductions.

3. Implement advanced turbulence-chemistry closures in the dual-time, preconditioned framework.

Using grids whose spatial resolution is determined by the turbulent flow features, as is often done for LES, results in grids too coarse to resolve the flame front. Numerous approaches exist to treat this disparity in resolution and to model the interactions between the turbulence and the flame. In this work, the dynamically thickened flame (DTF) model [67] is used with the dual-time, preconditioned scheme for practical applications. It is also shown that the ideal physical time step size for a non-reacting LES is too large for a chemically reacting flow, whose time step size requirement based on solution accuracy is driven by the reaction-diffusion processes in the flame front.

4. Apply the dual-time preconditioning scheme to a complex, reacting flows representative of industrial or aerospace applications.

Two configurations of practical interest are considered. The first is the Volvo Flyg-motor, a premixed, bluff-body stabilized combustor that is the subject of many experimental and numerical studies [68–78]. Simulations performed using an explicit, compressible scheme on three grid resolutions as part of the Model Validation for Propulsion workshop at the 2017 AIAA SciTech meeting [79] serve as a baseline for comparison with the preconditioned scheme on the coarsest grid [80]. The preconditioned scheme captures the recirculation zone in the experimental measurements more accurately than the explicit scheme, where all three grid levels predicted a wake-like structure. Other metrics are compared relative to the experimental data and explicit results. On the coarsest grid, the preconditioned scheme takes approximately ten times more computational expense than the explicit scheme on the same grid. However, consistent with the observations in the simpler cases, the explicit schemes requires more than a factor of ten increase in resources to capture the same solution quality.

This is not demonstrated directly, but inferred from the fine grid solution with the explicit scheme.

The second configuration is the Cambridge slot burner with a rod-stabilized V-flame [81–84]. This has available experimental datasets and the premixed and stratified conditions are simulated here. With inflow velocities of approximately 5 m s^{-1} and moderate levels of turbulence, this is an ideal target for evaluating the preconditioned scheme using a stiff chemical mechanism for the methane fuel. Due to the low Mach numbers in the flow, the preconditioned scheme is faster on a given grid resolution relative to the explicit scheme, in addition to the accuracy benefits of the preconditioned scheme. The preconditioned scheme’s success with this case demonstrates the improved capabilities enabled by this work.

1.4 Thesis outline

Chapter 2 outlines the physical governing equations used throughout the work, including the aspects of Large-Eddy Simulation and the required models. This is followed by a presentation of the original, 1971 form of the MacCormack scheme in Chapter 3. Analysis shows that the inherent dissipative behavior of this family of schemes is degraded in the low Mach number limit and modifications to the MacCormack scheme are needed to ensure Mach-number independent dissipation and optimal scheme convergence. Chapter 4 provides the changes needed to rectify the solution issues at low Mach numbers through the introduction of a preconditioned, dual-time formulation of the MacCormack scheme. The selection of time step size in the dual-time, preconditioned framework is discussed. Novel preconditioner definitions based on a Fourier analysis of perturbation responses improve the robustness to pressure disturbances by several orders of magnitude. Several canonical test cases building up to the LES of practical configurations in Chapter 5 highlight the improved performance of the proposed scheme.

With the fidelity of the non-reacting scheme established, the emphasis turns to the reacting scheme formulation and validation. Chapter 6 outlines the existing techniques used to couple external, stiff solvers for ordinary differential equations (ODEs) for the integration of chemical source terms in single-time schemes. These techniques are then extended to dual-time schemes and the importance of careful selection of operator splitting techniques is discussed. The reacting schemes are validated using the autoignition of high pressure, lean, premixed hydrogen mixtures with and without temperature and composition inhomogeneities. This configuration is selected because it is well studied using direct numerical simulation, the detailed chemical mechanism is stiff, and slight variations in temporal accuracy or in the coupling between reaction and diffusion manifolds results in large variations in the final solutions. This tests all of the critical aspects of the resulting schemes.

Based on what is learned during the autoignition studies in Chapter 7, the dual-time, preconditioned scheme is applied to two combustion rigs of industrial interest in Chapter 8. The first is a well-studied, premixed, bluff-body stabilized flame. Extensive numerical and experimental data is available for comparison with this rig. The second is a turbulent, rod-stabilized V-flame with both premixed and stratified configurations. Experimental data is available for this rig, but to date, this is the first application of LES on this configuration. Finally, the results are summarized and critiqued in Chapter 9 and avenues for future research are given.

CHAPTER 2

GOVERNING EQUATIONS

The types of flows of interest in this work are continuum flows with variable thermodynamic and transport properties in both space and time. Additionally, they are generally turbulent and therefore display a wide range of spatio-temporal scales. This chapter begins with an introduction of the equations for a three-dimensional, viscous flow. The formulation is separated into non-reacting, single component fluids and reacting, multi-component fluids. Although many equations of state exist, all cases considered are ideal gases and the description of this equation of state is provided. The details of the evaluation of the transport properties is also given.

Properly resolving every length and time scale in a turbulent flow is cost-prohibitive. To make the simulation of complex flows tractable, a technique based on spatially filtering the governing equations called large-eddy simulation is employed. The details of this approach as well as required models are given. The chapter ends with a description of the dual-time, preconditioned framework for chemically-reacting LES. Other than the extension of the preconditioning matrix to include the one-equation kinetic energy model, these developments are consistent with the existing literature [3, 5].

2.1 Non-reacting Navier-Stokes Equations

The vector form of the Navier-Stokes equations for a viscous fluid neglecting body forces is:

$$\frac{\partial}{\partial t} \iiint_{\Omega} \mathbf{W}_{NR} dV + \iint_{\partial\Omega} [\mathbf{F}_{NR} - \mathbf{G}_{NR}] \cdot d\mathbf{A} = \mathbf{S}_{NR} \quad (2.1)$$

where Ω is the domain, $\partial\Omega$ is the boundary of domain Ω , and

$$\begin{aligned}
 \mathbf{W}_{NR} &= \begin{pmatrix} \rho \\ \rho u \\ \rho v \\ \rho w \\ \rho E \end{pmatrix}, & \mathbf{F}_{NR} &= \begin{pmatrix} \rho \mathbf{u} \\ \rho \mathbf{u} u + p \hat{i} \\ \rho \mathbf{u} v + p \hat{j} \\ \rho \mathbf{u} w + p \hat{k} \\ \mathbf{u} (\rho E + p) \end{pmatrix}, \\
 \mathbf{G}_{NR} &= \begin{pmatrix} 0 \\ \tau_{xi} \\ \tau_{yi} \\ \tau_{zi} \\ u_j \tau_{ij} + \mathbf{q} \end{pmatrix}, & \mathbf{S}_{NR} &= \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}
 \end{aligned} \tag{2.2}$$

The density is ρ and $\mathbf{u} = u\hat{i} + v\hat{j} + w\hat{k}$ is the velocity vector. The density and pressure p are related through the equation of state presented in Section 2.3. The total energy is composed of the internal energy, e and the kinetic energy, $1/2 \mathbf{u} \cdot \mathbf{u}$:

$$E = e + 1/2 \mathbf{u} \cdot \mathbf{u} \tag{2.3}$$

The internal energy for the single-component fluid here is:

$$e = \int_{T_0}^T c_v(\tau) d\tau \tag{2.4}$$

where c_v is the heat capacity at constant volume and may be a function of temperature.

The viscous shear stress tensor for a Newtonian fluid is:

$$\tau_{ij} = \mu(T) \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) + \lambda(T) \frac{\partial u_k}{\partial x_k} \delta_{ij} \tag{2.5}$$

where $\mu(T)$ is the molecular viscosity as a function of temperature, $\lambda(T)$ is the bulk viscosity and δ_{ij} is the Kronecker delta. For the fluids considered here, Stokes' hypothesis is applied such that the trace of the stress tensor is zero. The result of this hypothesis is to relate the bulk viscosity to the molecular viscosity, $\lambda = -2/3\mu$. Using this assumption, the viscous stress tensor is now:

$$\tau_{ij} = 2\mu \left(S_{ij} - \frac{1}{3} S_{kk} \delta_{ij} \right) \quad (2.6)$$

where $S_{ij} = 1/2 \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right)$ is the rate-of-strain tensor. The heat flux $\mathbf{q} = -\kappa(T) \frac{\partial T}{\partial x_j}$ where $\kappa(T)$ is the thermal diffusivity coefficient. The details of computing the state variables p and T are given in Section 2.3 and the details for transport properties κ and μ are in Section 2.4.

2.2 Reacting Navier-Stokes Equations

Consideration of multi-component, chemically reacting fluids extends the system of equations to incorporate the transport and reactions of the individual chemical species. This adds a new conservation equation for each of the species partial densities, which each may contain non-linear, coupled source terms. The resulting set of equations requires additional models for transport properties and equations of state to account for the influence of the variable composition. The governing equations are:

$$\frac{\partial}{\partial t} \iiint_{\Omega} \mathbf{W}_R dV + \iint_{\partial\Omega} [\mathbf{F}_R - \mathbf{G}_R] \cdot d\mathbf{A} = \mathbf{S}_R \quad (2.7)$$

with the set of variables for a multi-component flow neglecting body forces:

$$\begin{aligned}
 \mathbf{W}_R &= \begin{pmatrix} \rho \\ \rho u \\ \rho v \\ \rho w \\ \rho E \\ \rho Y_k \end{pmatrix}, & \mathbf{F}_R &= \begin{pmatrix} \rho \mathbf{u} \\ \rho \mathbf{u} u + p \hat{i} \\ \rho \mathbf{u} v + p \hat{j} \\ \rho \mathbf{u} w + p \hat{k} \\ \mathbf{u} (\rho E + p) \\ \rho Y_k \mathbf{u} \end{pmatrix}, \\
 \mathbf{G}_R &= \begin{pmatrix} 0 \\ \tau_{xi} \\ \tau_{yi} \\ \tau_{zi} \\ u_j \tau_{ij} + \mathbf{q} \\ -\mathbf{J}_k \end{pmatrix}, & \mathbf{S}_R &= \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \dot{\omega}_k \end{pmatrix},
 \end{aligned} \tag{2.8}$$

where Y_k is the mass fraction of the k -th species and \mathbf{J}_k is the species diffusion vector. As before, the total energy is composed of the internal and kinetic energies. In the multi-component case, the internal energy is the mass-weighted sum of each component k :

$$e = \sum_k^{N_k} Y_k e_k \tag{2.9}$$

where each component's internal energy is computed as described in Section 2.3.

The heat flux vector in the energy equation is updated to include the change in energy due to the diffusion of species:

$$\mathbf{q} = -\kappa(T) \frac{\partial T}{\partial x_j} - \rho \sum_k^{N_k} h_k Y_k \mathbf{V}_k \tag{2.10}$$

where h_k is the internal enthalpy of the k -th species and \mathbf{V}_k is the diffusion velocity of the k -th species. To reduce the expense in computing the species diffusion, the Hirschfelder-Curtiss approximation is used:

$$\mathbf{V}_k X_k = -D_k \frac{\partial X_k}{\partial x_i} \quad (2.11)$$

where X_k are the species mole fractions and where D_k is an approximate diffusion coefficient for species k into the rest of the mixture defined in Section 2.4 [45]. This approximation for the species diffusion leads to a system that does not conserve mass. By adding a correction term to the velocity, mass conservation is recovered and the total species diffusion velocity is [45]:

$$\mathbf{V}_k = \frac{D_k W_k}{X_k W} \frac{\partial X_k}{\partial x_i} - \sum_k D_k \frac{W_k}{W} \frac{\partial X_k}{\partial x_i} \quad (2.12)$$

where W_k is the molecular weight of the k -th species and W is the molecular weight of the mixture.

The final new terms are the chemical source terms, $\dot{\omega}_k$. These source terms are computed from M reactions with forward and backward rate expressions that take a general form:



for the j -th reaction out of M . In this notation, the molar stoichiometric coefficients for the reactants and products are ν'_{kj} and ν''_{kj} respectively and S_k represents the particular species in the reaction. In order to conserve mass, each reaction must satisfy:

$$\sum_k^{N_k} (\nu''_{kj} - \nu'_{kj}) W_k = 0 \quad (2.14)$$

The species reaction rate $\dot{\omega}_k$ is:

$$\dot{\omega}_k = \sum_j^M \left[K_{fj} \prod_k^{N_k} X_k^{\nu'_{kj}} - K_{bj} \prod_k^{N_k} X_k^{\nu''_{kj}} \right] \quad (2.15)$$

where K_{fj} and K_{bj} are the forward and backward kinetic rate constants of each reaction respectively. These kinetic rate constants are computed differently depending on the chemical mechanism used. The most common form of reactions is the Arrhenius rate form:

$$K_{fj} = A_{fj} T^{\alpha_j} \exp\left(-\frac{E_j}{RT}\right) \quad (2.16)$$

where A_{fj} , α_j , and E_j are empirical parameters. Other expressions are possible but the details are not important for this work. For all of the chemically reacting cases performed, the chemical source terms are evaluated using the *Cantera* package, version 1.8 [85].

2.3 Equations of State

For the conditions considered in this work, all gases are assumed to be ideal. This assumption neglects any intermolecular forces, outside of collisions, and also neglects any particle diameter effects, resulting in the relationship between pressure, density and temperature:

$$p = \rho RT \quad (2.17)$$

where R is the mixture gas constant and T is the temperature. The gas constant uses a mass-based blending of the molecular weights:

$$R = R_u \sum_k^{N_k} \frac{Y_k}{W_k} \quad (2.18)$$

where R_u is the universal gas constant. The species-specific internal energy is a function of that species' sensible energy at a reference temperature, $e_{k,0}$, and temperature:

$$e_k = e_{k,0} + \int_{T_0}^T c_{v,k}(\tau) d\tau \quad (2.19)$$

where $c_{v,k}$ is the specific heat at constant volume for species k . The internal energy is closely related to the internal enthalpy,

$$h = e + \frac{p}{\rho} \quad (2.20)$$

that is defined for the k -th species as:

$$h_k = h_{k,0} + \int_{T_0}^T c_{p,k}(\tau) d\tau \quad (2.21)$$

where $h_{k,0}$ is the sensible enthalpy of the species and $c_{p,k}$ is the specific heat at constant pressure of the species.

In the general case, $c_{p,k}$, $c_{v,k}$ and $h_{k,0}$ are functions of temperature. These are implemented using polynomial curvefits, generally segmented into two temperature ranges, and each species requires 7 coefficients for each temperature range:

$$\begin{aligned} \frac{c_{p,k}}{R_u} &= a_{1,k} + a_{2,k}T + a_{3,k}T^2 + a_{4,k}T^3 + a_{5,k}T^4 \\ \frac{h_{k,0}}{R_u T} &= a_{1,k} + \frac{a_{2,k}}{2}T + \frac{a_{3,k}}{3}T^2 + \frac{a_{4,k}}{4}T^3 + \frac{a_{5,k}}{5}T^4 + \frac{a_{6,k}}{T} \end{aligned} \quad (2.22)$$

The seventh coefficient, $a_{7,k}$, is used to compute internal entropy of the species and not needed for the equation of state as implemented. The perfect gas equation of state that maintains this temperature dependency is called the thermally perfect equation of state and it is valid for a wide range of pressures, temperatures and compositions.

For simpler flows, and particularly non-reacting flows, the expense of computing the temperature-dependent properties and solving the integral for the internal energy is not needed. In these cases, where the change in properties is small with the change in temperatures

in the flow, the heat capacities and enthalpies may be held constant. When this happens, the relationship between internal energy and temperature is simplified:

$$e_k = e_{k,0} + c_{v,k}T \quad (2.23)$$

and likewise for internal enthalpy,

$$h_k = h_{k,0} + c_{p,k}T \quad (2.24)$$

This equation of state is simpler to convert between internal energy and temperature, requiring neither iteration nor evaluation of an integral. This reduces expense, at the drawback of only being approximately valid as temperatures vary. This form of the perfect gas equation of state is known as calorically perfect.

2.4 Transport Properties

Expressions for κ and μ are needed to close all of the remaining unknown terms in the governing equations. As with the equation of state, there are many models that may be used to compute these coefficients. It is possible to compute the binary diffusion coefficients and treat the multi-component diffusion in a precise fashion. However, this requires N_k^2 diffusion coefficient calculations in each cell at each step and for large simulations with more than a few species, it is prohibitively expensive. It is also unneeded as many flame cases can be adequately treated using simplified approximations [45, 86, 87].

A simplified formulation for cases where the differential diffusion effects are important, such as the hydrogen autoignition cases considered in Chapter 7, assumes that each species has its own different, yet constant, Lewis number:

$$\text{Le}_k = \frac{\kappa}{\rho c_p D_k} \quad (2.25)$$

which relates the mixture thermal conductivity coefficient κ to the species mass diffusion coefficient D_k using the mixture density and mixture heat capacity. The mixture averaged thermal conductivity is computed using the local gas composition and temperature:

$$\kappa = \frac{1}{2} \left(\sum_k^{N_k} X_k \kappa_k + \frac{1}{\sum_k^{N_k} X_k / \kappa_k} \right) \quad (2.26)$$

where the species-specific thermal conductivity coefficient κ_k is computed using a polynomial curvefit:

$$\kappa_k = b_{0,k} + T (b_{1,k} + T (b_{2,k} + T (b_{3,k} + T (b_{4,k} + T (b_{5,k} + T b_{6,k})))) \quad (2.27)$$

With the mixture averaged thermal conductivity and the user-specified, species-specific, constant Lewis numbers, the species diffusion coefficients can be found. A similar process is used to compute the mixture viscosity using Wilke's formula:

$$\mu = \sum_k^{N_k} \frac{X_k \mu_k}{\sum_j^{N_k} X_j \Pi_{kj}} \quad (2.28)$$

where the species-specific viscosities μ_k are computed using another polynomial curvefit:

$$\mu_k = c_{0,k} + T (c_{1,k} + T (c_{2,k} + T (c_{3,k} + T (c_{4,k} + T (c_{5,k} + T c_{6,k})))) \quad (2.29)$$

The final term needed to complete the expression is:

$$\Pi_{kj} = \frac{\left[1 + \left(\frac{\mu_k}{\mu_j} \right)^{1/2} \left(\frac{W_j}{W_k} \right)^{1/4} \right]^2}{\sqrt{8} \left[1 + \frac{W_k}{W_j} \right]^{1/2}} \quad (2.30)$$

When referenced later, this is known as constant-Lewis number transport.

As with the equation of state, it is possible to avoid the evaluation of the curvefits under certain conditions. For non-reacting flows, particularly those composed of air, and for

many methane or propane flows at atmospheric conditions, approximate calculations for the thermal conductivity and viscosity are acceptable. This is useful when simulations are using very simplified chemical kinetics with only a few species. In this case, the transport coefficients may be computed using Sutherland's law:

$$\mu = \mu_{ref} \left(\frac{T}{T_{ref}} \right)^{3/2} \frac{T_{ref} + S}{T + S} \quad (2.31)$$

where μ_{ref} is a reference viscosity at a given composition and temperature, T_{ref} , and S is an empirically derived constant. The thermal conductivity is related to the molecular viscosity:

$$\kappa = \frac{\mu C_p}{Pr} \quad (2.32)$$

where $Pr = \mu C_p / \kappa$ is the Prandtl number for a given gas mixture and assumed to be a user-specified constant. From this, the species diffusion coefficients can be found:

$$D_k = \frac{\mu}{\rho Le Pr} \quad (2.33)$$

When referenced later, this is known as Sutherland's transport.

2.5 Large-Eddy Simulation

Fully resolving all of the scales in a turbulent flow field is prohibitively expensive and will remain so in the foreseeable future. Fortunately, Kolmogorov's 1941 theory postulates that, at sufficiently high Reynolds numbers, there is a wide range of eddy sizes from the largest scales to the smallest. Energy cascades down from these large scales and is dissipated at the small scales. At some point in the range between the large and small scales, the eddies become isotropic and have properties that are universal for high-enough Reynolds numbers. It is this isotropy and universality that enables Large-Eddy Simulations (LES). In the LES approach, a low-pass filter is applied to the governing equations, removing the

scales smaller than the filter cutoff. The effect of the removed scales on the rest of the flow must be modeled. The universality of the small scales suggests that models may be used in all flows with Reynolds numbers high enough for the Kolmogorov theory to apply.

To derive the LES governing equations, each variable in the original governing equations is decomposed into a resolved component and an unresolved, or sub-grid, component, $\phi = \bar{\phi} + \phi''$ where the overline ($\bar{\cdot}$) indicates the resolved component and the double prime (ϕ'') is the unresolved component for every variable ϕ . This separation is achieved by convolving a spatial filter kernel $\mathcal{G}(t, \mathbf{x})$ and the variable $\phi(t, \mathbf{x})$ such that:

$$\bar{\phi}(t, \mathbf{x}) = \mathcal{G}(t, \mathbf{x}) \star \phi(t, \mathbf{x}) = \iint_{t, \Xi} \phi(t', \boldsymbol{\xi}) \mathcal{G}(t - t', \mathbf{x} - \boldsymbol{\xi}) dt' d^3 \boldsymbol{\xi} \quad (2.34)$$

This filter may be explicitly defined and the convolution may be performed at each time step in the simulation. This is known as explicitly filtered LES and is attractive because the filtering is independent of the mesh; therefore, grid convergence can be measured. From a mathematical standpoint, it also ensures the discretized LES equations are consistent and when solved with a stable numerical scheme, the approach is convergent. However, the explicit filtering operation is expensive and complicated on curvilinear grids, limiting the use of the technique to academic configurations.

Instead of defining an explicit filter, it may be assumed that the filter is defined by the mesh and numerical scheme. This is known as implicitly filtered LES, and rather than performing a filtering operation on the solution, one solves directly for the filtered variables. There is no additional cost due to the filtering in the scheme. Unfortunately, the discrete form of the implicitly filtered LES equations is not consistent and measuring grid independence through mesh refinement is an open challenge. Despite this drawback, implicitly filtered LES is a popular approach.

The filter definition is assumed to be a band-pass filter in space, also known as a top-hat filter, which for three dimensional filters may be approximated as the product of one

dimensional filters:

$$\mathcal{G}(t, \mathbf{x}) = \prod_i^3 g_i(t - t', x_i - x'_i) \quad (2.35)$$

These one dimensional filters are:

$$g_i(t, x_i) = \begin{cases} \frac{1}{\Delta_i} & \text{if } |x_i - x'_i| < \frac{\Delta_i}{2} \\ 0 & \text{otherwise} \end{cases} \quad (2.36)$$

where Δ_i is the length of the computational cell in the i -th direction. This gives a total filter size of $\bar{\Delta} = (\Delta_1 \Delta_2 \Delta_3)^{1/3}$. In reality, the filter shape is only approximately a top-hat as it relies on the numerical scheme for its actual definition and behavior [1]. Lastly, the filter sizes vary throughout the domain as the cell sizes change; for the cases considered here, the mesh is static in time and so $\mathcal{G}(t, \mathbf{x}) = \mathcal{G}(\mathbf{x})$ [88]. It is also assumed that the filter commutes with differentiation. This introduces third-order leading terms in the truncation error, making the scheme only second-order accurate [88]. The scheme used in this work is also second-order accurate, so these errors are comparable with the scheme truncation errors. It also requires that meshes vary smoothly in space – therefore, all meshes with non-uniform grids are created with less than 10% stretching in any given direction whenever possible.

As a final note on the filtering procedure, application of the filtering as defined to the continuity equation creates a source term. This source term is undesirable as the numerical scheme used to implement it is only approximate, leading to errors in mass conservation. This source term may be eliminated by considering the Favré averaged variables, $\tilde{\phi} = \overline{\rho\phi}/\bar{\rho}$. The resulting spatially-filtered, multi-component, reacting LES equations using Favré averaged variables are:

$$\frac{\partial}{\partial t} \iiint_{\Omega} \tilde{\mathbf{W}} \, dV + \iint_{\partial\Omega} [\tilde{\mathbf{F}} - \tilde{\mathbf{G}}] \cdot d\mathbf{A} = \tilde{\mathbf{S}} \quad (2.37)$$

where

$$\begin{aligned}
 \widetilde{\mathbf{W}} &= \begin{pmatrix} \bar{\rho} \\ \bar{\rho}\tilde{u} \\ \bar{\rho}\tilde{v} \\ \bar{\rho}\tilde{w} \\ \bar{\rho}\tilde{E} \\ \bar{\rho}k^{sgs} \\ \bar{\rho}\tilde{Y}_k \end{pmatrix}, & \widetilde{\mathbf{F}} &= \begin{pmatrix} \bar{\rho}\tilde{u} \\ \bar{\rho}\tilde{u}\tilde{u} + \bar{p}\hat{i} \\ \bar{\rho}\tilde{u}\tilde{v} + \bar{p}\hat{j} \\ \bar{\rho}\tilde{u}\tilde{w} + \bar{p}\hat{k} \\ \tilde{\mathbf{u}}(\bar{\rho}\tilde{E} + \bar{p}) + \mathbf{H}^{sgs} \\ \bar{\rho}\tilde{\mathbf{u}}k^{sgs} \\ \bar{\rho}\tilde{Y}_k\tilde{\mathbf{u}} \end{pmatrix}, \\
 \widetilde{\mathbf{G}} &= \begin{pmatrix} 0 \\ \bar{\tau}_{xi} - \tau_{xi}^{sgs} \\ \bar{\tau}_{yi} - \tau_{yi}^{sgs} \\ \bar{\tau}_{zi} - \tau_{zi}^{sgs} \\ \tilde{u}_j\bar{\tau}_{ij} + \tilde{\mathbf{q}} + \boldsymbol{\sigma}^{sgs} \\ \mathcal{T}_{k^{sgs}} \\ \tilde{\mathbf{J}}_k + \mathbf{Y}_k^{sgs} + \boldsymbol{\theta}_k^{sgs} \end{pmatrix}, & \widetilde{\mathbf{S}} &= \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ pd_{k^{sgs}} + P_{k^{sgs}} - D_{k^{sgs}} \\ \bar{\omega}_k \end{pmatrix}
 \end{aligned} \tag{2.38}$$

the description of which is in the following sections and is summarized from the work of Génin [62] and Masquelet [63].

2.5.1 Conservation of mass

The filtered conservation of mass equation has no additional terms and appears very similar to the unfiltered equation due to the Favré averaging. The spatially filtered density is $\bar{\rho}$ and the Favré averaged velocity is $\tilde{\mathbf{u}}$.

2.5.2 Conservation of momentum

The inviscid flux tensor in the filtered conservation of momentum equations also looks similar to its unfiltered counterpart. The pressure is now the spatially filtered pressure, \bar{p} .

The viscous flux tensor contains new terms, however. There is the resolved, filtered stress tensor $\bar{\tau}_{ij}$, which based on the assumptions made in the original Navier-Stokes equations, is:

$$\bar{\tau}_{ij} = 2\tilde{\mu} \left(\tilde{S}_{ij} - \frac{1}{3}\tilde{S}_{kk}\delta_{ij} \right) \quad (2.39)$$

where $\tilde{S}_{ij} = 1/2 \left(\frac{\partial \tilde{u}_i}{\partial x_j} + \frac{\partial \tilde{u}_j}{\partial x_i} \right)$ is the resolved rate-of-strain tensor and $\tilde{\mu} \left(\tilde{T}, \tilde{Y}_k \right)$ is the molecular viscosity as a function of the Favré averaged temperature and species composition. The new stress tensor, τ_{ij}^{sgs} , represents the stress introduced by the unresolved scales and is defined as:

$$\tau_{ij}^{sgs} = \bar{\rho} (\widetilde{u_i u_j} - \tilde{u}_i \tilde{u}_j) \quad (2.40)$$

This term contains the unknown filtered quantity $\widetilde{u_i u_j}$ and must be modeled. The model used here employs an eddy-viscosity closure using the subgrid kinetic energy and the assumed filter width:

$$\nu_t = c_\nu \bar{\Delta} \sqrt{k^{sgs}} \quad (2.41)$$

where k^{sgs} is the resolved subgrid kinetic energy. The coefficient c_ν may be held constant, as is done for all cases here, or may be determined dynamically using the localized dynamic k^{sgs} model (LDKM) [89]. The constant is $c_\nu = 0.067$. Using this eddy viscosity, the subgrid shear stress is:

$$\tau_{ij}^{sgs} = -2\bar{\rho}\nu_t \left(\tilde{S}_{ij} - \frac{1}{3}\tilde{S}_{kk}\delta_{ij} \right) + \frac{2}{3}\bar{\rho}k^{sgs}\delta_{ij} \quad (2.42)$$

2.5.3 Conservation of energy

The energy equation contains several modifications. The total energy is now defined as:

$$\tilde{E} = \tilde{e} + \frac{1}{2}\tilde{\mathbf{u}} \cdot \tilde{\mathbf{u}} + k^{sgs} \quad (2.43)$$

where \tilde{e} is the Favré averaged internal energy. The heat flux vector is now defined to be:

$$\tilde{\mathbf{q}} = \tilde{\kappa} \nabla \tilde{T} - \bar{\rho} \sum_k \tilde{h}_k \tilde{Y}_k \tilde{\mathbf{V}}_k - \sum_k \mathbf{q}_k^{sgs} \quad (2.44)$$

where $\tilde{\kappa}(\tilde{T}, \tilde{Y}_k)$ is the heat diffusion coefficient at a given temperature and composition and \mathbf{q}_k^{sgs} is the subgrid heat transfer due to unresolved turbulent diffusion of species. The species diffusion velocity $\tilde{\mathbf{V}}_k$ is computed using Equation (2.12) using the Favré averaged variables.

The exact expression for the subgrid heat transfer due to turbulent diffusion of species is:

$$\mathbf{q}_k^{sgs} = \bar{\rho} \left(\widetilde{Y_k h_k \mathbf{V}_k} - \tilde{Y}_k \tilde{h}_k \tilde{\mathbf{V}}_k \right) \quad (2.45)$$

This term, along with the subgrid fluctuations in state variables, are often neglected and this is the approach taken in this work as well [89, 90]. Two additional new terms appear in the energy equation, one in the inviscid flux tensor and another in the viscous. The inviscid, subgrid transport of enthalpy is defined as:

$$\mathbf{H}^{sgs} = -\bar{\rho} \left(\widetilde{E \mathbf{u}} - \tilde{E} \tilde{\mathbf{u}} \right) + (\overline{\mathbf{u} p} - \tilde{\mathbf{u}} \tilde{p}) \quad (2.46)$$

and the subgrid heating due to viscous dissipation is:

$$\boldsymbol{\sigma}^{sgs} = (\overline{u_j \tau_{ij}} - \tilde{u}_j \tilde{\tau}_{ij}) \quad (2.47)$$

These two terms are modeled together:

$$\mathbf{H}^{sgs} - \boldsymbol{\sigma}^{sgs} = (\bar{\rho} \nu_t + \tilde{\mu}) \nabla k^{sgs} + \frac{\bar{\rho} \nu_t c_p}{\text{Pr}_t} \nabla \tilde{T} - \tilde{u}_j \tilde{\tau}_{ij}^{sgs} \quad (2.48)$$

where Pr_t is the turbulent Prandtl number.

2.5.4 Conservation of subgrid kinetic energy

The subgrid kinetic energy equation contains four terms: $\mathcal{T}_{k^{sgs}}$ is the diffusion of k^{sgs} , $pd_{k^{sgs}}$ is the pressure-dilation correlation, $P_{k^{sgs}}$ is the production term and $D_{k^{sgs}}$ is the dissipation term. The exact form of the terms in the conservation equation for the subgrid kinetic energy are:

$$\begin{aligned}
 \mathcal{T}_{k^{sgs}} &= - \left(\left(\overline{\rho K \tilde{\mathbf{u}}} - \bar{\rho} \tilde{K} \tilde{\mathbf{u}} \right) - \tilde{\mathbf{u}} \tau_{ij}^{sgs} + (\overline{\mathbf{u} p} - \tilde{\mathbf{u}} \bar{p}) - (\overline{\mathbf{u} \tau_{ij}} - \tilde{\mathbf{u}} \bar{\tau}_{ij}) \right) \\
 pd_{k^{sgs}} &= \overline{p \nabla \cdot \mathbf{u}} - \bar{p} \nabla \cdot \tilde{\mathbf{u}} \\
 P_{k^{sgs}} &= -\tau_{ij}^{sgs} \nabla \tilde{\mathbf{u}} \\
 D_{k^{sgs}} &= (\overline{\tau_{ij} \nabla \mathbf{u}} - \bar{\tau}_{ij} \nabla \tilde{\mathbf{u}})
 \end{aligned} \tag{2.49}$$

where $K = 1/2 \mathbf{u} \cdot \mathbf{u}$ is the total kinetic energy. No additional modeling is required for the production term as all components are known from the solution or existing closures. The pressure-dilation correlation $pd_{k^{sgs}}$ is often neglected and is small even for supersonic flows with strong non-equilibrium and so it is neglected here as well [89].

The remaining terms require closures following the work of Génin and Menon [89]. The model for dissipation of subgrid kinetic energy relies on the filter size and a characteristic subgrid velocity represented by the subgrid kinetic energy itself:

$$D_{k^{sgs}} = \frac{\bar{\rho} c_\epsilon (k^{sgs})^{3/2}}{\Delta} \tag{2.50}$$

The coefficient c_ϵ is taken to be a constant for all simulations, with a value $c_\epsilon = 0.916$.

The diffusion of subgrid kinetic energy contains several unclosed terms. The contribution due to pressure fluctuations is often neglected in low Mach number and low compressibility situations [89]. However, in a general combustion case, some regions of the flow may not have negligible subgrid pressure fluctuations and the terms are retained. The model for these

terms is:

$$\overline{\mathbf{u}p} - \tilde{\mathbf{u}}\tilde{p} = -\frac{\bar{\rho}\nu_t\tilde{R}}{\text{Pr}_t}\nabla\tilde{T} \quad (2.51)$$

The velocity-kinetic energy correlation is modeled using a gradient diffusion hypothesis based on the subgrid kinetic energy. The resulting model for the diffusion of subgrid kinetic energy is:

$$\mathcal{T}_{k^{sgs}} = (\bar{\rho}\nu_t + \mu)\nabla k^{sgs} + \frac{\bar{\rho}\nu_t\tilde{R}}{\text{Pr}_t}\nabla\tilde{T} \quad (2.52)$$

2.5.5 Conservation of partial densities

The conservation equation for the partial densities introduces several new terms. The species diffusion flux is similar to its unfiltered analogue, where now the terms are evaluated using the Favré averaged variables:

$$\tilde{\mathbf{J}}_k = \bar{\rho}\tilde{Y}_k\tilde{\mathbf{V}}_k \quad (2.53)$$

Two new subgrid terms are introduced; the subgrid mass flux:

$$\mathbf{Y}_k^{sgs} = -\bar{\rho}\left(\widetilde{\mathbf{u}Y_k} - \tilde{\mathbf{u}}\tilde{Y}_k\right) \quad (2.54)$$

and the subgrid species diffusion flux:

$$\boldsymbol{\theta}_k^{sgs} = -\bar{\rho}\left(\widetilde{\mathbf{V}_k Y_k} - \tilde{\mathbf{V}}_k\tilde{Y}_k\right) \quad (2.55)$$

The subgrid species diffusion flux is neglected in this work while the subgrid mass flux is modeled using a gradient diffusion assumption:

$$\mathbf{Y}_k^{sgs} = \frac{\bar{\rho}\nu_t}{\text{Sc}_t}\nabla\tilde{Y}_k \quad (2.56)$$

where Sc_t is the turbulent Schmidt number. Finally, the reaction source terms are filtered and models for the interaction between the chemistry and the turbulence are described in the next section.

2.5.6 Chemistry-Turbulence Closure

Prior to detailing the closure model for chemistry-turbulence interactions, it is important to discuss the structure of flames and how turbulence may alter that structure. All of the reacting results in Chapter 7 and Chapter 8 are premixed, either completely premixed with no compositional variations, or partially premixed with compositional inhomogeneities or stratification. Therefore, discussion of the flames and interactions with turbulence are confined to premixed flames. Figure 2.1 illustrates the structure of a typical premixed, laminar flame. There are two primary regions between the burned and unburned gases. The first is the reaction zone, where a majority of the chemical reactions take place and the heat release rates are the largest. This zone is typically thin and its thickness is denoted δ_r . Within this zone, radical species are created and diffuse outwards from the reaction zone. On the burned gas side, these radicals recombine to form the stable product species.

In the preheat zone, the radicals transported from the reaction zone participate in the initiation and chain branching chemical reactions that provide the exothermic processes needed to sustain the reaction zone. The energy for these reactions is provided by thermal conductivity and radiation from the reaction zone. For the cases considered here, the radiative effects are not considered and radiative transport is neglected. The flame thickness, δ_F , is the size preheat zone and reaction zone together and is often called the thermal thickness. Simulations using reduced kinetics in Chapter 7 and Chapter 8 use premixed hydrogen or premixed methane as fuel sources. The main radicals in the reaction zone are atomic hydrogen and atomic oxygen, and their diffusion into the preheat zone begins the chain reactions that produce the hydroxyl radical. This is the primary source of heat release in this region. Capturing the correct flame evolution requires accurately capturing both the reaction




Figure 2.1: Schematic of a laminar, premixed flame structure where δ_F is thermal thickness, and δ_r is the thickness of the reaction zone; thermal conductivity and species diffusion transport energy from the primary reaction zone into the unburned gases, creating the preheat zone.

processes and diffusion processes and this requirement imposes constraints on the chemical integration schemes as discussed in Chapter 6. Also indicated in Figure 2.1 is a dashed line through the middle of the reaction zone. This line is artificially defined as the flame surface and is useful for several models. In the thickened flame model [91–96] used in this work and described below, this flame surface is computed based on user-specified criteria and used in sensors to detect the flame location. This applies the thickening dynamically and localized to the flame.

The interaction between the turbulence and the flame depends upon the relationship between the turbulent length scales and velocities and the flame length scales and velocities [45]. Again confining the discussion to premixed or partially premixed flames, there are two turbulent length scales of interest: the integral length scale representing the largest eddies in the flow, l_t , and the Kolmogorov length scale representing the smallest eddies in the flow, η . Likewise, the turbulent velocity scale is characterized by u' and the flame velocity scale is characterized by the laminar flame speed, S_L . These may be related to the flame thickness,

δ_F , through the Damköhler number to find the ratio of the largest turbulent time scales to the chemical time scales:

$$\text{Da} = \frac{l_t/u'}{\delta_F/S_L} \quad (2.57)$$

The relation between the smallest time scales and the chemical time scales is the Karlovitz number:

$$\text{Ka} = \left(\frac{\delta_F}{\eta} \right)^2 \quad (2.58)$$

Finally, a turbulent Reynolds number may be defined using these definitions:

$$\text{Re}_t = \text{Da}^2 \text{Ka}^2 \quad (2.59)$$

These non-dimensional numbers define several regimes of premixed combustion [45]. When the chemical time scale is smaller than any turbulent time scale and the flame is thinner than the smallest turbulent eddies ($\text{Ka} < 1$), the internal behavior reaction zone of the flame cannot be altered by the turbulence and behaves as a laminar flame. The entire flame, including the preheat zone, is thin and this regime is called the *flamelet* regime. The flame may be wrinkled by the turbulence and if $u' < S_L$, the wrinkling is minor and the regime is called the *wrinkled flamelet* regime; on the other hand, if $u' > S_L$, the largest turbulent eddies may transport burned or unburned gas into the reaction zone, creating a *corrugated flamelet*. Under the conditions that $\text{Ka} > 1$ and $\text{Da} > 1$, the largest turbulent time scales are larger than the chemical time scales, but $\eta < \delta_F$ and the small scale eddies can enter and interact with the reaction zone structure. In this *thin reaction zone* regime, the flame is still smaller than the largest eddies and may be wrinkled by them, but there is an additional interaction between the chemistry and the turbulence at the smallest scales, which requires modeling in LES. For larger values of Ka , the turbulence may be significant enough to lead to localized quenching and this regime is known as the *broken reaction zone*. Figure 2.2 illustrates these regimes in a modified Boghi diagram as proposed by Peters [45].

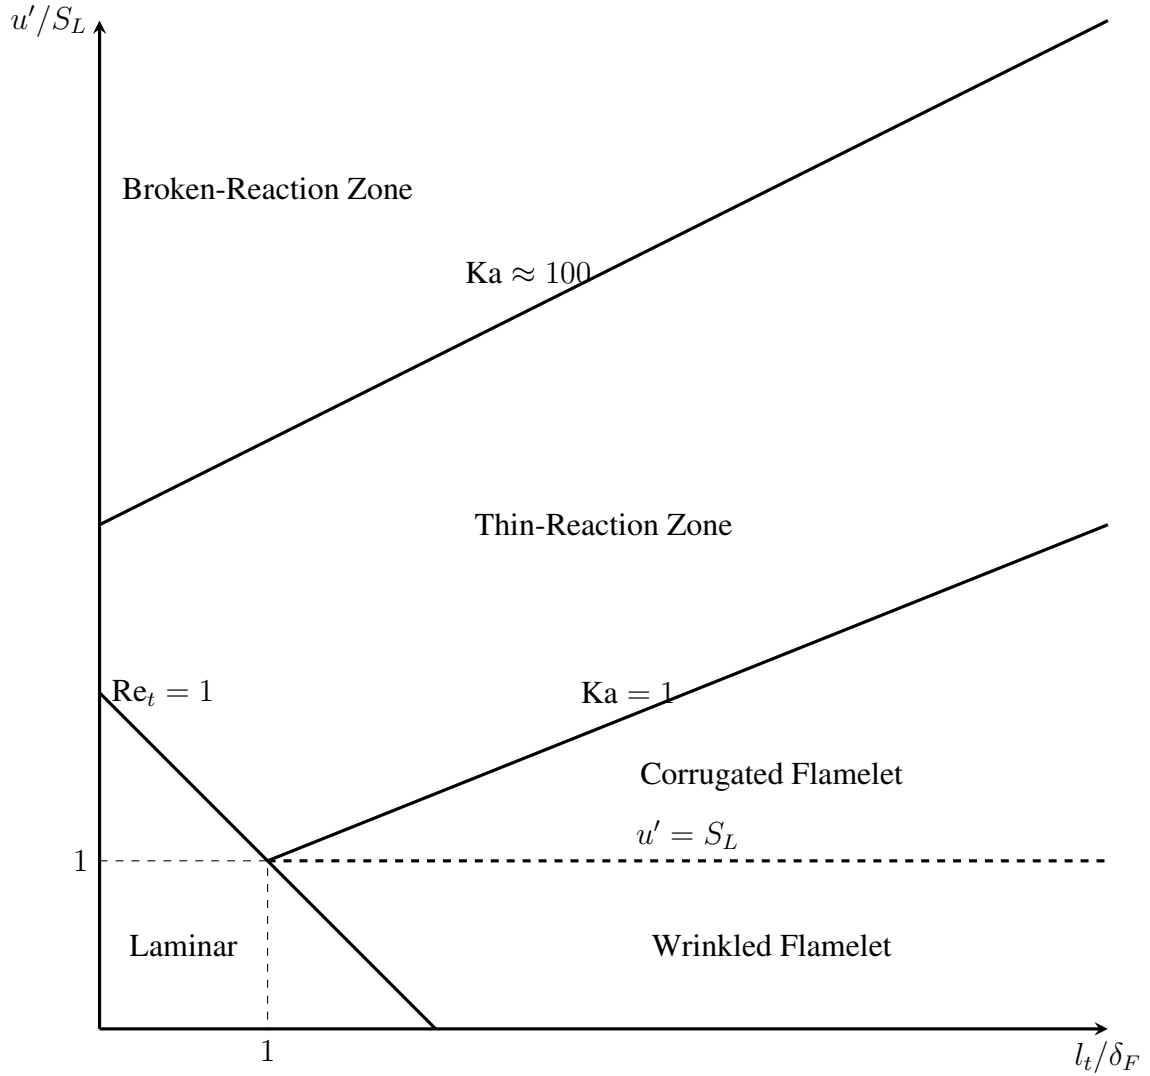


Figure 2.2: Peters' modified Borghi diagram showing the regimes of premixed, turbulent combustion in log-log scale [45].

In the flamelet regime, where the interaction between the turbulence and the chemistry is negligible, it is possible to omit any modeling terms at all and use:

$$\bar{\omega}_k = \dot{\omega}_k(\tilde{Y}_k, \bar{\rho}, \tilde{T}) \quad (2.60)$$

This no-modeling approach is frequently called the *quasi-laminar* chemistry approximation and may give acceptable results under certain conditions [97–99]. Alternatively, for regimes

where the quasi-laminar chemistry approximation is known not to hold, other methods to model the effects of turbulence on chemistry are required. The method used in this work to model these effects is the dynamically thickened flame (DTF) model [91, 92, 94–96].

From a numerical standpoint, the flame reaction zone presents several challenges, independent of its interaction with turbulence. It is often very thin and would require many grid points to properly resolve. It also contains the majority of the heat release in the flame, introducing strong, non-linear source terms that impact the robustness and efficiency of the scheme. These challenges can be overcome if the flame were thicker, such that it could be resolved on a typical grid designed for the underlying fluid mechanics using LES [45]. If a flame is assumed to be governed by a single, global reaction with rate $\dot{\omega}$, the laminar flame speed of that reaction is proportional to the square root of the thermal diffusivity D and the reaction rate [45]:

$$S_L \propto \sqrt{D\dot{\omega}} \quad (2.61)$$

and the flame thickness is:

$$\delta_F \propto \sqrt{\frac{D}{\dot{\omega}}} \quad (2.62)$$

If the thermal diffusivity is multiplied by a factor F and the reaction rate is divided by the same factor F , then the flame thickness is increased by F and the flame speed is unaltered.

However, changing the thickness of the flame also alters the relative ratios between the flame and the turbulent length scales. Specifically, $Da_{\text{thick}} = Da/F$ and the turbulence can influence the flame less than under the original conditions [45]. In order to restore the proper flame-turbulence interactions, an efficiency function E is introduced to account for the now-unresolved effects of the turbulence on the flame. This efficiency factor multiplies both the thermal diffusivity and the reaction rate, such that $\tilde{D} = EFD$ and $\tilde{\omega} = E\dot{\omega}/F$. In the full governing equations, this modification is done to all of the transport coefficients and

the reaction rates:

$$\begin{aligned}
\tilde{\kappa} &= \kappa \left(\tilde{Y}_k, \bar{\rho}, \tilde{T} \right) EF \\
\tilde{D}_k &= D_k \left(\tilde{Y}_k, \bar{\rho}, \tilde{T} \right) EF \\
\bar{\dot{\omega}}_k &= \frac{E\dot{\omega}_k \left(\tilde{Y}_k, \bar{\rho}, \tilde{T} \right)}{F}
\end{aligned} \tag{2.63}$$

Modifying the transport properties everywhere in the flow is undesirable, so sensors are used to detect the flame and apply the thickening and efficiency functions locally [93]. The sensor is designed for multi-step reactions and the user specifies the species to serve as flame indicators as well as the maximum reaction rate used for normalization. The sensor is defined as:

$$S = \sum_j^{N_s} \frac{|\dot{\omega}_j|}{\dot{\omega}_{j,\max}} \tag{2.64}$$

for N_s sensor species and the sensor is activated:

$$\min [S, 1] > \epsilon_{TF} \tag{2.65}$$

where ϵ_{TF} is a user-specified constant, taken to be 0.1 in all simulations. The thickening factor is determined based on the laminar flame thickness, the LES filter size (which, for implicitly-filtered LES is related to the local grid spacing), and a user-specified number of cells over which the flame should be resolved once thickened, n_F . The factor is computed at each cell as:

$$F = 1 + \left(\max \left[\frac{n_f \bar{\Delta}}{\delta_F}, 1 \right] - 1 \right) S \tag{2.66}$$

The efficiency function can take many different forms [91, 92]. The method recommended by Charlette et al. [92] known as the static power law function is the one used here. The

efficiency is given by:

$$E = \left(1 + \min \left[F - 1, \Gamma \left(\frac{u'}{S_L}, \frac{\Delta_e}{\delta_F} \right) \right] \frac{u'}{S_L} \right)^{1/2} \quad (2.67)$$

The function Γ takes as its arguments an estimate of the turbulence intensity relative to the laminar flame speed and the ratio of the effective filter size Δ_e and the flame thickness [92]:

$$\Gamma(K, D) = \left(\left((f_u^{-\alpha} + f_\delta^{-\alpha})^{-1/\alpha} \right)^{-\beta} + f_{\text{Re}}^{-\beta} \right)^{-1/\beta} \quad (2.68)$$

where the auxiliary functions are:

$$\begin{aligned} f_u &= 4 \left(\frac{27C_k}{110} \right)^{1/2} \frac{18C_k}{55} K^2 \\ f_\delta^2 &= \frac{27C_k\pi^{4/3}}{110} (D^{4/3} - 1) \\ f_{\text{Re}}^2 &= \text{Re}_\Delta \frac{9}{55} \exp \left[-\frac{3}{2} \frac{C_k\pi^{4/3}}{\text{Re}_\Delta} \right] \end{aligned} \quad (2.69)$$

where the Reynolds number at the effective filter size is:

$$\text{Re}_\Delta = 4KD \quad (2.70)$$

and the parameter α is:

$$\alpha = \frac{3}{5} + \frac{1}{5} \exp \left[-\frac{1}{10} K \right] - \frac{1}{5} \exp \left[-\frac{1}{10} D \right] \quad (2.71)$$

with constants $\beta = 1.4$ and $C_k = 1.5$. The velocity fluctuations at the effective filter level, u' , are computed using the subgrid kinetic energy as:

$$u' = \sqrt{2k^{sgs}} \quad (2.72)$$

To summarize, the DTF model will artificially thicken a flame so it is resolvable on a coarse (relative to the flame thickness) mesh. It does this by modifying the transport properties and reaction rate such that the flame speed is preserved; this, in turn, changes the effect of turbulence on the flame and an efficiency function is introduced to model the missing turbulence-chemistry effects. For this to work on a local sense, the user must specify the following:

- Marker species used to locate the flame surface; may be one or several.
- Maximum reaction rate for the sensor species.
- Cutoff size for sensor, above which the thickening is fully applied.
- Laminar flame speed thickness for the flame under the simulated conditions.
- Desired number of cells over which to spread the flame.

The flame speed, thickness and maximum reaction rates are functions of the equivalence ratio, and for partially premixed or stratified flows, there is not a single value that will work through all of the regimes. To account for this, these values may be interpolated from a lookup table indexed on equivalence ratio. Prior to computing the thickening or the efficiency functions, the local equivalence ratio is computed and the table is used to find suitable values of the flame speed, thickness, and maximum reaction rates.

2.6 Boundary Conditions

The full set of Navier-Stokes equations are second-order partial differential equations with five primary variables and the potential for many more depending on the LES models and the number of species tracked. The introduction of any physical geometry or the introduction of artificial surfaces used to truncate the domain into a finite region for numerical simulation requires an additional set of equations for these boundaries. At every boundary, each variable requires a specified condition and the velocity variables require two conditions, due to the

order of the equations. Unfortunately, unsteady, turbulent combustion processes are sensitive to acoustic, entropic, and vortical disturbances. The artificial truncation of the simulation domain will introduce errors and these errors may couple with the dynamics of the system, leading to the wrong solution or the complete failure of the scheme due to robustness issues. Modeling the effects of the physical geometry, such as walls, is relatively straight forward as will be shown. However, specification of inflow and outflow boundaries requires careful derivation. These boundary conditions must match, as closely as possible, the waves entering and departing the domain as the simulation evolves. This is achieved through analysis of the characteristic form of the governing equations whereby the appropriate waves may be imposed and the desired reflective and transmissive properties can be recovered.

There is flexibility in the choice of variables on which the boundary conditions are imposed. In this work, boundary conditions are imposed on a primitive variable set. Walls use the primitive set $\mathbf{U}_w = \{\bar{p}, \tilde{u}, \tilde{v}, \tilde{w}, \tilde{T}, k^{sgs}, \tilde{Y}_k\}^\top$ while inflows and outflows of all types use the primitive set $\mathbf{U}_{io} = \{\bar{p}, \tilde{u}, \tilde{v}, \tilde{w}, \bar{p}, k^{sgs}, \tilde{Y}_k\}^\top$. The distinction is made to facilitate the use of data from experiments, specifically mass-flow rates into or out of a rig and wall temperatures or heat flux rates measured. Generally speaking, a boundary or a portion of a boundary may impose three classes of conditions: Dirichlet, where the value of a variable is imposed on the boundary; Neumann, where the derivative of a variable is imposed on the boundary, or Robin, where there is a mix of derivatives and imposed values. The next sections will discuss how each type of boundary condition imposes values for the variables and the resulting implications.

2.6.1 Characteristic Inflows and Outflows

Thompson [100, 101] devised a method to formulate boundary conditions by decomposing the governing equations into their nonlinear, characteristic waves. Each wave propagates with a wave speed corresponding to the eigenvalues of the flux Jacobian. This information is used at inflows and outflows to determine, based on the local wave propagation direction,

whether a Neumann condition or a Dirichlet condition should be imposed on that particular characteristic variable. The analysis begins by writing the Navier-Stokes equations in a non-conservative, linearized form:

$$\frac{\partial \mathbf{U}_{io}}{\partial t} + \mathbf{A} \frac{\partial \mathbf{U}_{io}}{\partial x} + \mathbf{B} \frac{\partial \mathbf{U}_{io}}{\partial y} + \mathbf{C} \frac{\partial \mathbf{U}_{io}}{\partial z} + \mathbf{D} = 0 \quad (2.73)$$

where \mathbf{D} contains all of the source and viscous terms in the system. The flux Jacobians, \mathbf{A} , \mathbf{B} , and \mathbf{C} cannot be diagonalized simultaneously. Therefore, the system of equations is transformed into a coordinate system with a coordinate normal to the boundary, x_n , and transverse coordinates x_t and x_{tt} . The transverse flux terms are grouped together with \mathbf{D} :

$$\frac{\partial \mathbf{U}'_{io}}{\partial t} + \mathbf{A}' \frac{\partial \mathbf{U}'_{io}}{\partial x_n} + \mathbf{D}'' = 0 \quad (2.74)$$

where the rotated set of primitives is $\mathbf{U}'_{io} = \{\bar{\rho}, \tilde{u}_1, \tilde{u}_2, \tilde{u}_3, \bar{p}, k^{sgs}, \tilde{Y}_k\}^\top$ and $\mathbf{D}'' = \mathbf{B}' \frac{\partial \mathbf{U}'_{io}}{\partial x_t} + \mathbf{C}' \frac{\partial \mathbf{U}'_{io}}{\partial x_{tt}} + \mathbf{D}'$. The rotated flux Jacobian is:

$$\mathbf{A}' = \begin{bmatrix} \tilde{u}_1 & \bar{p} & 0 & 0 & 0 & 0 & 0 \\ 0 & \tilde{u}_1 & 0 & 0 & \frac{1}{\bar{\rho}} & 0 & 0 \\ 0 & 0 & \tilde{u}_1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \tilde{u}_1 & 0 & 0 & 0 \\ 0 & \bar{\rho}c & 0 & 0 & \tilde{u}_1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \tilde{u}_1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \tilde{u}_1 \end{bmatrix} \quad (2.75)$$

This Jacobian matrix may now be diagonalized with the right eigenvector matrix \mathbf{S} :

$$\mathbf{S}^{-1} \mathbf{A}' \mathbf{S} = \mathbf{\Lambda} \quad (2.76)$$

where Λ is the diagonal eigenvalue matrix with entries λ_i . There are two unique eigenvalues and the remaining eigenvalues are repeated, $\lambda_1 = \tilde{u}_1 - c$, $\lambda_5 = \tilde{u}_1 + c$, $\lambda_{i:i \notin \{1,5\}} = \tilde{u}_1$. This diagonalization may be substituted into Equation (2.74):

$$\mathbf{S}^{-1} \frac{\partial \mathbf{U}'_{io}}{\partial t} + \mathbf{L} + \mathbf{S}^{-1} \mathbf{D}'' = \mathbf{0} \quad (2.77)$$

with the wave amplitude vector $\mathbf{L} = \Lambda \mathbf{S}^{-1} \frac{\partial \mathbf{U}'_{iQ}}{\partial x_n}$:

$$\mathbf{L} = \begin{bmatrix} L_1 \\ L_2 \\ L_3 \\ L_4 \\ L_5 \\ L_6 \\ L_{6+k} \end{bmatrix} = \begin{bmatrix} \frac{\lambda_1}{2} \left(\frac{\partial \bar{p}}{\partial x_n} - \bar{\rho} c \frac{\partial \tilde{u}_1}{\partial x_n} \right) \\ \lambda_2 \left(\frac{\partial \bar{p}}{\partial x_n} - \frac{1}{c^2} \frac{\partial \bar{p}}{\partial x_n} \right) \\ \lambda_3 \frac{\partial \tilde{u}_2}{\partial x_n} \\ \lambda_4 \frac{\partial \tilde{u}_3}{\partial x_n} \\ \frac{\lambda_5}{2} \left(\frac{\partial \bar{p}}{\partial x_n} + \bar{\rho} c \frac{\partial \tilde{u}_1}{\partial x_n} \right) \\ \lambda_6 \frac{\partial k^{sgs}}{\partial x_n} \\ \lambda_{6+k} \frac{\partial \tilde{Y}_k}{\partial x_n} \end{bmatrix} \quad (2.78)$$

The original governing equations, Equation (2.37), are expressed in terms of the conservative variable vector $\tilde{\mathbf{W}}$ and the resulting boundary conditions need to be transformed into this variable set. Equation (2.77) is multiplied through by $\frac{\partial \tilde{\mathbf{W}}'}{\partial \mathbf{U}'_{io}} \mathbf{S}$ where $\tilde{\mathbf{W}}'$ is the conservative variable set in the rotated space and the transverse, viscous, and source terms

are neglected:

$$\frac{\partial \widetilde{\mathbf{W}}'}{\partial t} + \begin{bmatrix} d_4 \\ \tilde{u}_1 d_4 + \bar{\rho} d_1 \\ \tilde{u}_2 d_4 + \bar{\rho} d_2 \\ \tilde{u}_3 d_4 + \bar{\rho} d_3 \\ \frac{\partial \bar{\rho} \tilde{E}}{\partial \bar{\rho}} d_4 + \bar{\rho} \sum_{i=1}^3 \tilde{u}_i d_i + \frac{\partial \bar{\rho} \tilde{E}}{\partial \bar{\rho}} d_5 + \bar{\rho} d_6 \frac{\partial \tilde{E}}{\partial k^{sgs}} + \sum_k^{N_k} \bar{\rho} d_{6+k} \frac{\partial \tilde{E}}{\partial Y_k} \\ k^{sgs} d_4 + \bar{\rho} d_6 \\ \tilde{Y}_k d_4 + \bar{\rho} d_{6+k} \end{bmatrix} = \mathbf{0} \quad (2.79)$$

where:

$$\mathbf{d} = \mathbf{SL} = \begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \\ d_{6+k} \end{bmatrix} = \begin{bmatrix} \frac{L_5 - L_1}{\bar{\rho} c} \\ L_3 \\ L_4 \\ L_2 + \frac{L_5 + L_1}{c^2} \\ L_5 + L_1 \\ L_6 \\ L_{6+k} \end{bmatrix} \quad (2.80)$$

This system of boundary conditions are known as the Local One-Dimensional Inviscid (LODI) conditions [45, 102]. Based on this set of equations, multiple types of inflow and outflow may be specified.

Reflective Boundaries

Specifying the boundaries as reflective and providing target values for the velocity, temperature, subgrid kinetic energy, and gas composition on the inflow creates a boundary that

imposes the target values exactly. The resulting wave amplitudes are:

$$\begin{aligned} L_1 &= L_5 \\ L_3 &= L_4 = L_6 = L_{6+k} = 0 \end{aligned} \quad (2.81)$$

The L_2 wave is computed using a temperature form of the LODI equations:

$$L_2 = \frac{T\alpha_v^2}{c_p} (L_5 + L_1) \quad (2.82)$$

where α_v is the coefficient of thermal expansion. For an ideal gas such as those considered here, this simplifies to:

$$L_2 = \frac{\gamma - 1}{c^2} (L_1 + L_5) \quad (2.83)$$

where γ is the ratio of specific heats. For inflows, the L_5 wave is the only one exiting the domain and must therefore be computed using information from inside the domain:

$$L_5 = (\tilde{u}_1 + c) \left(\frac{\partial \bar{p}}{\partial x_n} + \bar{\rho} c \frac{\partial \tilde{u}_1}{\partial x_n} \right) \quad (2.84)$$

On the outflow, all waves except the L_1 wave are exiting the domain and must be computed based on the interior conditions:

$$\begin{aligned} L_2 &= \tilde{u}_1 \left(\frac{\partial \bar{p}}{\partial x_n} - \frac{1}{c^2} \frac{\partial \bar{p}}{\partial x_n} \right) \\ L_3 &= \tilde{u}_1 \frac{\partial \tilde{u}_2}{\partial x_t} \\ L_4 &= \tilde{u}_1 \frac{\partial \tilde{u}_3}{\partial x_{tt}} \\ L_5 &= \frac{\tilde{u}_1 + c}{2} \left(\frac{\partial \bar{p}}{\partial x_n} + \bar{\rho} (\tilde{u}_1 + c) \frac{\partial \tilde{u}_1}{\partial x_n} \right) \\ L_6 &= \tilde{u}_1 \frac{\partial k^{sgs}}{\partial x_n} \\ L_{6+k} &= \tilde{u}_1 \frac{\partial \tilde{Y}_k}{\partial x_n} \end{aligned} \quad (2.85)$$

with the incoming wave specified as:

$$L_1 = -L_5 \quad (2.86)$$

It is also possible to impose mass-flow rate values on the inflow rather than primitive values directly. This is done by requiring

$$L_5 (\tilde{u}_1 + c) + L_1 (\tilde{u}_1 - c) = -\tilde{u}_1 c^2 L_2 \quad (2.87)$$

The end result is:

$$L_5 = L_1 \frac{1 - M - (\tilde{u}_1 c T \alpha_V^2) / c_p}{1 + M + (\tilde{u}_1 c T \alpha_v^2) / c_p} \quad (2.88)$$

which for an ideal gas is:

$$L_5 = L_1 \frac{1 - \gamma M}{1 + \gamma M} \quad (2.89)$$

The transverse waves become:

$$\begin{aligned} L_3 &= -\frac{\tilde{u}_2}{\tilde{\rho}} \left(L_2 + \frac{L_5 + L_1}{c^2} \right) \\ L_4 &= -\frac{\tilde{u}_3}{\tilde{\rho}} \left(L_2 + \frac{L_5 + L_1}{c^2} \right) \end{aligned} \quad (2.90)$$

with all other waves computed as before.

As the name suggests, when a wave is incident on the boundary, these conditions will generate a reflected wave. This may interact with the rest of the flow and influence the solution in undesirable ways. The modifications presented in the next section allow the transmission and reflection properties of the boundaries to be tuned as desired.

Non-reflective/partially-reflective Boundaries

The amount of reflectivity can be adjusted by permitting the boundary condition to relax to the target value over time using a relaxation coefficient K . In this case, the desired response to impose a variable is $\phi_{bc} = \phi_{\text{target}} - e^{-Kt} (\phi - \phi_{\text{target}})$. The reflective conditions may be recovered as $K \rightarrow \infty$; conversely, a perfectly non-reflective boundary that permits waves to exit without any reflection occurs for $K = 0$. Values in between permit various amounts of wave transmission and reflection. The transverse and scalar waves are computed as before:

$$\begin{aligned} L_1 &= L_5 \\ L_3 &= L_4 = L_6 = L_{6+k} = 0 \end{aligned} \tag{2.91}$$

where now:

$$\begin{aligned} L_2 &= -\bar{\rho}\alpha_v K \left(\tilde{T} - \tilde{T}_{\text{target}} \right) \\ L_5 &= -\bar{\rho}c \frac{1 - M^2}{2} K \left(\tilde{u}_1 - \tilde{u}_{1,\text{target}} \right) \end{aligned} \tag{2.92}$$

As before, a constant mass version may be imposed:

$$L_5 = c \frac{1 - M^2}{2} K \left(\bar{\rho}\tilde{u}_1 - (\bar{\rho}\tilde{u}_1)_{\text{target}} \right) \tag{2.93}$$

The outflow is made non-reflective by setting:

$$L_1 = K \left(\bar{p} - \bar{p}_{\text{target}} \right) \tag{2.94}$$

2.6.2 Simple Inflows and Outflows

The simplest boundary conditions used for simulating fluid flows appear at what are often called supersonic inflows or supersonic outflows. In the traditional use, a supersonic inflow is a Dirichlet boundary, where the entire variable set U_{io} is imposed to the desired target

values. This type of boundary can lead to wave reflections and oscillations, but it is robust and ensures the desired values are rigidly imposed. On the other hand, supersonic outflows impose Neumann conditions on each variable such that $dU'_{io}/dx_n = 0$, where x_n is the coordinate normal to the boundary. Like the supersonic inflow, this boundary will generate reflections when a wave impinges upon it.

Although simple to impose, these boundaries do not follow the physics of the flow as well as the characteristic boundaries when used for subsonic flows. At inflows, imposing the Dirichlet condition on all of the variables is consistent with all of the characteristic waves entering the domain whereas the physics of the flow has one wave leaving the domain and the rest entering. Likewise, the Neumann condition at the outflow is consistent with all waves leaving the domain when in reality, one should be entering. The inflow discrepancy is not generally a problem aside from generating strong pressure wave reflections. The outflow discrepancy however leads to a loss of pressure in the domain when used for subsonic flows because there is no mechanism for the pressure to relax to the ambient values outside the domain.

This discrepancy may be addressed by making minor changes to the boundary conditions and making both inflows and outflows Robin boundaries. On the inflow, the velocity, temperature, and gas composition are imposed using Dirichlet conditions as before. However, the pressure is switched to a Neumann condition with $\partial\bar{p}/\partial x_n = 0$. This mimics the effect of the pressure waves exiting the domain, although it will still generate reflections from incident waves. Similarly, the outflow is modified so that all variables except pressure are treated as Neumann conditions with zero gradient. The pressure is treated as a Dirichlet boundary with the target value imposed rigidly.

Darmofal, Moinier and Giles performed an eigenmode analysis of this type of boundary condition to assess the performance with the preconditioned and non-preconditioned system [103]. The analysis decomposes the linearized Euler equations into their characteristic form and solves for the time-dependent eigenmodes. From these modes, the frequencies

are calculated and the behavior of the boundary conditions is determined. For the Robin boundary conditions with the non-preconditioned system of equations, the frequency is [103]:

$$\omega = (1 - M^2)(n + \frac{1}{2})\pi; \quad n \in \mathbb{Z} \quad (2.95)$$

This frequency has no imaginary component and therefore the boundary conditions do not dissipate any disturbances [103]. Both the Robin inflow and outflow are perfectly reflective with the non-preconditioned system.

2.6.3 Walls

The application of the wall boundary conditions is straight-forward, but must be done in the proper order such that all of the variables are consistent. The walls are applied by setting values in the ghost cells that are designed to impose the proper values on the cell face where the wall is located. Consistency in the values is maintained by setting the conditions in the following order:

1. Wall temperature is set based on the desired condition (e.g., adiabatic, isothermal, or heat flux).
2. Pressure is set assuming a zero-gradient normal to the wall.
3. Gas composition is set assuming a zero-gradient normal to the wall.
4. Density and internal energy are computed using the imposed temperature, pressure, and composition conditions.
5. Velocities are computed. The entire wall may or may not be moving in a specified direction. In addition to the prescribed wall motion, the velocities may have additional contributions based on the desired conditions:
 - The normal velocity is set based on several possible configurations. For most simulations, the normal velocity is zero. However, wall-effusion or surface bur-

ning may be modeled by imposing a wall-normal velocity. Likewise, vibrating walls may be created by imposing a wall-normal velocity based on vibration amplitudes and frequencies.

- The tangential velocities are set based on whether the wall is a slip or no-slip wall. For slip walls, the tangential velocities use a zero-gradient condition. For no-slip walls, the tangential velocities are mirrored such that they are zero on the wall and any prescribed motion is then added.
- The tangential velocities are modified based on any effusion or surface burning that may occur at an angle.

6. Update all of the conservative variables for use in the flux routines, with careful attention paid to the energy term to ensure everything is consistent with the imposed primitive values.

The methods for applying each condition are described in the following sections. The subscripts w , i and g indicate the wall value on the cell face, the interior value in the cell center next to the wall, and the ghost value in the cell center next to the wall.

Temperature conditions

In the simplest case of adiabatic walls, the condition imposed is $\tilde{T}_g = \tilde{T}_i$. However, most experiments do not have adiabatic boundary conditions on the walls and the heat exchange from the flow to the wall and into the ambient environment or into the cooling apparatus may dramatically change the dynamics of the flow. For long-running experiments, the walls may be assumed isothermal once the heat through the wall from the domain reaches a steady state. This can be imposed through the wall boundary conditions by specifying a wall temperature \tilde{T}_w and setting:

$$\tilde{T}_g = 2\tilde{T}_w - \tilde{T}_i \quad (2.96)$$

For rigs with more transient conditions, the heat flux through the walls may be specified. This is implemented by solving a simple heat conduction equation normal to the wall. The heat flux, q , may be constant or a profile in space and/or time:

$$\begin{aligned}\tilde{T}_w &= \frac{1}{2} \left[2\tilde{T}_i - \frac{qS_n}{\kappa_i} \right] \\ \tilde{T}_g &= 2\tilde{T}_w - \tilde{T}_i\end{aligned}\tag{2.97}$$

where S_n is the surface area of the cell normal to the wall.

Gas state conditions

With the temperature condition set appropriately, the remaining state of the gas is determined. The pressure is assumed to be zero-gradient:

$$\bar{p}_g = \bar{p}_i\tag{2.98}$$

and the same is true for the species mass fractions:

$$\tilde{Y}_{k,g} = \tilde{Y}_{k,i}\tag{2.99}$$

With the gas temperature, species and pressure set, the gas density is computed using the equation of state. At this stage, the internal energy is also computed as well as all of the gas thermodynamic properties.

Velocity conditions

The velocity conditions are set using either a slip or no-slip assumption for the transverse terms and the normal terms may be computed with or without vibrations or mass flux. The

transverse term behavior can be toggled by setting a factor f :

$$f = \begin{cases} 2 & \text{slip walls} \\ 0 & \text{no-slip walls} \end{cases} \quad (2.100)$$

In the case of vibrating walls, the normal velocity fluctuation is specified using a sum of user-defined periodic functions:

$$\tilde{V}'_n = \sum_{i=1}^{N_m} A_i \sin(2\pi f_i t) \quad (2.101)$$

where there are N_m user-specified modes with amplitudes A_i and frequencies f_i . The velocity components in the Cartesian reference frame are then computed:

$$\begin{aligned} \tilde{u}'_{\text{vib}} &= \tilde{V}'_n S_x \\ \tilde{v}'_{\text{vib}} &= \tilde{V}'_n S_y \\ \tilde{w}'_{\text{vib}} &= \tilde{V}'_n S_z \end{aligned} \quad (2.102)$$

where S_x , S_y , and S_z are the cell face areas in the X, Y and Z directions respectively. The velocity contributions from mass effusion or surface burning are computed using:

$$\begin{aligned} \tilde{u}'_{\text{effusion}} &= \tilde{V}'_{\text{effusion}} (S_x + T_x \cot \theta_{\text{effusion}}) \\ \tilde{v}'_{\text{effusion}} &= \tilde{V}'_{\text{effusion}} (S_y + T_y \cot \theta_{\text{effusion}}) \\ \tilde{w}'_{\text{effusion}} &= \tilde{V}'_{\text{effusion}} (S_z + T_z \cot \theta_{\text{effusion}}) \end{aligned} \quad (2.103)$$

where $\tilde{V}'_{\text{effusion}}$ is the prescribed effusion velocity or surface burning velocity, T_x , T_y , and T_z are the cell face areas tangential to their respective coordinate directions, and θ_{effusion} is the angle the effusion velocity makes to the wall. Two final sets of terms are needed for the wall

boundary conditions. First, the flow normal to the surface on the interior is needed:

$$\tilde{V}_n = \tilde{u}_i S_x + \tilde{v}_i S_y + \tilde{w}_i S_z \quad (2.104)$$

and the prescribed wall motion values \tilde{u}_w , \tilde{v}_w , and \tilde{w}_w , if any, are needed. Finally, the velocity components in the ghost cells to impose the desired wall boundary conditions are:

$$\begin{aligned} \tilde{u}_g &= 2\tilde{u}'_{\text{vib}} - 2\tilde{u}'_{\text{effusion}} + (f - 1)(-2\tilde{u}_w + \tilde{u}_i) - f\tilde{V}_n S_x \\ \tilde{v}_g &= 2\tilde{v}'_{\text{vib}} - 2\tilde{v}'_{\text{effusion}} + (f - 1)(-2\tilde{v}_w + \tilde{v}_i) - f\tilde{V}_n S_y \\ \tilde{w}_g &= 2\tilde{w}'_{\text{vib}} - 2\tilde{w}'_{\text{effusion}} + (f - 1)(-2\tilde{w}_w + \tilde{w}_i) - f\tilde{V}_n S_z \end{aligned} \quad (2.105)$$

The conservative variables are set based on the imposed gas state and velocities. The total energy in the ghost cell is recomputed based on the internal energy, subgrid kinetic energy, and resolved kinetic energy. With that complete, the entire system is determined and the wall boundary conditions are completely specified.

2.7 Preconditioning and Dual-Time

In Chapter 1, the numerical difficulties at low Mach numbers are traced to the disparity between the eigenvalues of the flux Jacobians in the system. Preconditioning is a technique used to modify the flux Jacobians in a way that reduces the ratio between the largest and smallest eigenvalues, thereby reducing stiffness and improving iterative performance. The physical governing equations are given in Equation (2.37), repeated here:

$$\frac{\partial}{\partial t} \iiint_{\Omega} \tilde{\mathbf{W}} \, dV + \iint_{\partial\Omega} [\tilde{\mathbf{F}} - \tilde{\mathbf{G}}] \cdot d\mathbf{A} = \tilde{\mathbf{S}} \quad (2.37)$$

The stiffness of the system may be altered by multiplying the temporal derivative by the preconditioning matrix based on the conservative variables Γ_C :

$$\Gamma_C \frac{\partial}{\partial \tau} \iiint_{\Omega} \widetilde{\mathbf{W}} \, dV + \iint_{\partial\Omega} [\widetilde{\mathbf{F}} - \widetilde{\mathbf{G}}] \cdot d\mathbf{A} = \widetilde{\mathbf{S}} \quad (2.106)$$

where now the temporal derivative is expressed in terms of a pseudo-time, τ , rather than a physical time, t . This change is required because altering the eigenvalues of the system changes the dynamics of the waves. In this case, it slows down the acoustic waves so their speed is comparable to the convective wave speed. The resulting system is no longer time accurate and indeed it is no longer the Navier-Stokes equations, except in the steady-state when the pseudo-time derivative is zero.

For the LES of turbulent, reacting flows, temporal accuracy and transient phenomenon are desired results and this loss of time accuracy is not suitable. Time accuracy can be restored by incorporating two temporal derivatives in the governing equations [5]:

$$\frac{\partial}{\partial \tau} \iint_{\Omega} \widetilde{\mathbf{W}} \, dV + \frac{\partial}{\partial t} \iiint_{\Omega} \widetilde{\mathbf{W}} \, dV + \iint_{\partial\Omega} [\widetilde{\mathbf{F}} - \widetilde{\mathbf{G}}] \cdot d\mathbf{A} = \widetilde{\mathbf{S}} \quad (2.107)$$

where a pseudo-time derivative has been added to Equation (2.37). In this system, the pseudo-time is iterated forward until $\tau \rightarrow \infty$, at which point the pseudo-time derivative is zero and Equation (2.107) is the same as Equation (2.37). Just as before, the eigenvalues of the system can be modified by multiplying the iterative time (in this case, pseudo-time) by the preconditioning matrix Γ_C :

$$\Gamma_C \frac{\partial}{\partial \tau} \iint_{\Omega} \widetilde{\mathbf{W}} \, dV + \frac{\partial}{\partial t} \iiint_{\Omega} \widetilde{\mathbf{W}} \, dV + \iint_{\partial\Omega} [\widetilde{\mathbf{F}} - \widetilde{\mathbf{G}}] \cdot d\mathbf{A} = \widetilde{\mathbf{S}} \quad (2.108)$$

Now the modified wave speeds are confined to the pseudo-time derivative, which as $\tau \rightarrow \infty$ becomes zero and the physical governing equations remain.

For low Mach number flows, it is often beneficial to work with a primitive variable set rather than the conservative variable set [3, 5]. The choice of primitive variables offers some flexibility. The pressure-temperature set, $\tilde{\mathbf{Q}} = \{\bar{p}, \tilde{u}, \tilde{v}, \tilde{w}, \tilde{T}, k^{sgs}, \tilde{Y}_k\}^\top$, is the most common choice and the one used here. This set of variables is related to the conservatives through the Jacobian matrix (written here for a multicomponent fluid with the one-equation subgrid kinetic energy model):

$$\frac{\partial \tilde{\mathbf{W}}}{\partial \tilde{\mathbf{Q}}} = \begin{bmatrix} \bar{\rho}_{\bar{p}} & 0 & 0 & 0 & \bar{\rho}_{\tilde{T}} & 0 & \bar{\rho}_{\tilde{Y}_j} \\ \bar{\rho}_{\bar{p}} \tilde{u} & \bar{\rho} & 0 & 0 & \bar{\rho}_{\tilde{T}} \tilde{u} & 0 & \bar{\rho}_{\tilde{Y}_j} \tilde{u} \\ \bar{\rho}_{\bar{p}} \tilde{v} & 0 & \bar{\rho} & 0 & \bar{\rho}_{\tilde{T}} \tilde{v} & 0 & \bar{\rho}_{\tilde{Y}_j} \tilde{v} \\ \bar{\rho}_{\bar{p}} \tilde{w} & 0 & 0 & \bar{\rho} & \bar{\rho}_{\tilde{T}} \tilde{w} & 0 & \bar{\rho}_{\tilde{Y}_j} \tilde{w} \\ \bar{\rho}_{\bar{p}} \tilde{H} - 1 & \bar{\rho} \tilde{u} & \bar{\rho} \tilde{v} & \bar{\rho} \tilde{w} & \bar{\rho}_{\tilde{T}} \tilde{H} + \bar{\rho} \tilde{H}_{\tilde{T}} & \bar{\rho} & \bar{\rho} \tilde{H}_{\tilde{Y}_j} + \bar{\rho}_{\tilde{Y}_j} \tilde{H} \\ \bar{\rho}_{\bar{p}} k^{sgs} & 0 & 0 & 0 & \bar{\rho}_{\tilde{T}} k^{sgs} & \bar{\rho} & \bar{\rho}_{\tilde{Y}_j} k^{sgs} \\ \bar{\rho}_{\bar{p}} \tilde{Y}_i & 0 & 0 & 0 & \bar{\rho}_{\tilde{T}} \tilde{Y}_i & 0 & \tilde{Y}_i \bar{\rho}_{\tilde{Y}_j} + \bar{\rho} \delta_{ij} \end{bmatrix} \quad (2.109)$$

where subscripts indicate differentiation, ie. $\bar{\rho}_{\bar{p}} = \frac{\partial \bar{\rho}}{\partial \bar{p}}$ and $\bar{\rho}_{\tilde{T}} = \frac{\partial \bar{\rho}}{\partial \tilde{T}}$, with the exception of i and j , which are index variables. Likewise, the source term vector in the reacting, LES governing equations, Equation (2.7), may be transformed from conservative to primitive variables by $\tilde{\mathbf{S}} = \mathbf{D} \tilde{\mathbf{Q}}$ where $\mathbf{D} = \frac{\partial \tilde{\mathbf{S}}}{\partial \tilde{\mathbf{Q}}}$:

$$\frac{\partial \tilde{\mathbf{S}}}{\partial \tilde{\mathbf{Q}}} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{\bar{\rho}_{\bar{p}}}{\bar{\rho}} (P_{k^{sgs}} - D_{k^{sgs}}) & 0 & 0 & 0 & \frac{\bar{\rho}_{\tilde{T}}}{\bar{\rho}} (P_{k^{sgs}} - D_{k^{sgs}}) & \frac{P_{k^{sgs}}}{2k^{sgs}} + \frac{1}{3} \bar{\rho} \nabla \cdot \tilde{\mathbf{u}} & 0 \\ \bar{\omega}_{\bar{p},i} & 0 & 0 & 0 & \bar{\omega}_{\tilde{T},i} & \bar{\omega}_{k^{sgs},i} & \bar{\omega}_{\tilde{Y}_j,i} \end{bmatrix} \quad (2.110)$$

where it should be noted the reaction rates do depend on the subgrid kinetic energy, k^{sgs} , because of the DTF model.

The decoupling of pressure from the equation of state causes the stiffness leading to the divergence of $\bar{\rho}_p$. The $\bar{\rho}_p$ terms in the transformation Jacobian are replaced with a new term designed to alter the eigenvalues of the system and reduce the condition number. This new matrix, the preconditioning matrix, is Γ and the governing equation becomes:

$$\Gamma \frac{\partial}{\partial \tau} \iiint \tilde{\mathbf{Q}} dV + \frac{\partial}{\partial t} \iiint \tilde{\mathbf{W}} dV + \iint [\tilde{\mathbf{F}} - \tilde{\mathbf{G}}] \cdot d\mathbf{A} = \tilde{\mathbf{S}} \quad (2.111)$$

The preconditioning matrix is given by [3]:

$$\Gamma = \begin{bmatrix} \Theta & 0 & 0 & 0 & \bar{\rho}_{\tilde{T}}\delta & 0 & \bar{\rho}_{\tilde{Y}_j} \\ \Theta \tilde{u} & \bar{\rho} & 0 & 0 & \bar{\rho}_{\tilde{T}}\tilde{u}\delta & 0 & \bar{\rho}_{\tilde{Y}_j}\tilde{u} \\ \Theta \tilde{v} & 0 & \bar{\rho} & 0 & \bar{\rho}_{\tilde{T}}\tilde{v}\delta & 0 & \bar{\rho}_{\tilde{Y}_j}\tilde{v} \\ \Theta \tilde{w} & 0 & 0 & \bar{\rho} & \bar{\rho}_{\tilde{T}}\tilde{w}\delta & 0 & \bar{\rho}_{\tilde{Y}_j}\tilde{w} \\ \Theta \tilde{H} - 1 & \bar{\rho}\tilde{u} & \bar{\rho}\tilde{v} & \bar{\rho}\tilde{w} & (\bar{\rho}_{\tilde{T}}\tilde{H} + \bar{\rho}\tilde{H}_{\tilde{T}})\delta & \bar{\rho} & \bar{\rho}\tilde{H}_{\tilde{Y}_j} + \bar{\rho}_{\tilde{Y}_j}\tilde{H} \\ \Theta k^{sgs} & 0 & 0 & 0 & \bar{\rho}_{\tilde{T}}k^{sgs}\delta & \bar{\rho} & \bar{\rho}_{\tilde{Y}_j}k^{sgs} \\ \Theta \tilde{Y}_i & 0 & 0 & 0 & \bar{\rho}_{\tilde{T}}\tilde{Y}_i\delta & 0 & \tilde{Y}_i\bar{\rho}_{\tilde{Y}_j} + \bar{\rho}\delta_{ij} \end{bmatrix} \quad (2.112)$$

where \tilde{H} is the total enthalpy and Θ is the parameter that rescales the eigenvalues:

$$\Theta = \frac{1}{\tilde{U}_r^2} - \frac{\bar{\rho}_{\tilde{T}}}{\bar{\rho}\tilde{h}_{\tilde{T}}} \quad (2.113)$$

When the preconditioning is turned off, $\Theta = \bar{\rho}_p$. The parameter δ is either 0 or 1; when $\delta = 0$, Γ is the original Choi-Merkle [20] preconditioning matrix and when $\delta = 1$, Γ is the Weiss-Smith [5] preconditioning matrix. The selection of the reference velocity, \tilde{U}_r , determines the amount and type of preconditioning used and also affects the selection of the time step. In general, the reference velocity may be defined using global or local criteria. Global criteria use information that is the same everywhere in the domain whereas local criteria use information from the cell and its neighbors. Local reference velocities are

desirable because they scale the required amount of preconditioning for each region of the flow. The local definitions of the reference velocity are found by performing a pressure perturbation analysis and finding values that balance the pressure gradient term in the momentum equation under various limits [3]:

$$U_{r,\text{local}} = \min \left(\max \left(|\mathbf{u}|, \frac{l}{\pi \Delta t}, \frac{\nu S_F}{V} \right), c \right) \quad (2.114)$$

where $|\mathbf{u}|$ is the steady preconditioning limit, $\frac{l}{\pi \Delta t}$ is the unsteady preconditioning limit with some appropriate reference length l and $\frac{\nu S_F}{V}$ is the viscous preconditioning limit. Unfortunately, a purely local definition like this fails in stagnation regions where $\mathbf{u} = 0$. One approach to counteract this is to introduce a global preconditioning definition [3, 5, 104]:

$$U_{r,\text{global}} = \min (\kappa V_\infty, c) \quad (2.115)$$

where V_∞ is some reference, or freestream, velocity and κ is a small coefficient. However, this global definition is not always easy to define for flows containing multiple flow regimes [3].

An alternative approach to dealing with stagnation regions was proposed by Weiss et al. [61] and analyzed by Darmofal and Siu [105]. Weiss et al. proposed adding a local velocity scale representative of the pressure fluctuations:

$$U_{r,\text{p,local}} = \sqrt{\frac{|\Delta p|}{\rho}} \quad (2.116)$$

where Δp is any local pressure difference that is easy to compute; most commonly, this is taken to be the maximum pressure difference across the cell faces [3, 61, 104]. It is also possible to introduce a global pressure fluctuation definition:

$$U_{r,\text{p,global}} = \sqrt{\frac{|p_g|}{\rho}} \quad (2.117)$$

Table 2.1: Definitions for the common limiting values of the reference velocity used in Equation (2.113) to determine the preconditioning parameter.

	Definition	Local/Global	Purpose
(A)	$ \mathbf{u} $	Local	Steady
(B)	$\frac{l}{\pi \Delta t}$	Local	Unsteady
(C)	$\frac{\nu S_F}{V}$	Local	Viscous
(D)	$\sqrt{\frac{ \Delta p }{\rho}}$	Local	Pressure perturbations
(E)	κV_∞	Global	Stagnation point
(F)	$\sqrt{\frac{ p_g }{\rho}}$	Global	Pressure perturbation
(G)	c	Global	Speed of sound

where $p_g = p - p_\infty$ is the gauge pressure. Naturally, all of the local and global terms may be combined to ensure that the reference velocity is defined everywhere. However, these definitions alone exhibit a lack of robustness when the flow experiences pressure perturbations and additional controls are required [3, 104, 106–108]. Table 2.1 summarizes the commonly used local and global definitions for the reference velocity.

Boundary Conditions

The equations for the boundary conditions described in Section 2.6 are applicable at low Mach numbers as well, without any modifications. In fact, walls and the traditional supersonic inflow and supersonic outflow conditions are all applied with no modifications. The Robin type simple inflows and outflows are also applied without modifications, although their behavior at low Mach numbers is substantially different in the preconditioned system than the original system [103]. The characteristic boundary conditions presented in Section 2.6.1 use the eigenvalues and eigenvectors of the system. When preconditioning

is applied, these have both been modified and therefore changes should be made to the boundary conditions as well [109–111]. Numerical studies indicate that failure to modify the boundary conditions degrades convergence rates, but the solution quality does not suffer [110, 111]. The following sections outline the differences required to account for the preconditioned system as well as the analysis demonstrating the efficiency of the simple boundaries.

Characteristic Inflows and Outflows

The governing equation used to begin the derivation of the characteristic boundary conditions, Equation (2.74), is modified to account for the dual-time stepping formulation:

$$\mathbf{\Gamma}'_{io} \frac{\partial \mathbf{U}'_{io}}{\partial \tau} + \frac{\partial \mathbf{U}'_{io}}{\partial t} + \mathbf{A}' \frac{\partial \mathbf{U}'_{io}}{\partial x_n} + \mathbf{D}'' = 0 \quad (2.118)$$

where $\mathbf{\Gamma}'_{io}$ is the preconditioning matrix formulated using the rotated primitives \mathbf{U}'_{io} . To find the system in the pseudo-time space, the equation is rearranged:

$$\frac{\partial \mathbf{U}'_{io}}{\partial \tau} + [\mathbf{\Gamma}'_{io}]^{-1} \frac{\partial \mathbf{U}'_{io}}{\partial t} + [\mathbf{\Gamma}'_{io}]^{-1} \mathbf{A}' \frac{\partial \mathbf{U}'_{io}}{\partial x_n} + [\mathbf{\Gamma}'_{io}]^{-1} \mathbf{D}'' = 0 \quad (2.119)$$

Just as in the non-preconditioned case, the system is diagonalized [109]:

$$\mathbf{S}_{PC}^{-1} \left([\mathbf{\Gamma}'_{io}]^{-1} \mathbf{A}' \right) \mathbf{S}_{PC} = \mathbf{\Lambda}_{PC} \quad (2.120)$$

where \mathbf{S}_{PC} is the right eigenvector matrix of the preconditioned system and $\mathbf{\Lambda}_{PC}$ is the diagonal matrix of eigenvalues for the preconditioned system. In this case, the eigenvalues

are [3, 109]:

$$\begin{aligned}
\lambda_{PC,1} &= \frac{1}{2} \left[(1 + \epsilon) \tilde{u}_1 - \sqrt{(1 - \epsilon)^2 \tilde{u}_1^2 + 4\epsilon c^2} \right] \\
\lambda_{PC,2} &= \tilde{u}_1 \\
\lambda_{PC,3} &= \tilde{u}_1 \\
\lambda_{PC,4} &= \tilde{u}_1 \\
\lambda_{PC,5} &= \frac{1}{2} \left[(1 + \epsilon) \tilde{u}_1 + \sqrt{(1 - \epsilon)^2 \tilde{u}_1^2 + 4\epsilon c^2} \right] \\
\lambda_{PC,5} &= \tilde{u}_1 \\
\lambda_{PC,6+k} &= \tilde{u}_1
\end{aligned} \tag{2.121}$$

where $\epsilon = U_r/c$ is the reference Mach number. All but two eigenvalues remain the same as the non-preconditioned case. The remaining two eigenvalues contain the preconditioning parameter and are written as $\lambda_{PC,1} = \tilde{u}' - c'$ and $\lambda_{PC,5} = \tilde{u}' + c'$. In the limit of no preconditioning, $\epsilon \rightarrow 1$ and these two eigenvalues revert to their normal values of $\tilde{u}_1 \pm c$.

The diagonalization is substituted into the system:

$$\mathbf{S}_{PC}^{-1} \frac{\partial \mathbf{U}'_{io}}{\partial \tau} + \mathbf{S}_{PC}^{-1} [\mathbf{\Gamma}'_{io}]^{-1} \frac{\partial \mathbf{U}'_{io}}{\partial t} + \mathbf{L}_{PC} + \mathbf{S}_{PC}^{-1} [\mathbf{\Gamma}'_{io}]^{-1} \mathbf{D}'' = 0 \tag{2.122}$$

with the wave amplitude vector $\mathbf{L}_{PC} = \mathbf{\Lambda}_{PC} \mathbf{S}_{PC}^{-1} \frac{\partial \mathbf{U}'_{io}}{\partial x_n}$ [110, 111]:

$$\mathbf{L}_{PC} = \begin{bmatrix} L_{PC,1} \\ L_{PC,2} \\ L_{PC,3} \\ L_{PC,4} \\ L_{PC,5} \\ L_{PC,6} \\ L_{PC,6+k} \end{bmatrix} = \begin{bmatrix} \frac{\lambda_{PC,1}}{2} \left(\frac{\partial \bar{p}}{\partial x_n} + \bar{\rho} (\tilde{u}_1 - c' - \tilde{u}') \frac{\partial \tilde{u}_1}{\partial x_n} \right) \\ \lambda_{PC,2} \left(\frac{\partial \bar{p}}{\partial x_n} - \frac{1}{U_r^2} \frac{\partial \bar{p}}{\partial x_n} \right) \\ \lambda_{PC,3} \frac{\partial \tilde{u}_2}{\partial x_n} \\ \lambda_{PC,4} \frac{\partial \tilde{u}_3}{\partial x_n} \\ \frac{\lambda_5}{2} \left(\frac{\partial \bar{p}}{\partial x_n} + \bar{\rho} (\tilde{u}_1 + c' - \tilde{u}') \frac{\partial \tilde{u}_1}{\partial x_n} \right) \\ \lambda_{PC,6} \frac{\partial k^{sgs}}{\partial x_n} \\ \lambda_{PC,6+k} \frac{\partial \tilde{Y}_k}{\partial x_n} \end{bmatrix} \tag{2.123}$$

and the vector $\mathbf{d} = \mathbf{S}_{PC} \mathbf{L}_{PC}$:

$$\mathbf{d} = \begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \\ d_{6+k} \end{bmatrix} = \begin{bmatrix} \frac{L_{PC,5} - L_{PC,1}}{\bar{\rho}c'} \\ L_{PC,3} \\ L_{PC,4} \\ L_{PC,2} + \frac{1}{\bar{U}_r^2} \left[\frac{(\tilde{u}_1 + c' - \tilde{u}')L_{PC,1} - (\tilde{u}_1 - c' - \tilde{u}')L_{PC,5}}{c'} \right] \\ \frac{(\tilde{u}_1 + c' - \tilde{u}')L_{PC,1} - (\tilde{u}_1 - c' - \tilde{u}')L_{PC,5}}{c'} \\ L_{PC,6} \\ L_{PC,6+k} \end{bmatrix} \quad (2.124)$$

such that:

$$\frac{\partial \mathbf{U}'_{io}}{\partial \tau} + [\mathbf{\Gamma}'_{io}]^{-1} \frac{\partial \mathbf{U}'_{io}}{\partial t} + \mathbf{d}_{PC} + [\mathbf{\Gamma}'_{io}]^{-1} \mathbf{D}'' = 0 \quad (2.125)$$

As with the non-preconditioned system, the transverse and viscous terms are neglected to form the LODI system. To ensure conservation in physical time, the system is multiplied by $\mathbf{\Gamma}'_{io}$

$$\mathbf{\Gamma}'_{io} \frac{\partial \mathbf{U}'_{io}}{\partial \tau} + \frac{\partial \mathbf{U}'_{io}}{\partial t} + \mathbf{\Gamma}'_{io} \mathbf{d}_{PC} = 0 \quad (2.126)$$

and it is assumed that the pseudo-time derivative will be driven to zero [109]:

$$\frac{\partial \mathbf{U}'_{io}}{\partial t} + \mathbf{\Gamma}'_{io} \mathbf{d}_{PC} = 0 \quad (2.127)$$

Simple Inflows and Outflows

Section 2.6.2 introduced the Robin inflow and outflow boundaries, where the velocities and temperature are imposed and pressure extrapolated on the inflow with the opposite occurring on the outflow. In that section, a summary of an eigenmode analysis performed by Darmofal, Moinier and Giles [103] indicated these boundaries are perfectly reflecting for the non-preconditioned system. The authors also performed the same analysis using the

preconditioned governing equations, for the same preconditioner described in Section 2.7. Using their notation for the system, the eigenvalues are $\lambda_1 = M$, $\lambda_2 = 1/2(M(1 + \epsilon) + \tau)$ and $\lambda_3 = 1/2(M(1 + \epsilon) - \tau)$, where if preconditioning is not active, $\lambda_2 = u + c$ and $\lambda_3 = u - c$, and with M the Mach number, ϵ the preconditioning parameter and τ is related to c' . From this, the frequencies are found [103]:

$$\begin{aligned}\mathcal{R}(\omega) &= \frac{\lambda_2 \lambda_3}{\lambda_3 - \lambda_2} (2n + 1) \pi; & n \in \mathbb{Z} \\ \mathcal{I}(\omega) &= -\frac{\lambda_2 \lambda_3}{\lambda_3 - \lambda_2} \log \left[\frac{\tau + M(1 - \epsilon)}{\tau - M(1 - \epsilon)} \right]\end{aligned}\tag{2.128}$$

Unlike the non-preconditioned case, the frequencies with the preconditioned system have an imaginary component. Additionally, the frequencies remain bounded as the Mach number approaches zero. Therefore, disturbances will decay exponentially and the boundaries are not reflective. This is in contrast to the behavior of the boundaries in the non-preconditioned system, where they are perfectly reflective. This difference makes these boundaries suitable for low Mach number flows with a preconditioned scheme.

CHAPTER 3

TRADITIONAL MACCORMACK SCHEME

The MacCormack scheme was first put forth in 1969 [38] and has endured the years since as a well-studied and often-implemented scheme. The scheme uses a unique combination of first-order stencils in space to create a second-order accurate scheme that is inherently stable. This approach is attractive as it maintains a minimal stencil size, which improves parallel performance, and requires few storage locations in memory, reducing the computational footprint. It also is free of tunable parameters, reducing the expertise needed to apply it. Because of this, many variations of the scheme have been proposed over the years to improve and extend it. Implicit versions of the scheme exist [112] as well as schemes suitable for shock capturing that employ corrections to ensure the scheme is total variation diminishing (TVD) [113–115].

Perhaps the most commonly used form of the scheme is the 1971 version proposed by MacCormack [39]. In this version, the direction of the first-order stencils is alternated in each step to eliminate hypothetical biasing errors that may arise. This version also contains a switch to the interface velocity calculation to address a weak, non-linear instability that may cause simulations to diverge. This non-linear instability correction is not well-known and not required in every case [116]. In practice, experience with this scheme has indicated that both corrections are important for stability and accuracy in chemically reacting, turbulent flows. For this reason, the 1971 form of the scheme is the version analyzed in this chapter for its performance at low Mach numbers. The deficiencies found are addressed in the next chapter. The analysis in this chapter is an expanded version of the analysis in Gallagher et al. [117].

3.1 1971 MacCormack Scheme

Integration with the MacCormack scheme uses one-dimensional sweeps for each computational coordinate in a predictor step and corrector step to yield a second-order accurate scheme in both space and time. The viscous flux vector in the Navier-Stokes equations is treated using second-order, central derivatives and is not unique to the MacCormack scheme. These terms are not included in the following analysis to render the distinctions in the MacCormack scheme clearer. Additionally, the bars and tildes are omitted from the variables and the integration works for both DNS and LES equations. The predictor step for the i computational direction is:

$$\Delta \mathbf{W}_{i,j,k}^{n+1/2} = -\frac{\Delta t}{V_{i,j,k}} \left[\mathbf{F}(\mathbf{W}^n)_{i+1/2,j,k} \cdot \mathbf{S}_{i+1/2,j,k} - \mathbf{F}(\mathbf{W}^n)_{i-1/2,j,k} \cdot \mathbf{S}_{i-1/2,j,k} \right] \quad (3.1)$$

where $V_{i,j,k}$ is the volume of the cell, \mathbf{S} is the cell face area vector and $\Delta \mathbf{W}$ is the incremental change in the conservative variable vector. Similar sweeps are then performed in the other computational directions prior to the corrector step:

$$\begin{aligned} \Delta \mathbf{W}_{i,j,k}^{n+1/2} &= \Delta \mathbf{W}_{i,j,k}^{n+1/2} - \frac{\Delta t}{V_{i,j,k}} \left[\mathbf{F}(\mathbf{W}^n)_{i,j+1/2,k} \cdot \mathbf{S}_{i,j+1/2,k} - \mathbf{F}(\mathbf{W}^n)_{i,j-1/2,k} \cdot \mathbf{S}_{i,j-1/2,k} \right] \\ \Delta \mathbf{W}_{i,j,k}^{n+1/2} &= \Delta \mathbf{W}_{i,j,k}^{n+1/2} - \frac{\Delta t}{V_{i,j,k}} \left[\mathbf{F}(\mathbf{W}^n)_{i,j,k+1/2} \cdot \mathbf{S}_{i,j,k+1/2} - \mathbf{F}(\mathbf{W}^n)_{i,j,k-1/2} \cdot \mathbf{S}_{i,j,k-1/2} \right] \end{aligned} \quad (3.2)$$

where the final time update is then:

$$\mathbf{W}_{i,j,k}^{n+1/2} = \mathbf{W}_{i,j,k}^n + \Delta \mathbf{W}_{i,j,k}^{n+1/2} \quad (3.3)$$

The same coordinate-by-coordinate sweeps are performed for the corrector step using the updated values:

$$\begin{aligned}
\Delta \mathbf{W}_{i,j,k}^{n+1} &= -\frac{\Delta t}{V_{i,j,k}} \left[\mathbf{F}(\mathbf{W}^{n+1/2})_{i+1/2,j,k} \cdot \mathbf{S}_{i+1/2,j,k} - \mathbf{F}(\mathbf{W}^{n+1/2})_{i-1/2,j,k} \cdot \mathbf{S}_{i-1/2,j,k} \right] \\
\Delta \mathbf{W}_{i,j,k}^{n+1} &= \Delta \mathbf{W}_{i,j,k}^{n+1} - \frac{\Delta t}{V_{i,j,k}} \left[\mathbf{F}(\mathbf{W}^{n+1/2})_{i,j+1/2,k} \cdot \mathbf{S}_{i,j+1/2,k} - \mathbf{F}(\mathbf{W}^{n+1/2})_{i,j-1/2,k} \cdot \mathbf{S}_{i,j-1/2,k} \right] \\
\Delta \mathbf{W}_{i,j,k}^{n+1} &= \Delta \mathbf{W}_{i,j,k}^{n+1} - \frac{\Delta t}{V_{i,j,k}} \left[\mathbf{F}(\mathbf{W}^{n+1/2})_{i,j,k+1/2} \cdot \mathbf{S}_{i,j,k+1/2} - \mathbf{F}(\mathbf{W}^{n+1/2})_{i,j,k-1/2} \cdot \mathbf{S}_{i,j,k-1/2} \right] \\
\mathbf{W}_{i,j,k}^{n+1} &= \frac{1}{2} \left(\mathbf{W}_{i,j,k}^{n+1/2} + \mathbf{W}_{i,j,k}^n + \Delta \mathbf{W}_{i,j,k}^{n+1} \right)
\end{aligned} \tag{3.4}$$

With no loss of generality, the scheme will be expressed using only the i directional sweep for clarity.

The flux vector may be split into convective and pressure fluxes for notational purposes useful later:

$$\begin{aligned}
\mathbf{W}_{i,j,k}^{n+1/2} &= \mathbf{W}_{i,j,k}^n - \frac{\Delta t}{V_{i,j,k}} \left(S_{i+1/2,j,k} U_{i+1/2,j,k}^n \mathbf{W}_{i+1/2,j,k}^n + S_{i+1/2,j,k} p_{i+1/2,j,k}^n \mathbf{P}_{i+1/2,j,k}^n \right. \\
&\quad \left. - S_{i-1/2,j,k} U_{i-1/2,j,k}^n \mathbf{W}_{i-1/2,j,k}^n + S_{i-1/2,j,k} p_{i-1/2,j,k}^n \mathbf{P}_{i-1/2,j,k}^n \right)
\end{aligned} \tag{3.5}$$

where U is the contravariant velocity at the cell face with normal vector \hat{n} , $U = \mathbf{u} \cdot \hat{n}$, and $\mathbf{P} = \{0, 1, 0, 0, U\}^\top$. The same splitting of fluxes is done for the corrector step. The scheme as written is a generic second-order Runge-Kutta time integration method and it is the spatial discretization of the fluxes across each cell face that distinguishes the MacCormack scheme from other methods.

In the 1969 MacCormack scheme, the interface values are computed using upwind or downwind values [38], for example taking $(\cdot)_{i+1/2} = (\cdot)_{i+1}$ and $(\cdot)_{i-1/2} = (\cdot)_i$ in the predictor step:

$$\begin{aligned}
\mathbf{W}_{i,j,k}^{n+1/2} &= \mathbf{W}_{i,j,k}^n - \frac{\Delta t}{V_{i,j,k}} \left(S_{i+1/2,j,k} U_{i+1,j,k}^n \mathbf{W}_{i+1,j,k}^n + S_{i+1/2,j,k} p_{i+1,j,k}^n \mathbf{P}_{i+1,j,k}^n \right. \\
&\quad \left. - S_{i-1/2,j,k} U_{i,j,k}^n \mathbf{W}_{i,j,k}^n + S_{i-1/2,j,k} p_{i,j,k}^n \mathbf{P}_{i,j,k}^n \right)
\end{aligned} \tag{3.6}$$

Table 3.1: Alternating difference direction pattern for MacCormack scheme [117].

Iteration	Predictor			Corrector		
	$i + 1/2$	$j + 1/2$	$k + 1/2$	$i + 1/2$	$j + 1/2$	$k + 1/2$
1	$i + 1$	$j + 1$	$k + 1$	i	j	k
2	$i + 1$	$j + 1$	k	i	j	$k + 1$
3	$i + 1$	j	$k + 1$	i	$j + 1$	k
4	i	$j + 1$	$k + 1$	$i + 1$	j	k
5	$i + 1$	j	k	i	$j + 1$	$k + 1$
6	i	$j + 1$	k	$i + 1$	j	$k + 1$
7	i	j	$k + 1$	$i + 1$	$j + 1$	k
8	i	j	k	$i + 1$	$j + 1$	$k + 1$

where $S_{i\pm 1/2,j,k}$ are the areas of the cell faces. For the corrector step, the $i + 1/2$ interface uses the values at i and the $i - 1/2$ interface uses the values at $i - 1$. This approach always uses the same difference direction for each of the predictor and corrector steps; for example, $u_{i+1/2,j,k}^n = u_{i+1,j,k}^n$ during the predictor step and $u_{i+1/2,j,k}^{n+1/2} = u_{i,j,k}^{n+1/2}$ for the corrector step. It was hypothesized that this fixed ordering introduced spatial biases in the solution and a more accurate solution is possible by alternating the difference directions between iterations [39]. An example of the alternating order is given in Table 3.1 and a “complete” cycle occurs over 2^D iterations, where D is the number of dimensions. MacCormack [39] noted that there was no analytical way to verify or refute the hypothesis about bias in the solution. MacCormack does note, however, that introducing the alternating directions may allow for a higher stable CFL number in some situations.

The 1971 analysis of the MacCormack scheme [39] notes that the original scheme as given above may develop a weak, non-linear instability when $U_{i,j,k} < 0$ and $U_{i+1,j,k} > 0$. This flow configuration results in a loss of information regarding the sign of the momentum and velocity in the momentum term $(\rho u)_{i+1/2,j,k} u_{i+1/2,j,k}$. Because of this, both the positive

and negative momentum on each side of the face may grow unbounded despite the sum remaining finite. This type of instability also occurs for the leapfrog scheme [116] and may cause simulations to crash unexpectedly, particularly in the neighborhood of stagnation or separation points. MacCormack [39] proposed altering the calculation of the interface velocity when this is detected:

$$U_{i+1/2,j,k} = \begin{cases} \frac{U_{i+1,j,k} + U_{i,j,k}}{2} & \text{if } U_{i,j,k} < 0 \text{ and } U_{i+1,j,k} > 0 \\ U_{i+1,j,k} & \text{else} \end{cases} \quad (3.7)$$

This is equivalent to central interpolation for the face velocity. The resulting discretization retains the sign of the momentum term and all terms remain bounded. Discussions of this non-linearity correction are often omitted and in some cases the correction may not be required for stability as in the work by Mendez-Nunez and Carroll [116]. As noted earlier, experience with complex, turbulent flows indicates the correction is required for complex flows. It is possible that using artificial viscosity eliminates the need for the correction in many instances by dissipating the resulting discontinuities.

Unfortunately, this non-linearity correction may lead to pressure oscillations in the solution. Figure 3.1a demonstrates the oscillations in the pressure field that arise from the original non-linearity correction in a simple temporal mixing layer used as a standard test case in LESLIE (for demonstration only, the details of the case are omitted here and can be found in [63]). The oscillations are relatively minor in this case as the Reynolds number is low enough for viscous dissipation to keep them under control. Figure 3.1b is the pressure field when the non-linearity switch is disabled. The pressure field is smooth and the simulation is stable, although the stability is not guaranteed in all cases when the switch is off. Lastly, Figure 3.1c is the pressure field using a new form of the non-linearity switch proposed in this work. At low Mach numbers, these oscillations are generally larger than the expected solution variation and the original formulation of the switch will lead to degraded solution quality. The formulation of the new switch is given with the full formulation for

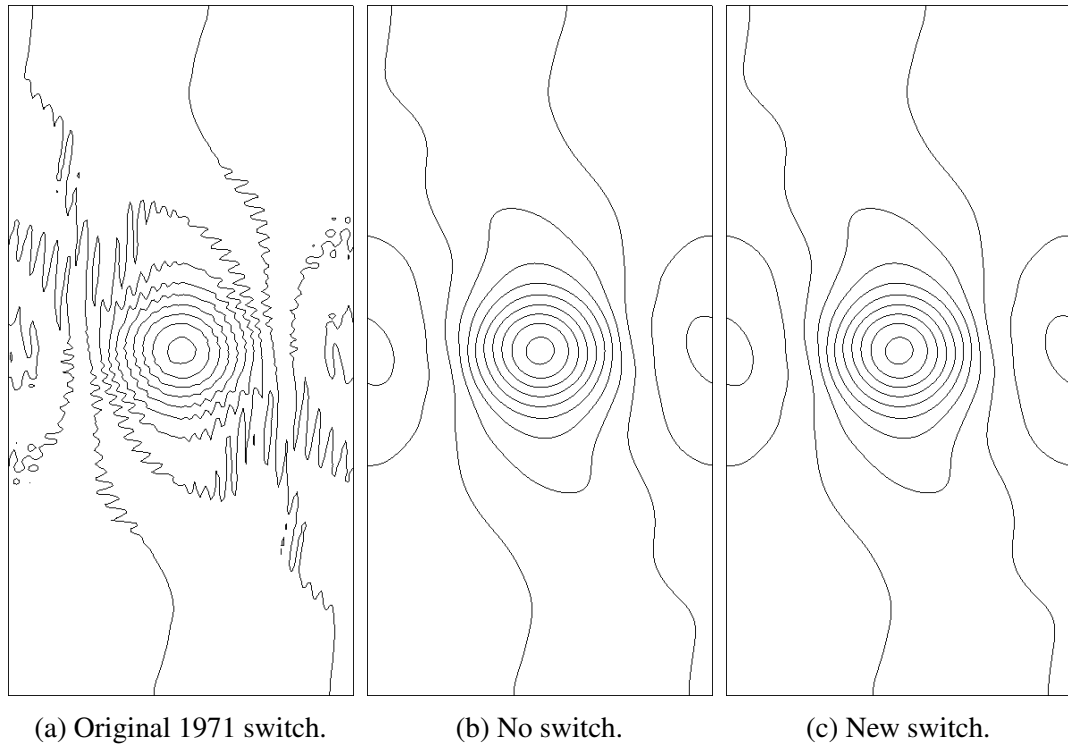


Figure 3.1: Pressure oscillations in a smooth flow using various forms of the non-linearity switch. Oscillations with the original switch are large enough to degrade solution quality at low Mach numbers. The details of the new switch are in Chapter 4 and Appendix A.

low Mach numbers in Chapter 4, while more details of the problem in the original scheme are in Appendix A.

The addition of artificial viscosity, particularly a second-order form, permits the use of the scheme in non-smooth flows. With pressure-based sensors for the artificial viscosity, shocks are detected and the artificial viscosity activates to reduce the oscillations that ordinarily arise. For low-speed flows, density or temperature based sensors are useful to detect flame fronts. Without the artificial viscosity in the neighborhood of the flame, the large gradients in density, temperature, and species mass fractions can cause numerical instabilities. In the next section, a commonly used form of artificial dissipation is introduced for use with the MacCormack scheme in reacting flows.

3.2 Artificial Dissipation

Traditionally, extra dissipation is required to stabilize central schemes to damp out dispersive errors. The first-order stencils used in the MacCormack scheme add dissipation that is generally sufficient for smooth flows and artificial dissipation is not needed unless shocks or flames are present. The Jameson-Schmidt-Turkel (JST) artificial dissipation scheme adds a second- and fourth-order dissipation term to stabilize central schemes in the presence of non-smooth flow features and to dissipate the dispersive errors present in smooth flows, respectively [24]. The artificial dissipation flux for a finite-volume scheme in generalized coordinates is given by:

$$AD = d_{i+\frac{1}{2},j,k} - d_{i-\frac{1}{2},j,k} + d_{i,j+\frac{1}{2},k} - d_{i,j-\frac{1}{2},k} + d_{i,j,k+\frac{1}{2}} - d_{i,j,k-\frac{1}{2}} \quad (3.8a)$$

$$d_{i+\frac{1}{2}} = \lambda_{i+\frac{1}{2}} \left(\epsilon_{i+\frac{1}{2}}^{(2)} [\mathbf{W}_{i+1} - \mathbf{W}_i] - \epsilon_{i+\frac{1}{2}}^{(4)} [\mathbf{W}_{i+2} - 3\mathbf{W}_{i+1} + 3\mathbf{W}_i - \mathbf{W}_{i-1}] \right) \quad (3.8b)$$

where the indices j, k are omitted for clarity in Equation 3.8b. The values of λ , $\epsilon^{(2)}$ and $\epsilon^{(4)}$ at the cell faces are computed as arithmetic averages of the neighboring cell-centered values. The eigenvalue takes the usual value of the fastest-traveling wave, $\lambda_{i,j,k} = (|U|+c) A$ where $|U|$ is the contravariant velocity normal to the cell face and A is the face area.

The second-order dissipation coefficient is designed to drop the order of the scheme to first-order near discontinuities in a variable. To detect both shocks and flames, a pressure sensor and a density sensor are used:

$$\epsilon_{i,j,k}^{(2)} = \kappa^{(2)} \max_q (S_q^p, S_q^\rho); \quad \forall q \in \{\text{cell and its neighbors}\} \quad (3.9)$$

where $\kappa^{(2)}$ is a user-defined dissipation coefficient. The sensor for any variable ϕ is computed in each flux direction; for example, in the i -direction:

$$S_{i,j,k}^{\phi} = \frac{|\phi_{i+1,j,k} - 2\phi_{i,j,k} + \phi_{i-1,j,k}|}{|\phi_{i+1,j,k}| + 2|\phi_{i,j,k}| + |\phi_{i-1,j,k}|} \quad (3.10)$$

The fourth-order dissipation is designed to be effective only in smooth regions of the flow and is defined:

$$\epsilon_{i,j,k}^{(4)} = \max\left(0, \kappa^{(4)} - \epsilon_{i,j,k}^{(2)}\right) \quad (3.11)$$

where $\kappa^{(4)}$ is a user-defined dissipation coefficient.

It is worth noting that the artificial dissipation scheme used here is not guaranteed to eliminate the oscillations that may appear around discontinuities and is therefore not total variation diminishing (TVD). This implies the dissipation scheme is suitable only for weak discontinuities that may appear in a flow, such as flames or weak shocks in transonic conditions. In those cases, new minima or maxima may still appear and the user must determine if this compromises the simulation accuracy. Strong shocks may still lead to code divergence due to the size of pressure or temperature oscillations, despite the artificial dissipation.

3.3 Low Mach Number Analysis

One of the most important aspects of the preconditioning approach is the modification of the artificial dissipation terms [2, 3]. At low Mach numbers, the original JST scheme has sub-optimal dissipative properties. The momentum equation is overly dissipated while the pressure and temperature fields are insufficiently dissipated. Modifications are required for the dissipation scheme to behave correctly in the low Mach number limit [3]. The following analysis shows the dissipation built into the MacCormack scheme suffers from the same problems in the low Mach number limit as the artificial dissipation schemes do. If preconditioning is used, the dissipative properties of the scheme are improved for the

steady preconditioning case. However, unsteady preconditioning has the same problems as the original scheme without preconditioning and the direct application of preconditioning to the scheme is not optimal.

The analysis begins with the one-dimensional, linearized Euler equations:

$$\mathbf{u}_t + \mathbf{A}\mathbf{u}_x = 0 \quad (3.12)$$

where \mathbf{u} is the vector of unknowns, \mathbf{A} is the flux Jacobian and subscripts indicate differentiation. Provided the flux Jacobian is constant, the equations may be diagonalized such that:

$$\mathbf{w}_t + \boldsymbol{\lambda}^\top \mathbf{w}_x = 0 \quad (3.13)$$

where $\mathbf{w} = \mathbf{T}^{-1}\mathbf{u}$ and $\boldsymbol{\lambda} = \boldsymbol{\Lambda}_{ii} = \mathbf{T}^{-1}\mathbf{A}\mathbf{T}$ [118]. At this stage, the equations are still in their exact form. When discretized, the equations will be approximate due to truncation errors associated with neglected terms. These errors mean the discrete equations are no longer identical to the exact equations. However, it is possible that the discrete form of the equations does match some other, exact governing equation.

Warming and Hyatt's modified equation analysis method [119] uses Taylor series approximations of the unknown vector \mathbf{w} to find the exact governing equation that matches the discrete approximation to Equation (3.13). In the case of the MacCormack scheme, the resulting modified equation is identical to that of the Lax-Wendroff method and is [118, 119]:

$$\mathbf{w}_t + \boldsymbol{\lambda}^\top \mathbf{w}_x = -\frac{\boldsymbol{\lambda}^\top \Delta x^2}{6} \left(1 - \frac{\boldsymbol{\lambda}^{2\top} \Delta t^2}{\Delta x^2}\right) \mathbf{w}_{xxx} - \frac{\boldsymbol{\lambda}^{2\top} \Delta t \Delta x^2}{8} \left(1 - \frac{\boldsymbol{\lambda}^{2\top} \Delta t^2}{\Delta x^2}\right) \mathbf{w}_{xxxx} + \text{H.O.T} \quad (3.14)$$

where $\lambda^{2\tau} = \{\lambda_i^2\}$. For the i -th component of the unknown vector w , the modified equation is:

$$w_{i,t} + \lambda_i w_{i,x} = -\frac{\lambda_i \Delta x^2}{6} (1 - \text{CFL}_i^2) w_{i,xxx} - \frac{\lambda_i \Delta x^3}{8} \text{CFL}_i (1 - \text{CFL}_i^2) w_{i,xxxx} + \text{H.O.T} \quad (3.15)$$

The first term on the right hand side of Equation (3.15) is an odd-ordered derivative of the unknown vector and therefore indicates a dispersive term. The second term on the right hand side is even-ordered and is the dissipative term, and this is the reason for the inherent stability of the MacCormack scheme. The built in dispersion and dissipation of the MacCormack scheme is dependent on the CFL number corresponding to each wave. These terms use the wave speed of each equation, similar to a matrix artificial dissipation used in purely central schemes, and suffer from the same problems at low Mach numbers. The fourth-order dissipation term is the same as the Jameson-Schmidt-Turkel (JST) artificial dissipation [24] if the JST dissipation coefficient is $\epsilon^{(4)} = \frac{\text{CFL}(1-\text{CFL}^2)}{8}$.

With a preconditioning scheme, the eigenvalues of the flux Jacobian are rescaled using a preconditioning matrix Γ . The preconditioned form of the linearized Euler equations is found by multiplying the temporal derivative with the preconditioning matrix Γ :

$$\begin{aligned} \Gamma \mathbf{u}_\tau + \mathbf{A} \mathbf{u}_x &= 0 \\ \mathbf{u}_\tau + \Gamma^{-1} \mathbf{A} \mathbf{u}_x &= 0 \end{aligned} \quad (3.16)$$

where Γ is the appropriate preconditioning matrix and τ is the pseudo-time coordinate. The transformation from physical time in Equation (3.13) to a pseudo-time in Equation (3.16) is due to the loss of time accuracy from scaling the eigenvalues – more details are provided in Chapter 4.

This equation is diagonalized by $\lambda' = \Lambda'_{ii} = \mathbf{T}^{-1} (\mathbf{\Gamma}^{-1} \mathbf{A}) \mathbf{T}$ and $\mathbf{w} = \mathbf{T}^{-1} \mathbf{u}$, yielding the modified equation for the preconditioned MacCormack scheme:

$$\mathbf{w}_\tau + \lambda'^\top \mathbf{w}_x = -\frac{\lambda'^\top \Delta x^2}{6} \left(1 - \frac{\lambda'^\top \Delta \tau^2}{\Delta x^2} \right) \mathbf{w}_{xxx} - \frac{\lambda'^\top \Delta \tau \Delta x^2}{8} \left(1 - \frac{\lambda'^\top \Delta \tau^2}{\Delta x^2} \right) \mathbf{w}_{xxxx} + \text{H.O.T} \quad (3.17)$$

This form of the equation is not conservative in the steady-state unless $\mathbf{\Gamma}$ is constant and therefore the conservative equation is [120]:

$$\mathbf{\Gamma} \mathbf{w}_\tau + \mathbf{\Gamma} \lambda'^\top \mathbf{w}_x = -\frac{\mathbf{\Gamma} \lambda'^\top \Delta x^2}{6} \left(1 - \frac{\lambda'^\top \Delta \tau^2}{\Delta x^2} \right) \mathbf{w}_{xxx} - \frac{\mathbf{\Gamma} \lambda'^\top \Delta \tau \Delta x^2}{8} \left(1 - \frac{\lambda'^\top \Delta \tau^2}{\Delta x^2} \right) \mathbf{w}_{xxxx} + \text{H.O.T} \quad (3.18)$$

In this method, both the dispersion and the dissipation terms use the preconditioned wave speeds. Depending on the form of the preconditioning matrix, this may or may not be suitable. Using a steady preconditioning matrix, the terms are well-behaved at the low Mach number limit and no further modifications are required. However, in the unsteady, low Mach limit that arises in high Strouhal number flows or acoustic problems, the resolution of the acoustic waves is important. In this case, an unsteady preconditioner should be used. This preconditioner approaches the unpreconditioned limit as will be discussed in Chapter 4. When this occurs, the inherent dissipation in the MacCormack scheme reverts to the dissipation of the scheme without preconditioning and becomes overly dissipative. The straightforward application of preconditioning to the MacCormack scheme shall be called the MC-PC scheme. Testing of this scheme demonstrates the overly dissipative nature when used with the unsteady preconditioner and is reported in Chapter 5. Modifications to the scheme are required to ensure proper behavior at the unsteady, low Mach number limit. These modifications result in the generalized MacCormack scheme for low Mach number flows introduced in the next chapter.

CHAPTER 4

GENERALIZED MACCORMACK SCHEME

The 1971 form of the MacCormack scheme is inherently stable, but analysis of the method at low Mach numbers indicates the inherent dissipation is overly dissipative in that limit when used with the unsteady preconditioner. During the analysis of the scheme in Chapter 3, a preconditioning matrix Γ was introduced but no details provided. With the dual-time form of the preconditioned governing equations detailed in Chapter 2, the issues surrounding the inherent dissipation in the MacCormack scheme may be addressed. This is done by modifying the discretization procedure for the pressure fluxes and incorporating an artificial dissipation on the pressure field alone to ensure the correct scaling of the eigenvalues in the unsteady, low Mach number limit. An additional correction to the scheme eliminates the instabilities observed when using the original non-linearity switch recommended by MacCormack [39]. For completeness, the single-time MacCormack scheme without preconditioning is shown with the new non-linearity switch, and this scheme is called the Generalized MacCormack-Single Time scheme (GMC-ST). Validation of the non-linearity switch for the GMC-ST scheme is provided in Appendix A. The dual-time form with preconditioning is called the Generalized MacCormack-Preconditioned (GMC-PC) scheme and the changes required relative to the GMC-ST scheme are provided. The description of the GMC-ST and GMC-PC schemes is an extension of that already published by Gallagher et al. [117].

Lastly, the time step selection method for the pseudo-time is addressed. The time step selection depends on the choice of preconditioners used and there are many global and local preconditioner definitions from which to choose. Preconditioned schemes typically exhibit robustness issues in complex flows. This is most frequent near stagnation points [3, 5, 104] and when the flow is nonlinear, and techniques exist to limit the time step when this occurs

[60]. The non-linear time step controls adjust the time step after the fluxes are computed to limit the amount of change in the solution. Because the time step is selected after the fluxes are computed, these are not easily generalizable to multi-step, explicit schemes such as the one used here, and when using an external, stiff source term solver the controls may lead to an expensive loop to negotiate robust time steps. A new procedure based on the Fourier analysis technique of Darmofal and Siu [105] shows how to design reference velocities for use in the preconditioning definition that limit the solution change in the same way as the non-linear time step controls. By modifying the preconditioner prior to the flux calculation, no problems arise for multi-step schemes or for external, stiff solvers. A detailed study of the existing and new preconditioner definitions indicates global preconditioners are not needed and the new preconditioner definitions improve code stability in the presence of strong pressure waves by several orders of magnitude.

4.1 Generalized MacCormack Scheme

The 1971 MacCormack scheme is suitable for single-time, explicit integration of moderate Mach numbers. However, modifications are required for accuracy in low Mach number, unsteady flows. Additionally, due to the central nature of the final inviscid flux, the scheme requires additional artificial dissipation to properly handle discontinuities that arise in flows with shocks and/or flames. For these reasons, a modified scheme is introduced in this section that can function as both a single-time, explicit integration scheme and as the pseudo-time iterator in dual-time with or without preconditioning. For clarity, the tildes and bars above the variables have been omitted.

The non-linearity correction in the original scheme may produce large-amplitude pressure oscillations that become worse as the grid is refined and are particularly damaging to the solution quality at low Mach numbers. A new non-linearity switch is proposed, drawing on inspiration from the AUSM scheme. Equation (3.5) is written with the convective and the pressure fluxes separated, similar to the derivation of the AUSM scheme [32]. When

written in this fashion, the solution vector is treated as a passive scalar convected by an interface velocity. Equation (3.6) defines the interface values using the upwind and downwind approach from MacCormack [39]. To address the situation where the information about the sign of the momentum terms is lost, MacCormack proposed averaging the interface velocity when the velocities on each side of the face are moving away from it. This leads to insufficient dissipation when the original non-linearity correction is used. In the proposed non-linearity switch, the value of the solution vector on the interface is averaged while the interface velocities are treated with the upwind-downwind approach of the MacCormack scheme:

$$\mathbf{W}_{i+1/2,j,k} = \begin{cases} \frac{\mathbf{W}_{i+1,j,k} + \mathbf{W}_{i,j,k}}{2} & \text{if } U_{i,j,k} < 0 \text{ and } U_{i+1,j,k} > 0 \\ \mathbf{W}_{i+1,j,k} \text{ or } \mathbf{W}_{i,j,k} & \text{else} \end{cases} \quad (4.1)$$

This preserves information about the sign of the momentum term while also maintaining sufficient dissipation, thereby removing the pressure oscillations. Additional details of this change are discussed in Appendix A.

4.1.1 Single-time integration

In single-time form, the governing equation, Equation (2.37), is solved using a variant of the 1971 MacCormack scheme. The predictor-corrector steps written for a single computational direction only are:

$$\begin{aligned} \mathbf{W}_{i,j,k}^{n+1/2} = \mathbf{W}_{i,j,k}^n - \frac{\Delta t}{V_{i,j,k}} & (S_{i+1/2,j,k} U_{i+1,j,k}^n \mathbf{W}_{i+1/2,j,k}^n - S_{i-1/2,j,k} U_{i,j,k}^n \mathbf{W}_{i-1/2,j,k}^n \\ & + S_{i+1/2,j,k} P_{i+1,j,k}^n \mathbf{P}_{i+1,j,k}^n - S_{i-1/2,j,k} P_{i,j,k}^n \mathbf{P}_{i,j,k}^n) \\ & + \Delta t \mathbf{S}_{i,j,k}^n \end{aligned} \quad (4.2)$$

$$\begin{aligned}
\mathbf{W}_{i,j,k}^{n+1} = & \frac{1}{2} \left[\mathbf{W}_{i,j,k}^{n+1/2} + \mathbf{W}_{i,j,k}^n \right. \\
& - \frac{\Delta t}{V_{i,j,k}} \left(S_{i+1/2,j,k} U_{i,j,k}^{n+1/2} \mathbf{W}_{i+1/2,j,k}^{n+1/2} - S_{i-1/2,j,k} U_{i-1,j,k}^{n+1/2} \mathbf{W}_{i-1/2,j,k}^{n+1/2} \right. \\
& \left. \left. + S_{i+1/2,j,k} P_{i,j,k}^{n+1/2} \mathbf{P}_{i,j,k}^{n+1/2} - S_{i-1/2,j,k} P_{i-1,j,k}^{n+1/2} \mathbf{P}_{i-1,j,k}^{n+1/2} \right) \right. \\
& \left. + \Delta t \mathbf{S}_{i,j,k}^{n+1/2} \right]
\end{aligned} \tag{4.3}$$

where in both steps, the term $\mathbf{W}_{i\pm 1/2,j,k}$ is computed using Equation (4.1). The source term vector $\mathbf{S}_{i,j,k}$ contains the production and dissipation terms for the subgrid kinetic energy as well as the chemical source terms if the flow is reacting. For non-reacting flows, the chemical source terms are zero and the entire source vector is assumed not to be stiff, allowing for direct integration of the sources with the rest of the fluxes. For reacting flows, this is no longer assumed to be the case and more details are provided in Chapter 6. The only difference between this scheme and the original scheme is the change in the nonlinearity switch from Equation (3.7) to Equation (4.1). The physical time step, Δt , is the global minimum of [4]:

$$\Delta t_{i,j,k} = V_{i,j,k} \frac{\text{CFL}}{|\mathbf{u} \cdot \mathbf{S}_I| + |\mathbf{u} \cdot \mathbf{S}_J| + |\mathbf{u} \cdot \mathbf{S}_K| + c S_F + 2\gamma \frac{\nu}{\text{Pr}} \frac{S_F}{V_{i,j,k}}} \tag{4.4}$$

where $\mathbf{S}_I, \mathbf{S}_J, \mathbf{S}_K$ are representative cell face areas projected onto the I, J , and K computational coordinates and where the total cell face area is S_F . ν is the kinematic viscosity and the CFL number is taken to be 0.5. For simplicity, this scheme is referred to as the Generalized MacCormack-Single Time (GMC-ST) scheme.

4.1.2 Low-Mach Number Formulation

The GMC-ST scheme is overly dissipative when used for low Mach numbers. For acoustic or highly unsteady low Mach number problems, even with preconditioning, MC-PC is overly dissipative, while around discontinuities the solution exhibits large oscillations due

to its central nature. Both problems can be fixed by adding particular forms of artificial dissipation to the scheme to allow better control over its characteristics. To reduce the dissipation in a central derivative scheme at the unsteady, low Mach number limit, Potsdam et al. [121] introduced a blended matrix artificial dissipation designed to use the unsteady preconditioned eigenvalues for the pressure field and the steady preconditioned eigenvalues for the other fields. By doing so, the approach ensures optimal dissipation for all variables. A similar concept may be introduced into the MacCormack scheme. The interface pressure is computed using central differencing and a JST-type fourth-order dissipation term is added to the pressure field only:

$$\Delta \mathbf{P}^{AD4} = \begin{cases} \epsilon_4 \Gamma \sigma_p \mathbf{L}_p \frac{\partial^3 \mathbf{Q}}{\partial x^3} & \text{with preconditioning} \\ \epsilon_4 \frac{\partial \mathbf{W}}{\partial \mathbf{Q}} \sigma_e \mathbf{L}_p \frac{\partial^3 \mathbf{Q}}{\partial x^3} & \text{without preconditioning} \end{cases} \quad (4.5)$$

ϵ_4 is a user-specified constant [24], σ_p is the spectral radius of $|\Gamma^{-1} \mathbf{A}|$ and σ_e is the spectral radius of $|(\frac{\partial \mathbf{W}}{\partial \mathbf{Q}})^{-1} \mathbf{A}|$. \mathbf{A} is the flux Jacobian and \mathbf{L}_p is the selector matrix designed to apply the dissipation only to the pressure term:

$$\mathbf{L}_p = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \quad (4.6)$$

This is equivalent to the approach used in Rhie-Chow's collocated scheme [4, 14] and maintains ideal dissipative characteristics in the unsteady, low Mach number limit [34]. A traditional second-order artificial dissipation term is added to reduce oscillations around discontinuities. More advanced versions of the MacCormack scheme add a term to the corrector step to enforce a total variation diminishing property and completely eliminate oscillations [113–115].

4.1.3 Dual-time integration

The pseudo-time derivative in Equation (2.111) is marched forward in pseudo-time until the temporal derivative is sufficiently small and may be assumed to be zero. To translate this into the physical-time solution, a discretization must be chosen for the physical-time derivative. An implicit discretization permits the selection of physical time step sizes that are as large as desired to capture the important transient features. In the case of non-reacting flows, the physical time derivative is discretized using a second-order backward-difference formula (BDF) [3, 5]. The physical time index is given by n while the pseudo-time index is given by m :

$$\begin{aligned}
& \left(\Gamma^m + \frac{3}{2} \frac{\Delta\tau_{i,j,k}}{\Delta t} \frac{\partial \mathbf{W}}{\partial \mathbf{Q}} - \Delta\tau_{i,j,k} \mathbf{D}^m \right) \Delta \mathbf{Q}_{i,j,k}^{m+1/2} = \\
& - \frac{\Delta\tau_{i,j,k}}{V_{i,j,k}} \left(S_{i+1/2,j,k} U_{i+1/2,j,k}^m \mathbf{W}_{i+1/2,j,k}^m - S_{i-1/2,j,k} U_{i-1/2,j,k}^m \mathbf{W}_{i-1/2,j,k}^m \right. \\
& + S_{i+1/2,j,k} \frac{p_{i+1,j,k}^m + p_{i,j,k}^m}{2} \mathbf{P}_{i+1,j,k}^m - S_{i-1/2,j,k} \frac{p_{i,j,k}^m + p_{i-1,j,k}^m}{2} \mathbf{P}_{i,j,k}^m \\
& + S_{i+1/2,j,k} (\Delta \mathbf{P}^{AD4})_{i+1/2,j,k}^m - S_{i-1/2,j,k} (\Delta \mathbf{P}^{AD4})_{i-1/2,j,k}^m \left. \right) \\
& - \frac{\Delta\tau_{i,j,k}}{2\Delta t} [3\mathbf{W}^m - 4\mathbf{W}^n + \mathbf{W}^{n-1}] + \Delta\tau_{i,j,k} \mathbf{S}_{i,j,k}^m
\end{aligned} \tag{4.7}$$

$$\begin{aligned}
& \left(\Gamma^m + \frac{3}{4} \frac{\Delta\tau_{i,j,k}}{\Delta t} \frac{\partial \mathbf{W}}{\partial \mathbf{Q}} - \frac{1}{2} \Delta\tau_{i,j,k} \mathbf{D}^{m+1/2} \right) \Delta \mathbf{Q}_{i,j,k}^{m+1} = \frac{1}{2} \left[\Gamma^m \left(\mathbf{Q}_{i,j,k}^m - \mathbf{Q}_{i,j,k}^{m+1/2} \right) \right. \\
& - \frac{\Delta\tau_{i,j,k}}{V_{i,j,k}} \left(S_{i+1/2,j,k} U_{i,j,k}^{m+1/2} \mathbf{W}_{i+1/2,j,k}^{m+1/2} - S_{i-1/2,j,k} U_{i-1,j,k}^{m+1/2} \mathbf{W}_{i-1/2,j,k}^{m+1/2} \right. \\
& + S_{i+1/2,j,k} \frac{p_{i+1,j,k}^{m+1/2} + p_{i,j,k}^{m+1/2}}{2} \mathbf{P}_{i,j,k}^{m+1/2} - S_{i-1/2,j,k} \frac{p_{i,j,k}^{m+1/2} + p_{i-1,j,k}^{m+1/2}}{2} \mathbf{P}_{i-1,j,k}^{m+1/2} \\
& + S_{i+1/2,j,k} (\Delta \mathbf{P}^{AD4})_{i+1/2,j,k}^{m+1/2} - S_{i-1/2,j,k} (\Delta \mathbf{P}^{AD4})_{i-1/2,j,k}^{m+1/2} \left. \right) \\
& \left. - \frac{\Delta\tau_{i,j,k}}{2\Delta t} \left(3\mathbf{W}^{m+1/2} - 4\mathbf{W}^n + \mathbf{W}^{n-1} \right) + \Delta\tau_{i,j,k} \mathbf{S}_{i,j,k}^{m+1/2} \right]
\end{aligned} \tag{4.8}$$

where $\Delta Q^{m+1/2} = Q^{m+1/2} - Q^m$ and $\Delta Q^{m+1} = Q^{m+1} - Q^{m+1/2}$, and Δt is the physical time step. The physical time step may be chosen arbitrarily, including infinitely large values for steady-state solutions.

The left-hand side of these expressions is a matrix whose size is $(5 + N_{LES} + N_k) \times (5 + N_{LES} + N_k)$ entries where N_{LES} is the number of LES equations (either 0 or 1 depending on whether running in DNS or LES modes) and N_k is the number of species. This matrix is dense and has no unique structure to exploit when solving the system of equations (ie. it is not banded, not positive-definite, not symmetric). It also must be updated at each sub-iteration when run using LES because of the source term Jacobian or any case when physical time accuracy is required because of the conservative to primitive Jacobian. The solution of the system of equations is performed using the Linear Algebra PACKage (LAPACK) [122] and represents a large overhead on the scheme. Each pseudo-iteration with the GMC-PC scheme in non-reacting flows takes between 1.5 and 3 times as long as each iteration with the GMC-ST scheme. However, this increased cost per iteration is more than offset in a reduced number of iterations or reduced grid size required for accurate results at low Mach numbers as discussed in Chapter 5.

The pseudo-time iterations are repeated until the change in Q is small. At this point $Q^{n+1} = Q^m$ and W^{n+1} can be computed from Q^{n+1} . Convergence is stopped when the residual meets one of the following criteria:

- (i) Simulation reaches a maximum number of iterations for that physical time step;
- (ii) Residual reduces below a maximum norm value;
- (iii) Residual has decayed a specified number of decades;
- (iv) Change in residual is less than a specified amount for several consecutive iterations.

These criteria are user-specified and are configured for each case.

Another aspect of the usual MacCormack scheme must be changed when it is used as a pseudo-time iterator. Alternating the difference directions prohibits convergence to

steady state as each time the difference direction changes between iterations, the solution is slightly different. Figure 4.1 is an example showing the convergence history of the GMC-PC for the lid-driven cavity problem at a Reynolds number of 100 and a Mach number of 0.3 (see Section 5.3 for setup details) with and without alternating the difference directions. The simulation with the alternating difference directions never converges to the steady state solution (although the mean of the unsteady solution does converge); rather, it oscillates around the steady state solution with a period of 8 iterations corresponding to the completion of the differencing cycle in two dimensions. This is in contrast with the convergence history when the difference directions are fixed and do not change with each iteration. In that situation, the convergence is rapid. In order to ensure convergence in pseudo-time when using this scheme, the alternating of the difference directions is disabled throughout the pseudo-time iterations. As a guard against potential biasing, the difference directions are alternated for each physical-time iteration consistent with the traditional application of the MacCormack scheme. This scheme is referred to as the Generalized MacCormack-Preconditioned (GMC-PC) scheme.

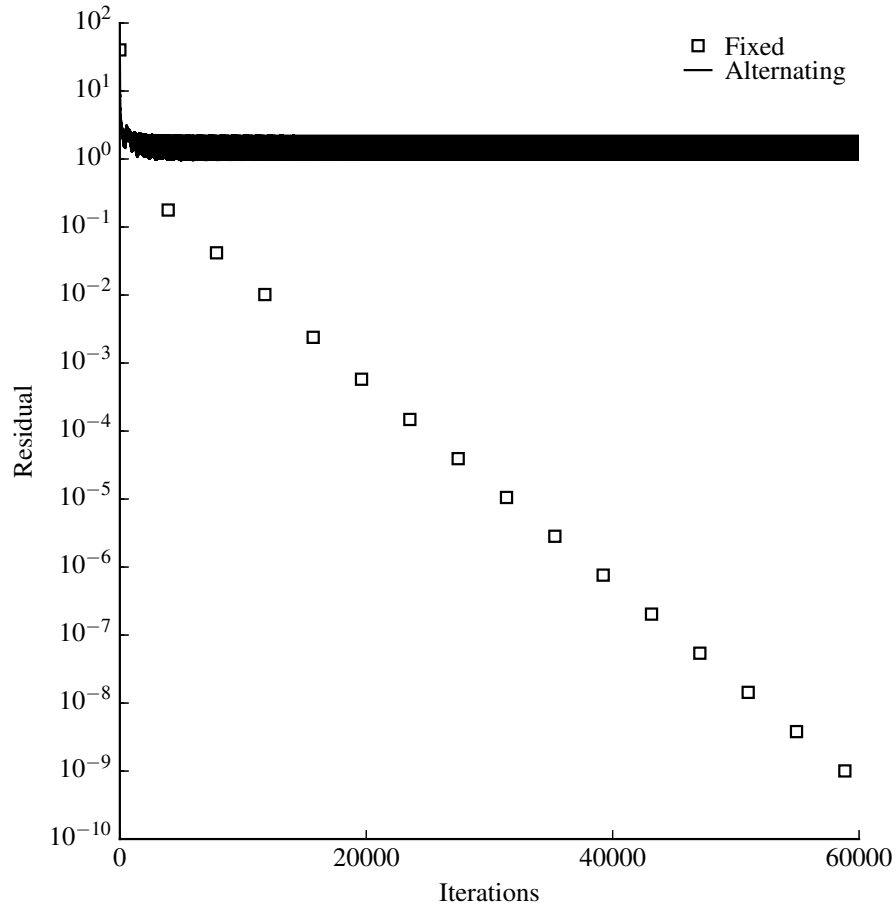


Figure 4.1: Convergence history of the GMC-PC for the lid-driven cavity at Reynolds number of 1000 and Mach number of 0.3 with/without alternating the difference directions.

4.2 Time Step Selection and Robustness

The modification of the eigenvalues destroys the time accuracy within the pseudo-time iterations. Therefore, it is possible to use various acceleration techniques to improve the rate of convergence in pseudo-time. The simplest acceleration technique is to permit the time step in each cell to vary so each cell is always evolving at its maximum stable time step. This is known as local time stepping and is the approach taken in this work. The selection of the time step value for each cell is determined by the type of preconditioner, and these have been designed to maximize robustness while still allowing rapid convergence.

4.2.1 Local Time Stepping

The following equations are used to compute the pseudo-timestep, $\Delta\tau$, in computational coordinates [5]:

$$\begin{aligned}
 u_A &= \mathbf{u} \cdot \mathbf{S}_I + \mathbf{u} \cdot \mathbf{S}_J + \mathbf{u} \cdot \mathbf{S}_K \\
 U_{r,A} &= U_r S_F \\
 \beta &= \rho_p + \frac{\rho_T}{\rho h_T} \\
 \alpha &= \frac{1}{2} (1 - \beta U_r^2) \\
 u' &= u_A (1 - \alpha) \\
 c' &= \sqrt{\alpha^2 u_A^2 + U_{r,A}^2} \\
 \frac{\Delta\tau}{V} &= \frac{\text{CFL}}{u' + c' + 2\gamma \frac{\nu}{Pr} \frac{S_F}{V}}
 \end{aligned} \tag{4.9}$$

Each cell is then permitted to evolve in pseudo-time using its local time step value. This generates instabilities when there are large spatial gradients in the local time step values. Such instabilities can be eliminated when an explicit smoothing filter is applied to the local time step to reduce the gradients. This is done by applying an explicit filter iteratively:

$$\Delta\tau_q = \min \left[\Delta\tau_q, (1 - N\omega)\Delta\tau_q + \omega \sum_{p=1}^N \Delta\tau_p \right] \tag{4.10}$$

where the subscript q corresponds to each cell center with a total number of neighbors N and p is each neighboring cell center around q . A user-specified weight, ω , is used to control the strength of the smoothing. This smoothing step reduces local maxima and increases local minima. The increase of the local minima may violate stability constraints, so the minimum value of the smoothed time step and the time step based on the stability limit in Equation (4.9) is taken. In cases where large spatial time step gradients do not occur, this smoothing is not needed and may delay convergence if applied. The user must specify whether to use the time step smoothing on a case-by-case basis, with it being unneeded in most smoothly-varying flows.

4.2.2 Robust Preconditioning

This parameter, Θ , is controlled through the selection of the reference velocity, U_r in Equation (2.113), and as $U_r \rightarrow c$, $\Theta \rightarrow \rho_p$ and the preconditioning turns off. By carefully designing the reference velocity, optimal choices minimize the stiffness. Pressure fluctuations larger than the dynamic pressure are related to the instabilities found near stagnation regions but may trigger instability away from stagnation points as well [3, 104]. Darmofal and Siu [105] analyze the pressure gradient preconditioner, definition (D) in Table 2.1 and Table 4.1, proposed by Weiss et al. [61] by considering the one-dimensional preconditioned Euler equations with symmetrizing primitive variables:

$$\frac{\partial \mathbf{Q}_s}{\partial t} + \mathbf{\Gamma}^{-1} \mathbf{A} \frac{\partial \mathbf{Q}_s}{\partial x} = 0 \quad (4.11)$$

where $\mathbf{Q}_s = \{p/\rho c, u, s\}^T$ with s being the entropy; \mathbf{A} , the flux Jacobian:

$$\mathbf{A} = \begin{bmatrix} M & 1 & 0 \\ 1 & M & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (4.12)$$

and the Weiss–Smith preconditioning matrix [5, 105]:

$$\mathbf{\Gamma}^{-1} = \begin{bmatrix} \epsilon & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (4.13)$$

where the preconditioning is controlled through the value of ϵ . The entropy equation is decoupled and not needed for the analysis. Performing a Fourier transformation of the system and introducing $\hat{\mathbf{Q}}_s = \{\hat{p}/\bar{\rho c}, \hat{u}\}^T$ where $(\hat{\cdot})$ is a perturbation and $(\bar{\cdot})$ is the mean

value gives the exact solution:

$$\hat{\mathbf{Q}}_s(t, k) = \exp(-ikt\mathbf{\Gamma}^{-1}\mathbf{A})\hat{\mathbf{Q}}_s(0, k) = \mathbf{G}(t, k)\hat{\mathbf{Q}}_s(0, k) \quad (4.14)$$

at time t and wavenumber k . The growth or decay of the initial conditions in time is the amplification matrix, $\mathbf{G} = \exp(-ikt\mathbf{\Gamma}^{-1}\mathbf{A})$. From this analysis, one can ascertain the pressure and velocity response to perturbations in the initial conditions [105]. These responses are:

$$\begin{bmatrix} \frac{\hat{p}}{\rho c} \\ \hat{u} \end{bmatrix} = \begin{bmatrix} \cos \pi \alpha & -i\beta \sin \pi \alpha \\ -\frac{i}{\beta} \sin \pi \alpha & \cos \pi \alpha \end{bmatrix} \begin{bmatrix} \frac{\hat{p}}{\rho c} \\ \hat{u} \end{bmatrix}_0 \quad (4.15)$$

where $\alpha = ckt\beta/\pi$ and

$$\beta = \begin{cases} \sqrt{\epsilon} & \epsilon \leq 1 \\ \frac{1}{\sqrt{\epsilon}} & \epsilon > 1 \end{cases} \quad (4.16)$$

The preconditioning parameter in this analysis is related to the reference velocity in Equation (4.9) through $U_r = c\beta$. When $\epsilon = 1$, no preconditioning is applied to the solution and $U_r = c$. Based on this, Darmofal and Siu [105] find the velocity perturbation response to an initial pressure perturbation when using the preconditioner $U_r = \sqrt{|\Delta p|/\rho}$ proposed by Weiss et al. [61] is:

$$|\hat{u}|^2 \leq \frac{|\hat{p}_0|}{\bar{\rho}} \quad (4.17)$$

Figure 4.2 shows the reference Mach number, U_r/c , and the perturbation Mach number response, $|\hat{u}|/\bar{c}$, for increasing values of initial pressure perturbations $|\hat{p}_0|/\bar{\rho}c^2$ using definitions (A) and (D) from Table 2.1. The changes in slope indicate when one term takes control of the reference velocity; for small perturbation strengths, the preconditioner is controlled by definition (A) and $|\hat{u}|/\bar{c}$ is three orders of magnitude larger than $|\hat{p}|/\bar{\rho}c^2$. As the perturbation strength increases, definition (D) becomes the controlling term and “undoes” the effect of

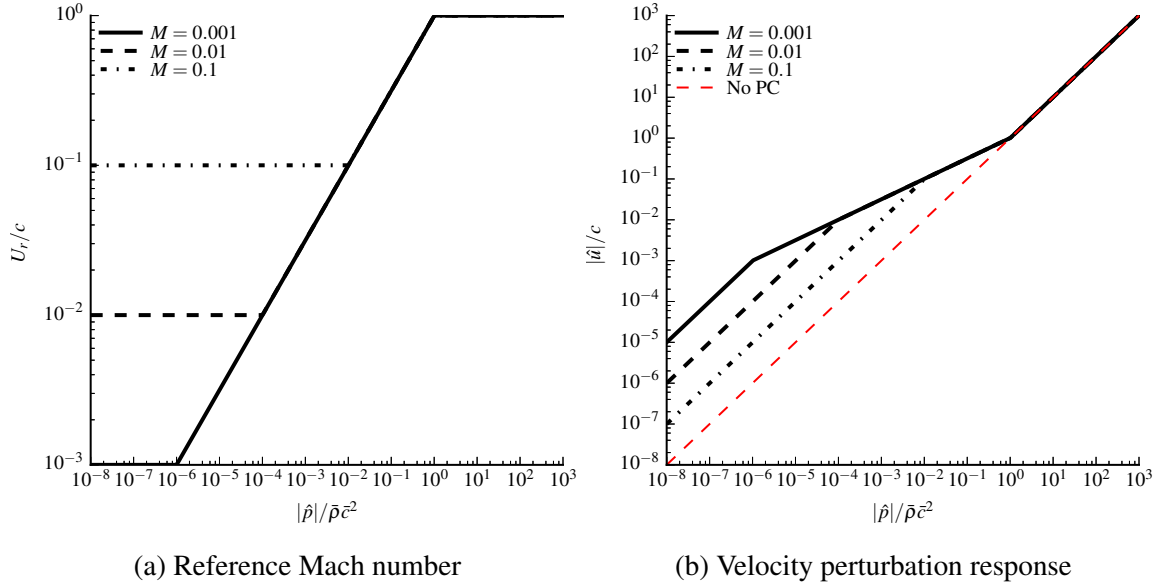


Figure 4.2: Reference Mach number and the velocity perturbation response for increasing pressure perturbation strengths when using definitions (A), (D), and (G) in Table 2.1 for flows with different mean Mach numbers, M .

preconditioning by driving the reference velocity towards \bar{c} . At pressure perturbations of 1 atm, definition (D) exceeds \bar{c} and the preconditioning is completely turned off through (G).

Sankaran et al. [104] showed the definition for the reference velocity using definitions (A) and (D) still exhibits a lack of robustness near stagnation regions for steady, inviscid flows. They proposed a solution to this issue by setting the local reference velocity for each cell equal to the maximum reference velocity of the local cell and all of its face neighbors. Tests reported later in this section with pressure perturbations suggest that this correction does not enhance the robustness of the scheme away from stagnation points. However, by including regional effects into the local definition, it does eliminate the need for global definitions of the reference velocity.

Solution limiting in the presence of non-linearities may be achieved using the non-linear time step controls of Lian et al. [60]. The local reference quantities used to limit the response

for the pressure and velocity are [60]:

$$\begin{aligned} p_{ref} &= \min \left(\frac{\rho u^2}{2}, |\Delta p| \right) \\ u_{ref} &= \min \left(|u|, \frac{c}{\gamma} \frac{|\Delta p|}{p} \right) \end{aligned} \quad (4.18)$$

where γ is the ratio of specific heats. Lian et al. [60] limited the timestep to ensure that $|\delta p| \leq \alpha p_{ref}$ and $|\delta u| \leq \alpha u_{ref}$, where α is a user-defined parameter taken to be 0.1 for the pressure term and the $c|\Delta p|/(p\gamma)$ component of the velocity limit and 2.0 for the $|u|$ component of the velocity limit. New preconditioning definitions are derived here by using the Fourier analysis of Darmofal and Siu [105] to find values of the reference velocity that will enforce these limits prior to computing the time step rather than after. In this analysis, the limits from Lian et al. are used as upper bounds on the velocity and pressure perturbation responses in Equation (4.15).

Limiting the velocity perturbation response such that $|\hat{u}| = \alpha_{u1} |\bar{u}|$ gives a preconditioner reference velocity $U_r = |\Delta p| / (\alpha_{u1} \rho |u|)$. For small pressure perturbations, $\Delta p \propto \rho u^2$, this choice gives $U_r \propto u/\alpha_{u1}$ and this term behaves similarly to the existing steady preconditioning term (A); for large pressure perturbations, $\Delta p \propto \rho c^2$ and thus $U_r \propto c/(\alpha_{u1} M)$ and the preconditioner is disabled as required. Limiting the pressure perturbation response to initial velocity perturbations, $|\hat{p}| = \alpha_{p1} \rho |\hat{u}|/2$ gives a unique type of preconditioner based on the local velocity gradient, $U_r = |\Delta u| \alpha_{p1}/2$. The performance of this term is not evaluated and testing its usefulness is a future task. From Equation (4.15), the limitation $p_{ref} \leq |\Delta p|$ cannot be enforced by varying the preconditioning definition and is left unenforced.

The final limit to consider is $|\hat{u}| = \alpha_{u2} c |\Delta p| / (p\gamma)$ and the resulting reference velocity is $U_r = \alpha_{u2} c$. When $M \leq \alpha_{u2}$ and the pressure perturbations are small, this term controls the preconditioner and it is larger than the typical steady preconditioner, $U_r = |u|$; this is not optimal for these conditions and degrades performance at low Mach numbers. The reference Mach number and the velocity perturbation response using these new definitions

in combination with definitions (A), (D), and (G) are shown in Figure 4.3a and Figure 4.3b when $\alpha_{u1} = 2.0$ and $\alpha_{u2} = 0.1$. The new definitions reduce the amount of preconditioning at lower pressure perturbations, indicated by the shift to the left relative to the original definitions in Figure 4.3a. They also reduce the amount of preconditioning for lower Mach numbers due to the $\alpha_{u2}c$ term indicated by the shift upwards relative to the original definitions, and undo the preconditioning more rapidly as the pressure gradient rises, indicated by the increased slope relative to the original. When $\alpha_{u1} > 1$, there is a small region as the Mach number increases where definition (D) still activates; otherwise, definition (D) is superseded. Figure 4.3c and Figure 4.3d show the reference Mach number and velocity perturbation response when the $U_r = \alpha_{u2}c$ term is omitted and $\alpha_{u1} = 1$, which represents the typical use of these terms. These new definitions for the reference velocity are tested using simulations of pressure disturbances and the choice of included terms and coefficient values is justified. Table 4.1 provides a complete listing of all terms to be tested and is an extension of Table 2.1 including the newly developed reference velocities.

4.2.3 Pressure Perturbation Tests

A one-dimensional domain 10 cm in length is discretized with 100 uniform cells and the left and right boundaries impose a zero-gradient condition on all variables. The flow is initially uniform with a temperature of 300 K and a velocity $\mathbf{u} = \{Mc, 0, 0\}^T$ where c is the speed of sound and the convective Mach number is $M \in \{0.001, 0.01, 0.1\}$. The pressure is initialized as a Gaussian pulse centered in the domain ($x_c = 5$ cm) according to $p(x) = p_\infty + \delta p \exp [(-(x - x_c)/\sigma)^2]$ where δp is the magnitude of the perturbation, $p_\infty = 1$ atm and $\sigma = 1$ cm. The pressure perturbation δp is increased by factors of ten times ρu^2 until the preconditioned scheme is unstable for each Mach number. All other conditions are held the same. The objective is not to find the exact pressure perturbation that causes instability; rather, the order of magnitude of the maximum stable pulse is sufficient to understand the robustness limits and this is the number reported for each case. In addition

Table 4.1: Definitions for the common limiting values of the reference velocity used in Equation (2.113) to determine the preconditioning parameter.

	Definition	Local/Global	Purpose
(A)	$ \mathbf{u} $	Local	Steady
(B)	$\frac{l}{\pi\Delta t}$	Local	Unsteady
(C)	$\frac{\nu S_F}{V}$	Local	Viscous
(D)	$\sqrt{\frac{ \Delta p }{\rho}}$	Local	Pressure perturbations
(E)	κV_∞	Global	Stagnation point
(F)	$\sqrt{\frac{ p_g }{\rho}}$	Global	Pressure perturbation
(G)	c	Global	Speed of sound
(H)	$\frac{ \Delta p }{\alpha_{u1}\rho u }$	Local	Pressure perturbations
(I)	$\alpha_{u2}c$	Local	Pressure perturbations

to varying the magnitude of the perturbation, the terms included in the preconditioning definition, listed in Table 4.1, are turned on and off in combinations to understand the influence of each. The smoothing of the local time step size and the modification of the local reference velocity based on the values in the neighborhood of a cell are referred to as the $\Delta\tau_s$ modification and the maximum U_r modification, respectively.

The GMC-PC scheme is run to steady-state by setting the physical time step to $\Delta t = 1 \times 10^{20}$ s and the pseudo-time iterations are continued until the residual reduces below 1×10^{-9} . The residual used to measure convergence is the L_2 -norm of the change in the solution vector at each time step:

$$\text{Resid.} = \left[\frac{1}{L \times I \times J \times K} \sum_l^L \sum_i^I \sum_j^J \sum_k^K (Q_{i,j,k,l}^{m+1} - Q_{i,j,k,l}^m)^2 \right]^{1/2} \quad (4.19)$$

where I, J, K are the number of cells in each computational coordinate, L is the number of variables in the solution vector, m is the pseudo-time step index, and Q_l is the l -th variable.


Table 4.2 reports the order of magnitude of the maximum stable pulse strength for all of the combinations of implementation details and preconditioning definitions. Without using the $\Delta\tau_s$ modification or the maximum U_r modification, definition (A) on its own is only stable for perturbations on the order of ρu^2 as demonstrated by the first row of Table 4.2; this is consistent with the findings reported by others [3, 61, 104, 105]. Definition (D), in concert with definition (A) and (G), does not enhance the stability in the presence of pressure perturbations. On the other hand, the global pressure definition (F) combined with definition (A) and (G) does enhance the stability by an order of magnitude for each Mach number, although the results are still less than desirable. The combination of definitions (A), (D), (F), and (G) improves the stability for the $M = 0.1$ case by another order of magnitude.

Taking the maximum U_r makes the robustness independent of the choice between local and global pressure limits. This allows the removal of the global limits from Table 4.1. Combining both the $\Delta\tau_s$ modification and the maximum U_r gives the advantages of both eliminating the need for global preconditioning definitions while also providing Mach independence.


The new possible definitions for the reference velocity, (H) and (I) are added to the original definitions (A), (D), and (G). This results in an order of magnitude improvement in the maximum stable pulse strength relative to the original definitions and the behavior is the same for all Mach numbers tested as demonstrated in Table 4.3. For $M \leq \alpha_{u2} = 0.1$ and small pressure perturbations, the new definition (I) dominates the traditional definition (A) and results in less preconditioning and slower convergence. Figure 4.4 illustrates the difference in convergence that results from including that term or not including it for the $M = 0.001$ case with a pressure pulse of $O(10)$. Initially the pressure perturbation is within the domain and the pressure terms control the convergence, resulting in similar behavior with and without the term. Once the pulse has left the domain and the remaining perturbations

are small, the new definition causes a marked flattening of the residual. Stability is not affected while the convergence is orders of magnitude faster by omitting this term. Lastly, when $\alpha_{u1} = 1$, the new pressure definition (H) is always larger than the original pressure definition (D) and the latter may be omitted. The final reference velocity definition that enhances robustness without the need for additional, more complex controls required by others [60, 104], is:


$$U_r = \min \left(\max \left(|\mathbf{u}|, \frac{l}{\pi \Delta t}, \frac{\nu S_F}{V}, \frac{|\Delta p|}{\rho |\mathbf{u}|} \right), c \right) \quad (4.20)$$




(a) Reference Mach number using Lian et al.'s values for α_{u1} and α_{u2} .



(b) Velocity perturbation response using Lian et al.'s values for α_{u1} and α_{u2} .



(c) Reference Mach number using $\alpha_{u1} = 1$ and neglecting the $U_r = \alpha_{u2}c$ limit.



(d) Velocity perturbation response using $\alpha_{u1} = 1$ and neglecting the $U_r = \alpha_{u2}c$ limit.

Figure 4.3: Reference Mach number and the velocity perturbation response for increasing pressure perturbation strengths, comparing original preconditioner (black) and the preconditioner augmented with limits based on the non-linear time step control from Lian et al. [60] (blue) with different mean Mach numbers, M .

Table 4.2: Pulse strength order of magnitude for the maximum stable pressure pulse measured in atmospheres for the various combinations of preconditioning definitions from Table 2.1

(a) Without $\Delta\tau_s$ or max. U_r

U_r	$M = 0.001$	$M = 0.01$	$M = 0.1$
(A), (G)	$O(1e - 6)$	$O(1e - 4)$	$O(1e - 2)$
max (A), (D), (G)	$O(1e - 6)$	$O(1e - 4)$	$O(1e - 2)$
max (A), (F), (G)	$O(1e - 5)$	$O(1e - 3)$	$O(1e - 1)$
max (A), (D), (F), (G)	$O(1e - 5)$	$O(1e - 3)$	$O(1)$

(b) Without $\Delta\tau_s$; with max. U_r

U_r	$M = 0.001$	$M = 0.01$	$M = 0.1$
(A), (G)	$O(1e - 6)$	$O(1e - 4)$	$O(1e - 2)$
max (A), (D), (G)	$O(1e - 5)$	$O(1e - 3)$	$O(1e - 1)$
max (A), (F), (G)	$O(1e - 5)$	$O(1e - 3)$	$O(1e - 1)$
max (A), (D), (F), (G)	$O(1e - 5)$	$O(1e - 3)$	$O(1e - 1)$

(c) With both $\Delta\tau_s$ and max. U_r

U_r	$M = 0.001$	$M = 0.01$	$M = 0.1$
(A), (G)	$O(1e - 5)$	$O(1e - 3)$	$O(1e - 1)$
max (A), (D), (G)	$O(1)$	$O(1)$	$O(1)$
max (A), (F), (G)	$O(1)$	$O(1)$	$O(1)$
max (A), (D), (F), (G)	$O(1)$	$O(1)$	$O(1)$

Table 4.3: Comparison in maximum stable pulse strength between traditional preconditioning using definitions (A) and (D) from Table 4.1 and the new preconditioning definitions designed to enforce the same limits as the non-linear time step control from Lian et al. [60].

U_r	$M = 0.001$	$M = 0.01$	$M = 0.1$
(A), (D), (G)	O(1)	O(1)	O(1)
(A), (D), (G), (H), (I)	O(10)	O(10)	O(10)




Figure 4.4: Residual for the $M = 0.001$ case at the maximum stable pressure pulse with and without the $U_r = \alpha_{u2c}$ preconditioner definition.

CHAPTER 5

NON-REACTING SCHEME VALIDATION

Numerous test cases may be used to validate the generalized MacCormack scheme and provide an assessment of its performance with and without preconditioning. This chapter presents a set of the most interesting cases that have been selected to highlight various aspects of the scheme. The chapter begins with a simple, demonstrative case chosen to visually indicate the different behaviors of the generalized MacCormack-Single Time (GMC-ST) and generalized MacCormack-Preconditioned (GMC-PC) schemes. The inviscid, steady-state, low Mach number flow over a bump highlights the deficiencies of the GMC-ST scheme due to excessive numerical dissipation in the velocity field and the improved solution using the GMC-PC scheme. The periodic convection of an inviscid, time-accurate vortex at both low and moderate Mach numbers demonstrates the convergence and dissipative properties of the GMC-PC scheme for different physical time steps and other scheme settings. This test also verifies the theoretical prediction that the inherent dissipation of the 1971 version of the MacCormack scheme with preconditioning (MC-PC) reverts to the unpreconditioned behavior in the unsteady limit. The classic lid-driven cavity test case from Ghia et al. [123] verifies the steady-state performance of the GMC-ST and GMC-PC schemes for a viscous flow, while decaying isotropic turbulence evaluates the unsteady performance for a viscous flow. These cases are an expanded presentation of the results published in Gallagher et al. [117].

Time-accurate simulations require the specification of the physical CFL number and the computation of the physical time step, Δt . Wherever referenced in the text, the physical CFL number has a subscript to indicate the velocity scale chosen as a reference for computing the time step. For example, $CFL_V = V\Delta t/\Delta x$ for some velocity scale V . In all cases, the GMC-ST scheme uses a CFL number based on the largest eigenvalue of the system as

required for stability in an explicit scheme, indicated by a subscript $u + c$, ie. $\text{CFL}_{u+c} = 0.5$ for all GMC-ST cases. The GMC-PC scheme is run using either the acoustic velocity, c , or the convective velocity, u and the CFL numbers are reported in terms of CFL_c or CFL_u .

5.1 Flow over Bump

The inviscid flow over a bump in a channel is a configuration used to validate numerical schemes [9, 124–126] and is chosen here, not for quantitative comparisons, but because it clearly illustrates the degradation of solution quality at low Mach numbers with compressible schemes. Figure 5.1 diagrams the computational domain for this configuration. A circular segment with a height equal to 10% of its chord length is located in the center of a channel. In this case, the chord length is 1 m. The top and bottom domain boundaries are slip walls and the left boundary is the Robin inflow and the right, the Robin outflow. The inflow Mach number is $M_\infty = 0.001$ and the outflow pressure is 1 atm. This case is run with the GMC-ST and GMC-PC schemes and both schemes are run until they reach steady state. The mesh is 512×256 cells, with a uniform grid in the streamwise direction over the bump and smooth stretching from the boundaries to the leading and trailing edges as shown in Figure 5.2.

The expected solution is smooth and symmetric about the center of the bump [124]. Figure 5.3 shows the velocity field with isolines of Mach number superimposed. Asymmetries appear in the solution with the GMC-ST scheme in Figure 5.3a, although the largest sources of error in the solution are the oscillations in the flow illustrated by the isolines of Mach number. This is in strong contrast to the solution using the GMC-PC scheme in Figure 5.3b, which is both smooth and symmetric. Despite the simple nature of the case and the qualitative comparisons, the GMC-PC scheme exhibits improved solution quality relative to the GMC-ST scheme due to the proper scaling of the dissipation on the velocity and pressure fields in the GMC-PC scheme.




Figure 5.1: Schematic of the domain for the inviscid flow through a channel with a bump.






Figure 5.2: Mesh used for the inviscid flow over a bump in a channel case containing 512×256 cells.



(a) GMC-ST



(b) GMC-PC

Figure 5.3: Steady-state solutions for the inviscid flow over a bump at $M_\infty = 0.001$, showing the degradation that occurs when compressible schemes (GMC-ST) are applied at low Mach numbers. Isolines of Mach number and velocity contours highlight the smoothness and symmetry of the solution using the preconditioned scheme, GMC-PC.

5.2 Inviscid Vortex Convection

The previous case indicates the dissipative properties of the GMC-ST scheme are undesirable when applied to a low Mach number, inviscid, steady flow. The periodic convection of an inviscid vortex assesses these properties for an unsteady flow and provides more flexibility to explore how different aspects of the GMC-PC scheme influence the results. The computational domain is a square measuring 2 mm per side, shown in Figure 5.4, and contains a uniform grid with 256 cells in each direction. Each of the four computational boundaries are periodic and a vortex is initialized in the center of the domain using the velocity and pressure profiles [127]:

$$\begin{aligned}
 u(x, y) &= \frac{U_r - V_s(x - x_c)}{R^2 \exp(-[(x - x_c)^2 + (y - y_c)^2] / 2R^2)} \\
 v(x, y) &= \frac{V_s(y - y_c)}{R^2 \exp(-[(x - x_c)^2 + (y - y_c)^2] / 2R^2)} \\
 P(x, y) &= P_r \exp\left(-\frac{\gamma}{2} \left(\frac{V_s}{cR}\right)^2 \exp[-2((x - x_c)^2 + (y - y_c)^2)]\right)
 \end{aligned} \tag{5.1}$$

where $U_r = M_b c$ is the convection velocity, V_s is the vortex strength, and (x_c, y_c) are coordinates of the vortex center. The vortex radius, $R = 0.2$ mm, $P_r = 101\,325$ Pa is the reference pressure and $\gamma = 1.4$ is the ratio of specific heats. The vortex convects to the right, through the right boundary and into the domain again on the left. The simulations are stopped when it has returned to its starting position after convecting once through the domain. Two convective Mach numbers are simulated to show the behavior of the GMC-PC scheme relative to the GMC-ST scheme, $M_b = 0.001$ and $M_b = 0.5$.

Moderate Mach number behavior

When the convection speed of the vortex is $M_b = 0.5$, the reference velocity used to define the preconditioning parameter in the GMC-PC scheme is the speed of sound. When this occurs, preconditioning is disabled and the system is unmodified in pseudo-time. In this




Figure 5.4: Computational domain for the inviscid vortex convection case.

configuration, the primary differences between the GMC-ST and GMC-PC schemes are that the latter is an implicit scheme in physical time and that the latter also has additional dissipation on the pressure field as described in Chapter 4. As such, the differences in the solution between the two schemes should be minor.

Figure 5.5 compares the transverse velocity along the streamwise centerline of the vortex for the GMC-ST and GMC-PC schemes with the theoretical solution. Both schemes exhibit minor dissipation in the peak values of velocity relative to the theoretical values, but overall the profiles match well. Minor differences between the schemes are apparent at the leading and trailing edges of the vortex, where the GMC-PC scheme shows an increased amount of numerical dispersion in the solution. This dispersive error is consistent with implicit temporal discretizations. Similar trends are observed in the pressure profile through the center of the vortex in the streamwise direction shown in Figure 5.6. The GMC-PC and GMC-ST schemes match one another well in the core of the vortex. However, as with the velocity profile, the GMC-PC scheme shows more dispersion in the pressure profile than the GMC-ST scheme. Also evident are numerical oscillations in the GMC-ST scheme




Figure 5.5: Velocity profiles across the vortex convecting at $M = 0.5$, using the GMC-ST and the GMC-PC schemes. Reprinted with permission from Gallagher et al.[117].

that are not present in the GMC-PC scheme and are attributed to the new non-linearity switch discussed in Chapter 4 (for a discussion of the pressure field using the original non-linearity switch, see Appendix A). The numerical dissipation added to the pressure field in the GMC-PC scheme suppresses these weak oscillations.

Effect of physical time step selection

At low Mach numbers, the precise behavior of the GMC-PC scheme depends on the physical time step selection through the unsteady preconditioner definition. For small physical time steps, the unsteady preconditioner is dominant and this has the effect of disabling the preconditioning by driving the reference velocity to the speed of sound. For large physical time steps, the steady preconditioner is dominant and the reference velocity converges to the convective velocity. These two limits are represented by physical CFL numbers of $CFL_c = 1$ and $CFL_u = 1$ respectively. With the convective Mach number set to 0.001, $CFL_u = 1 = 1000 CFL_c$.




Figure 5.6: Pressure profiles across the vortex convecting at $M = 0.5$, using the GMC-ST and the GMC-PC schemes.

In the analysis of the traditional MacCormack scheme at low Mach numbers presented in Chapter 3, it is noted that the preconditioning approach could be applied directly to the 1971 MacCormack scheme and this scheme is called MC-PC. The modified equation analysis showed that the inherent dissipation scales with the preconditioned eigenvalues and therefore, when preconditioning is active, the inherent dissipation is well-behaved. However, in the case of small physical time steps, the preconditioning is deactivated and the dissipation properties of the MC-PC scheme are predicted to revert back to those of the unmodified 1971 MacCormack scheme. Running the low Mach number vortex using the MC-PC scheme with a physical CFL number of $CFL_c = 1$ verifies this to be the case as shown in Figure 5.7. There is little difference in the solution using the GMC-ST and the MC-PC schemes, and both show a reduction in the windward side peak of approximately 50% and a reduction in the leeward side peak of approximately 30% relative to the theoretical profile. On the other hand, the modifications made to the GMC-PC scheme – specifically the change from biasing the pressure fluxes to a central interpolation coupled with an artificial dissipation on the pressure field – greatly enhances the solution quality. The GMC-PC scheme again




Figure 5.7: Velocity profiles across the vortex convecting at $M = 0.001$ illustrating the limitations of applying preconditioning to the original MacCormack scheme (MC-PC) when run at a $CFL_c = 1$. The improved scheme (GMC-PC) shows little numerical dissipation under these conditions. Reprinted with permission from Gallagher et al.[117].

shows minor dispersion, but otherwise matches the theoretical profile very well and exhibits little numerical dissipation. This suggests the GMC-PC scheme properties are independent of Mach number.

The GMC-PC scheme is also run at a larger time step, $CFL_u = 1$, and the resulting velocity profile relative to the GMC-ST scheme and GMC-PC scheme at $CFL_c = 1$ are in Figure 5.8. With a physical time step 1000 times larger than when run at $CFL_c = 1$, the results of the GMC-PC scheme at $CFL_u = 1$ show an increased amount of numerical dissipation, as would be expected. However, the numerical dissipation is far less than that seen in the GMC-ST scheme. The leeward side peak shows little dissipation relative to the theoretical, while the windward side peak exhibits less than half of the dissipation seen in the GMC-ST scheme.




Figure 5.8: Velocity profiles across the vortex convecting at $M = 0.001$ for the GMC-PC scheme at two different physical CFL numbers compared to the GMC-ST and theoretical profiles. Both limiting cases of the GMC-PC scheme show improved dissipative properties relative to the GMC-ST scheme at these conditions.

Effect of convergence criterion

The dual-time iteration procedure stops based on user-defined convergence criteria, at which point the converged solution is the solution at the next physical time level (see Chapter 4 for details). There are multiple possible criteria from which the user may select, however experience has shown the most reliable and practical is stopping convergence when the residual reduces by a user-specified number of decades (orders of magnitude). Assessment of the effect this criterion has on the solution quality indicates the solution for this case is insensitive to the number of decades specified, provided the pseudo-time derivative is converged *enough*. Figure 5.9 shows the velocity profiles for the GMC-PC scheme at $CFL_u = 1$ for convergence criteria set to 2, 3 and 4 decades of residual reduction. The solutions are indistinguishable from one another in all cases.

However, simulations run with the convergence criterion under 2 decades diverged. At 0.5 decades, the simulation crashes within 4 physical time steps while at 1.5 decades, it

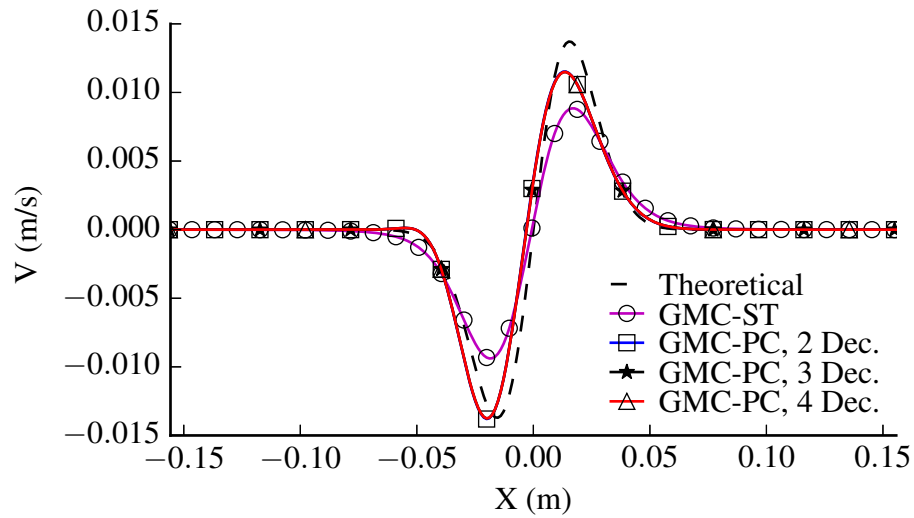


Figure 5.9: Velocity profiles across the vortex convecting at $M = 0.001$, using the GMC-ST and GMC-PC schemes with various convergence criteria. Reprinted with permission from Gallagher et al.[117].

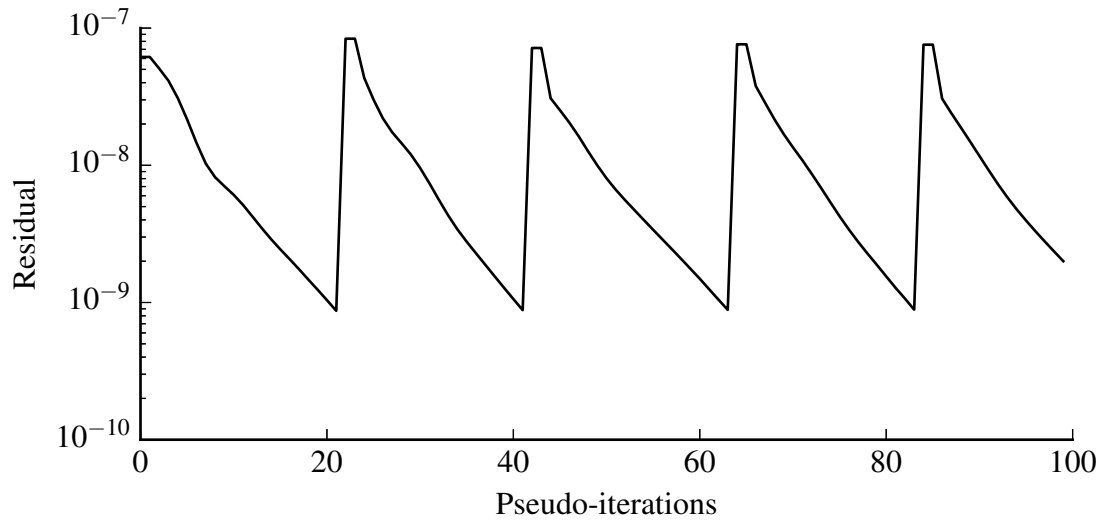
crashes after 200 physical time steps out of the 256 steps required. Experience with other, more complex cases also seems to indicate that insufficient convergence leads to divergence in physical time that may take a long time to manifest if only weakly insufficient.

Scheme efficiency


The previous results indicate the GMC-PC scheme has superior dissipative properties relative to the GMC-ST scheme but no consideration has been given to the cost. On this grid, the GMC-PC scheme at $CFL_c = 1$ needs approximately 20 pseudo-time iterations per physical time iteration to converge 4 decades. The convergence behavior is shown in Figure 5.10a where the discrete jumps in the residual indicate the transition to the next physical time step and the restarting of the dual-time iterations. At this CFL number, the GMC-PC scheme requires approximately 10 times the number of iterations as the GMC-ST scheme, which runs at $CFL_{u+c} \approx CFL_c = 0.5$. If the convergence criterion were reduced by half, the number of pseudo-iterations would be reduced by approximately half and the workload

reduced to a factor of 5 times the GMC-ST scheme. Likewise, when run at a CFL number of $CFL_u = 1$, the GMC-PC scheme requires approximately 2000 pseudo-iterations to converge 4 decades as shown in Figure 5.10b. This is consistent with the ratio of CFL numbers and the scheme when run with this criterion is approximately 10 times more expensive. However, if the convergence criterion is reduced to 2 decades, this requires approximately 500 pseudo-iterations and the overall cost is reduced to approximately 2.5 times the number of iterations as the GMC-ST scheme.

Despite the increased number of iterations in the GMC-PC scheme relative to the GMC-ST scheme, the solution quality is significantly improved at low Mach numbers. For a simple configuration such as this, it is possible to find a grid resolution that would provide comparable results using the GMC-ST scheme. In this instance, it requires a grid with three times as many cells in each direction – a resolution of 768×768 cells. This also reduces the time step size for the GMC-ST scheme by the same factor, resulting in 27 times the computational work relative to the scheme on the original grid. Figure 5.11 shows the GMC-ST scheme on this refined grid relative to the GMC-PC scheme at $CFL_c = 1$ and the results are comparable. From this, the GMC-PC scheme is approximately 3 to 6 times more efficient (based on convergence criteria) than the GMC-ST scheme for the same solution quality and it is capable of producing a Mach number independent solution on the original grid. When run at $CFL_u = 1$, the GMC-PC scheme has similar efficiency as when run at $CFL_c = 1$, but reducing the convergence criteria has a larger effect on the reduction of iterations. Therefore, this scheme can be up to an order of magnitude more efficient than the GMC-ST scheme.



(a) $CFL_c = 1$



(b) $CFL_u = 1$

Figure 5.10: Convergence history for several physical time iterations of the vortex convection case at two physical CFL numbers. Discontinuities in the residual indicate the transition to the next physical-time iteration. Reprinted with permission from Gallagher et al.[117].

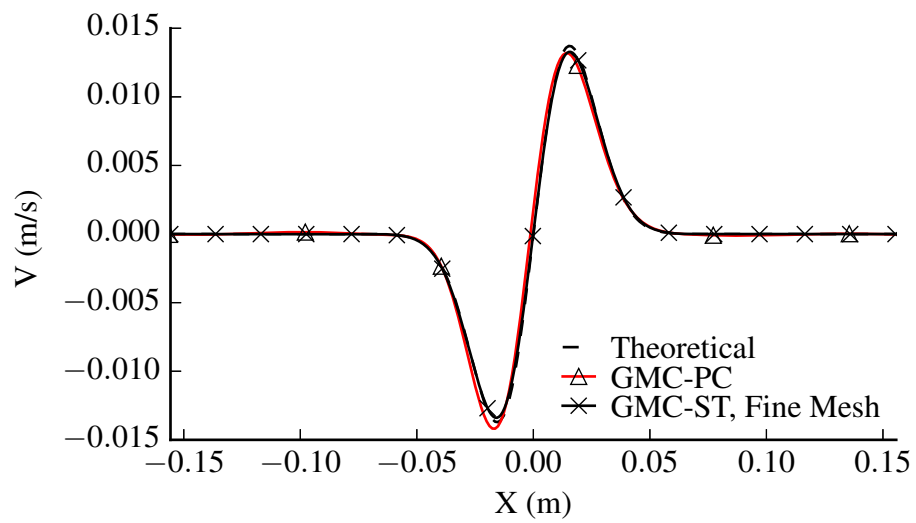


Figure 5.11: Comparison of the solutions using the GMC-ST and GMC-PC schemes, where the GMC-ST solution requires three times finer mesh in both directions for comparable accuracy to the GMC-PC solution on the original mesh. Reprinted with permission from Gallagher et al.[117].

5.3 Lid-Driven Cavity

The lid-driven cavity is a classic case designed to test the accuracy of numerical schemes for viscous fluids. Numerical data are available for multiple Reynolds numbers and the resulting flows range from simple, laminar flows to highly unsteady, turbulent flows with multiple eddies and separation points [123, 128]. The bump flow and vortex flows considered previously focus on the behavior of the GMC-PC scheme for inviscid steady and unsteady flows respectively. The lid-driven cavity cases in this section focus on the viscous steady-state performance of the GMC-PC scheme.

The test employs a square cavity, bounded on all sides by no-slip walls. The top wall is impulsively set into motion, dragging the fluid along with it. There are four main features that develop [123]. The primary vortex forms towards the top-right corner (for a rightward moving top wall), with the core moving towards the center of the domain as the Reynolds number increases. The two corner eddies at the bottom left and bottom right of the domain also increase in size as the Reynolds number increases, with the bottom right corner eddy growing at a larger rate and to a larger size than the left corner eddy. For Reynolds numbers less than or equal to 1000, the corner eddies are each a single, large eddy. At higher Reynolds numbers, the lower right eddy bifurcates with a second, small eddy forming in the corner. In the original Ghia et al. dataset, the lower right eddy bifurcates at a Reynolds number of 3200 and the size of the secondary eddy increases with the Reynolds number [123]. The left corner eddy also bifurcates at Reynolds numbers greater than or equal to 7500. A separation bubble on the top of the left wall is the remaining topological feature. For Reynolds numbers under 1000, this bubble does not form, while at a Reynolds number of 1000 the bubble begins to develop, yet remains weak. Its size and strength increase as the Reynolds number increases.

In addition to different topological evolution with increasing Reynolds numbers, the amount of unsteadiness in the flow increases as well. For $Re \leq 1000$, the fully-developed

flow has a steady-state solution. At larger Reynolds numbers, the flow is unsteady and periodic in time, so no steady-state solution exists (although the time-averaged flow field does converge). Because the interest is in evaluating the steady-state performance of the GMC-PC scheme, two Reynolds numbers are tested: $Re = 100$ and $Re = 1000$. The Mach number, based on the velocity of the top wall, is also varied and $M = 0.3$ and $M = 0.001$ configurations are tested. Provided the Mach number is low enough that compressibility in the flow is not significant, the results should be independent of the Mach number.

The Reynolds number and Mach number are defined to be $Re = u_{\text{wall}}L/\nu$, where the length scale is the size of the domain, $L = 1$ m, and $M = u_{\text{wall}}/c$, respectively. The solution is initialized with stagnant conditions at $p = 1$ atm and $T = 298.15$ K. The wall boundary conditions are no-slip, with an isothermal condition for temperature such that $T_{\text{wall}} = 298.15$ K. For incompressible solvers, this temperature boundary condition is not important as the temperature equation is decoupled from the system and there is no increase in temperature. For compressible solvers at low Mach numbers, the kinetic energy from the wall dissipated through viscosity in the fluid is generally small enough that the temperature increases in the domain are not enough to change the velocity field. In that instance, adiabatic wall boundary conditions are possible. However, at higher Mach numbers, including $M = 0.3$, the amount of kinetic energy dissipated is enough to cause a large temperature increase throughout the fluid. This will change the density as well as the transport properties and generate a different flow field than the original data from Ghia et al. [123], who used an incompressible solver. The isothermal wall conditions act as an energy sink to remove the dissipated kinetic energy. To ensure consistency in the simulation details, both Mach numbers and Reynolds numbers use isothermal walls.

Iterative convergence properties

The solution is considered to have reached steady state when one of the convergence criteria given in Chapter 4, repeated here for clarity, is satisfied:

- (i) Simulation reaches a maximum number of iterations for that physical time step;
- (ii) Residual reduces below a maximum norm value;
- (iii) Residual has decayed a specified number of decades;
- (iv) Change in residual is less than a specified amount for several consecutive iterations.

Criterion (i) limits the number of pseudo-iterations in the GMC-PC scheme to 1×10^6 and limits the maximum physical time for the simulation to 20 s for the GMC-ST scheme. For both schemes, Criterion (ii) terminates the simulation when the norm of the residual drops below 1×10^{-9} . Criterion (iii) is removed from consideration by choosing a large number of decades, 100 in this case. Lastly, the simulation is considered stalled through Criterion (iv) when the difference in the residual is less than 1×10^{-12} for more than 10 consecutive iterations.

Table 5.1 provides the final residual reached and the termination used to halt the convergence for each of the grids and configurations tested. On all grids and for all conditions, the GMC-ST scheme converges close to the desired residual. At $M = 0.3$, the scheme reaches the target residual on all grids for both Reynolds numbers. However, at low Mach numbers, convergence is generally stopped due to the residual stalling. The opposite happens for the GMC-PC scheme. Specifically, convergence stalls for all grids and both Reynolds numbers at $M = 0.3$, while the desired residual is reached for all cases at $M = 0.001$. Unlike the GMC-ST scheme, the cases where the GMC-PC scheme convergence stalled have a final residual much larger than the target. The cause of this discrepancy is not known, but as will be shown, the final flow field is correct despite the relatively large final residual.

Figure 5.12 shows the convergence history on the finest grids used for each condition and scheme. At $M = 0.3$, the GMC-PC and GMC-ST schemes converge at similar rates as the simulation progresses, although the GMC-PC scheme stalls at a relatively large residual. This difficulty converging is less at $Re = 1000$. On the other hand, convergence for the low Mach number case is markedly different. The GMC-PC scheme converges rapidly to the

target residual, while the GMC-ST scheme is slow to converge. For comparison, on the finest grid, the GMC-PC scheme reached the steady-state solution with a wall-time approximately 8.5 times faster than the GMC-ST scheme at $Re = 100$ and required approximately 24 times fewer iterations. At $Re = 1000$, the GMC-PC scheme required approximately 6 times less wall-time and approximately 11 times fewer iterations. This suggests that the cost per iteration is just under 3 times more for the GMC-PC scheme at $Re = 100$ and just under 2 times more at $Re = 1000$. However, as will be shown in the next sections, the GMC-ST scheme is not capable of capturing the correct solution on these grids, making the GMC-PC scheme even more efficient at low Mach numbers.

Quality of solution

Ghia et al. [123] provide velocity contours through the center of the domain to facilitate comparison of numerical schemes. Comparisons between the two schemes and these velocity contours are shown in Figure 5.13. At $M = 0.3$, both the GMC-ST and GMC-PC schemes match the reference profiles well with little difference between the schemes and the expected results. At $M = 0.001$, this is not the case.

Figure 5.14 shows the velocity profiles for both schemes at the low Mach number condition. For the $Re = 100$ case, the GMC-ST scheme predicts a velocity field close to the reference data, but shows an increased amount of dissipation near the peaks while the GMC-PC data matches the reference data well. At $Re = 1000$, the GMC-PC scheme matches the reference data while the GMC-ST scheme deviates from the expected values. These profiles correspond to the coarsest grid resolutions where the GMC-PC scheme first became grid independent. These correspond to the 64×64 and 128×128 resolutions for $Re = 100$ and 1000 , respectively. The steady state velocity contours for $Re = 1000$ at $M = 0.001$ in Figure 5.15 indicate both schemes predict the formation of a primary vortex, but this vortex is located closer to the upper-right corner for the GMC-ST scheme. This predicted location is similar to the primary vortex at a much lower Reynolds number,

indicating the increased numerical dissipation at low Mach numbers in the GMC-ST scheme lowered the effective Reynolds number of the flow.

Grid convergence properties

Multiple grids are used to measure the grid independence and spatial convergence rates of the two schemes. The velocity profiles from Ghia et al. are used to create an error measure for each result. The error is defined as $\epsilon = |u(y_{G,min}) - U_{G,min}|$, where $U_{G,min}$ is the minimum value of the X-component of the velocity profile along the Y-axis according to the reference data from Ghia et al. [123]. The coordinate $y_{G,min}$ is the y -coordinate of that minimum velocity, and $u(y_{G,min}) = u_{min}$ is the velocity component from the present work at that coordinate. This error measure is shown in Figure 5.16, where it should be noted that the grid independent error is seemingly large. The data from Ghia et al. is numerical, using a second order accurate scheme, and therefore convergence to their data within machine precision is not to be expected.

At $M = 0.3$, both the GMC-ST and GMC-PC schemes converge at the same rate to the target data. For reference, the blue dashed line has a slope of 2, indicating second-order accurate convergence rates are achieved for both schemes at $M = 0.3$ as expected. This is true for both Reynolds numbers. At $M = 0.001$, the GMC-PC scheme shows nearly identical convergence rates as it did at $M = 0.3$, again suggesting the performance of the scheme is independent of Mach number. However, the GMC-ST scheme does not perform well at either Reynolds number in the low Mach number configuration and at $Re = 1000$, the error has not yet begun to decrease even on the finest grid tested.

Unlike the vortex test case, it is not practical to find a grid such that the GMC-ST scheme is capable of giving the same quality solution as the GMC-PC scheme for both Reynolds numbers. In the $Re = 100$ case, it appears the GMC-ST scheme enters the convergent range between the 64×64 and 128×128 grids. From this, a simple extrapolation suggests the grid would need to be somewhere between 256×256 and 512×512 to reach the

same solution quality. For the $Re = 1000$ case, the GMC-ST scheme has not yet started to converge. Assuming the convergence were to start immediately for grids larger than 256×256 , simple extrapolation indicates the required grid would be 2048×2048 . Not only does the GMC-PC scheme require less wall-time on a given grid, but it is many orders of magnitude more efficient than the GMC-ST scheme when the accuracy of the predicted solution is considered.

Table 5.1: Convergence information for the lid-driven cavity for each combination of flow conditions, meshes and schemes.

Mesh Size	Reynolds	Mach	GMC-ST		GMC-PC	
			Final Residual	Criterion	Final Residual	Criterion
16 × 16	100	0.3	1.00×10^{-9}	(ii)	2.57×10^{-5}	(iv)
32 × 32	100	0.3	1.00×10^{-9}	(ii)	2.51×10^{-5}	(iv)
64 × 64	100	0.3	1.00×10^{-9}	(ii)	2.62×10^{-5}	(iv)
128 × 128	100	0.3	1.00×10^{-9}	(ii)	2.27×10^{-5}	(iv)
16 × 16	100	0.001	2.07×10^{-9}	(iv)	1.00×10^{-9}	(ii)
32 × 32	100	0.001	1.29×10^{-9}	(iv)	1.00×10^{-9}	(ii)
64 × 64	100	0.001	1.12×10^{-9}	(iv)	1.00×10^{-9}	(ii)
128 × 128	100	0.001	1.00×10^{-9}	(ii)	1.00×10^{-9}	(ii)
16 × 16	1000	0.3	1.00×10^{-9}	(ii)	6.99×10^{-6}	(iv)
32 × 32	1000	0.3	1.00×10^{-9}	(ii)	3.60×10^{-6}	(iv)
64 × 64	1000	0.3	1.00×10^{-9}	(ii)	1.25×10^{-6}	(iv)
128 × 128	1000	0.3	1.00×10^{-9}	(ii)	3.69×10^{-7}	(iv)
256 × 256	1000	0.3	1.00×10^{-9}	(ii)	2.50×10^{-7}	(iv)
16 × 16	1000	0.001	9.66×10^{-9}	(iv)	1.00×10^{-9}	(ii)
32 × 32	1000	0.001	4.66×10^{-9}	(iv)	1.00×10^{-9}	(ii)
64 × 64	1000	0.001	8.56×10^{-9}	(iv)	1.00×10^{-9}	(ii)
128 × 128	1000	0.001	8.47×10^{-9}	(iv)	1.00×10^{-9}	(ii)
256 × 256	1000	0.001	6.46×10^{-9}	(iv)	1.00×10^{-9}	(ii)

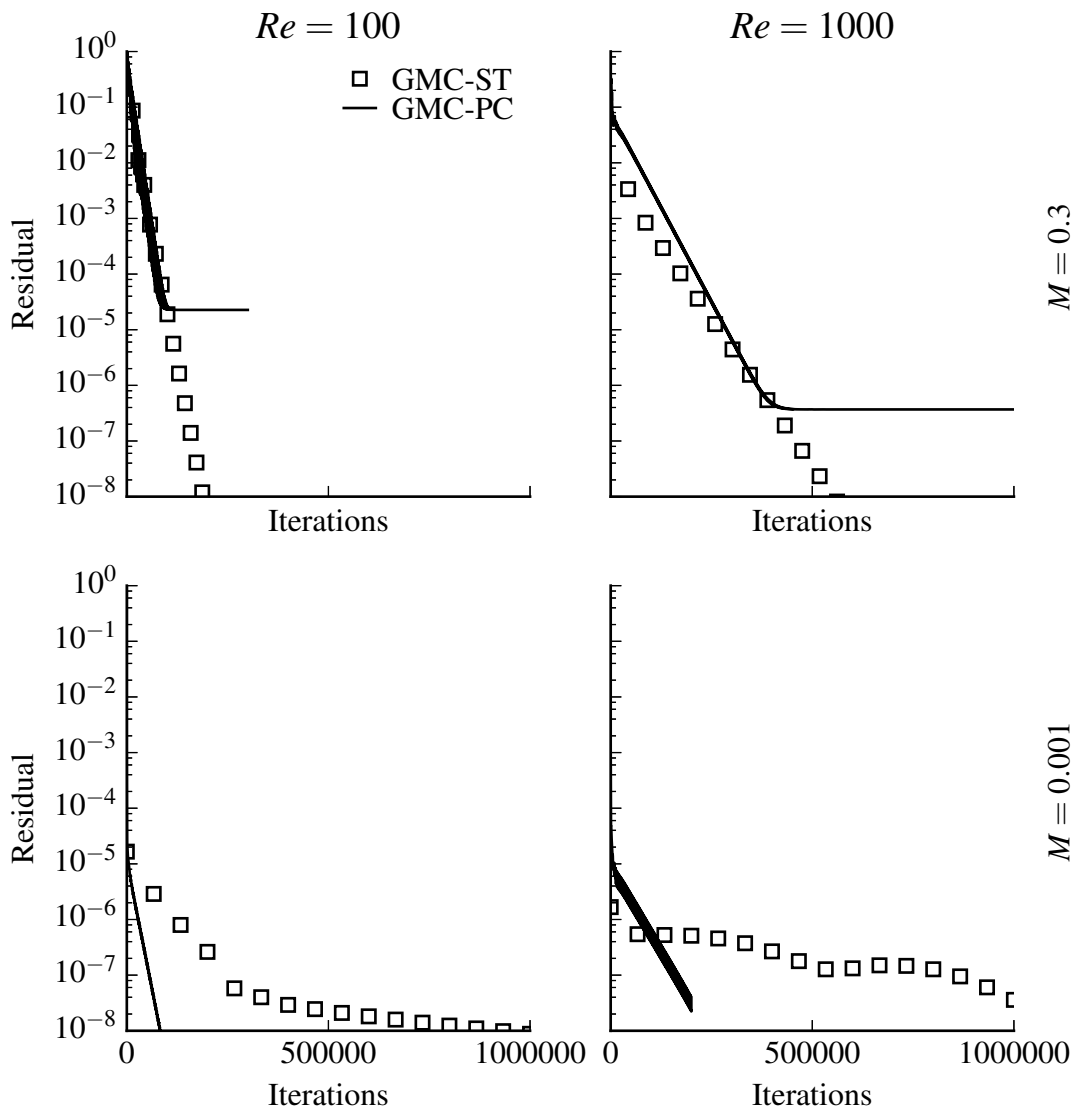




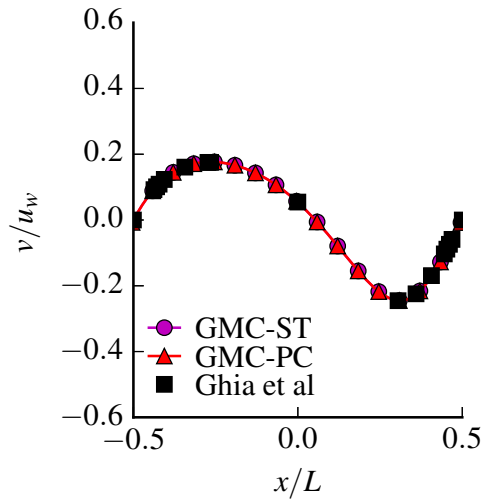
Figure 5.12: Convergence histories for the GMC-ST and GMC-PC schemes on the finest grids for each configuration. Reprinted with permission from Gallagher et al.[117].




(a) U-velocity profiles at $Re = 100$



(b) U-velocity profiles at $Re = 1000$




(c) V-velocity profiles at $Re = 100$




(d) V-velocity profiles at $Re = 1000$


Figure 5.13: U- and V-velocity profiles for both Reynolds numbers at $M = 0.3$ indicating both the GMC-PC and GMC-ST schemes are capable of capturing the correct solution.




(a) U-velocity profiles at $Re = 100$



(b) U-velocity profiles at $Re = 1000$



(c) V-velocity profiles at $Re = 100$



(d) V-velocity profiles at $Re = 1000$

Figure 5.14: U- and V-velocity profiles for both Reynolds numbers at $M = 0.001$ demonstrating the inability of the GMC-ST scheme to capture the correct results. Reprinted with permission from Gallagher et al.[117].




Figure 5.15: Steady-state velocity contours in the lid-driven cavity for the $M = 0.001$ at $Re = 1000$ using the GMC-PC scheme (filled contour) and the GMC-ST scheme (lines). Contours are at 17 levels from $|u| = 0$ (light) to $|u| = u_{\text{wall}}$ (dark). Reprinted with permission from Gallagher et al.[117].

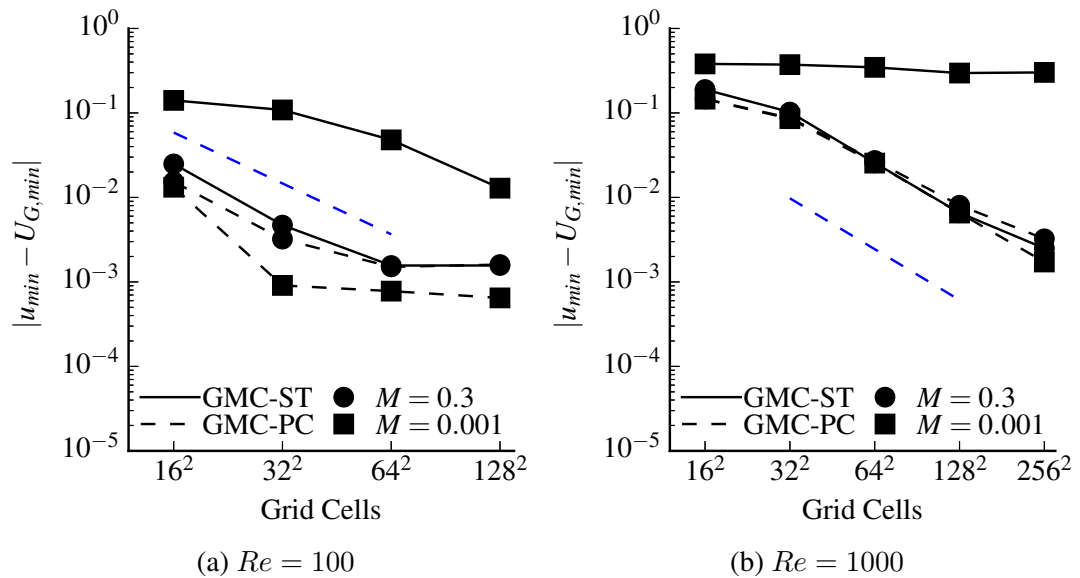


Figure 5.16: Error in velocity relative to reference data [123] for each mesh tested. The blue dashed line represents second-order accurate convergence. Reprinted and adapted with permission from Gallagher et al.[117].

5.4 Decaying Isotropic Turbulence

The final non-reacting flow considered is a viscous, unsteady problem and represents a final building block towards the understanding of the dissipative characteristics of the schemes. Decaying isotropic turbulence is an initial value problem selected to highlight the differences between the GMC-ST and GMC-PC schemes when run at a low turbulent Mach number configuration. By running this case with no turbulent closure models, any differences in the results are due solely to the differences in the numerical dissipation between the two schemes. An initial turbulent field is generated on a uniform box grid with 64^3 cells measuring 6.28 m per side. The von Karman-Pao model spectrum [129] is used to generate the initial turbulence:

$$E(k) = 16 (u')^2 \sqrt{\frac{2}{\pi}} \left(\frac{k^4}{k_0^5} \right) \exp \left(-2 \frac{k^2}{k_0^2} \right) \quad (5.2)$$

where $E(k)$ is the kinetic energy at wave number k and k_0 is the wave number with the peak energy content, set here to $k_0 = 4$. The RMS velocity, u' , is chosen such that the resulting turbulent Mach number is $M_t = 0.001$ and the kinematic viscosity of the fluid is set to $100 \text{ mm}^2 \text{ s}^{-1}$. Based on these conditions, the turbulence is under-resolved and the ratio of the grid cell size to the Kolmogorov length scale is $\Delta x/\eta \approx 20$. Since this is an unsteady problem, the GMC-PC scheme is run at two physical CFL numbers, $\text{CFL}_c = 1$ and $\text{CFL}_u = 1$.

Two metrics are used to indicate the dissipative properties of the schemes. The first is the temporal evolution of the large scale structures. This is represented through the RMS of the vorticity fluctuations, ω_{RMS} . The second metric is the temporal evolution of the Kolmogorov length scale at the smallest structures of the flow. This is defined as $\eta = (\nu^3/\epsilon)^{1/4}$, where $\epsilon = 2\nu \langle s_{ij} s_{ij} \rangle$ with $s_{ij} = 1/2 \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right)$ [129]. Both metrics are tracked as a function of the eddy-turnover time, $\tau_t = L_{11}/\langle u \rangle$ where L_{11} is the integral length scale.

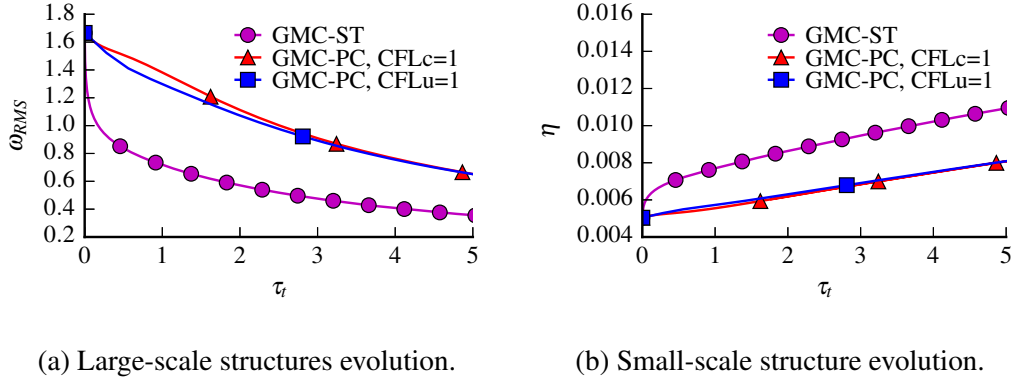


Figure 5.17: Time variation of the vorticity fluctuations and Kolmogorov length scale in decaying isotropic turbulence using the GMC-ST and GMC-PC schemes. Reprinted with permission from Gallagher et al.[117].

The GMC-ST scheme generates a rapid initial decay in the large scale structures of the flow represented by the vorticity fluctuations as shown in Figure 5.17a. The GMC-PC scheme, however, shows no such rapid initial decay. Furthermore, the large scale features of the flow dissipate slightly faster at $CFL_u = 1$, but the differences are minor relative to $CFL_c = 1$. In a similar fashion, the GMC-ST scheme shows an initial rapid increase in the size of the dissipative scales η in Figure 5.17b whereas the GMC-PC scheme does not. The sensitivity to the physical CFL number is even less evident in the size of the dissipative scales. Overall, both the large and small scale features of the flow are primarily insensitive to the physical CFL number when using the GMC-PC scheme.

At an intermediate eddy-turnover time of $\tau_t = 3.5$, the kinetic energy spectra in Figure 5.18 show the presence of a pronounced inertial range in the GMC-PC results while the GMC-ST shows a much smaller range that could be classified as inertial. The kinetic energy content of the largest scales is similar in all three cases, but the energy content in the intermediate scales to the smallest resolved scales is higher with the GMC-PC scheme. This observation is consistent with previous observations of the behavior of preconditioned schemes at low Mach numbers [130]. As with the changes in the characteristic flow sizes in Figure 5.17, the kinetic energy spectra using the GMC-PC in Figure 5.18 are independent




Figure 5.18: Turbulent kinetic energy spectra at $\tau_t = 3.5$ for the GMC-ST and GMC-PC schemes. Reprinted with permission from Gallagher et al.[117].

of the physical CFL number. The increased richness of the turbulent scales is evident in the Z-vorticity contours in the flow (Figure 5.19). Minor differences are noticeable in the GMC-PC scheme depending on the chosen physical CFL number, while there are significant differences in the size and intensity of the structures compared to the GMC-ST scheme. These differences between the GMC-ST and GMC-PC schemes are directly attributed to the improved dissipative characteristics of the GMC-PC scheme. Again, attempts to grid-refine the GMC-ST calculations to match the solution quality of the GMC-PC scheme were not explicitly performed, but we can again conclude that the GMC-PC scheme would provide significant savings in the associated computational expense.

All of the metrics considered thus far indicate the GMC-PC scheme exhibits less numerical dissipation than the GMC-ST scheme. As a result, the velocity correlations in Figure 5.20 show a much larger range in the velocity values for the GMC-PC scheme relative to the GMC-ST scheme, as would be expected. However, the GMC-ST scheme

appears to exhibit some anisotropy in the distribution of velocities. The amount of anisotropy in the solution can be measured through the variance of the velocity components and the covariance between them. At an eddy-turnover time of $\tau_t = 3.5$, the covariance matrices for the u and v velocity components are:

$$\begin{aligned}
 \text{COV}(u, v)_{\text{GMC-ST}} &= \begin{bmatrix} 0.862 \times 10^{-2} & 34.24 \times 10^{-5} \\ 34.24 \times 10^{-5} & 0.829 \times 10^{-2} \end{bmatrix} \\
 \text{COV}(u, v)_{\text{GMC-PC, CFL}_c=1} &= \begin{bmatrix} 1.610 \times 10^{-2} & 5.936 \times 10^{-5} \\ 5.936 \times 10^{-5} & 1.581 \times 10^{-2} \end{bmatrix} \\
 \text{COV}(u, v)_{\text{GMC-PC, CFL}_u=1} &= \begin{bmatrix} 1.691 \times 10^{-2} & 5.647 \times 10^{-5} \\ 5.647 \times 10^{-5} & 1.700 \times 10^{-2} \end{bmatrix}
 \end{aligned} \tag{5.3}$$

with similar matrices for the other components. The diagonal terms in the GMC-ST matrix are approximately half those in both GMC-PC matrices, again confirming the GMC-ST scheme is more dissipative. Any deviation in the off-diagonal components is indicative of anisotropy in the solution [129]. The GMC-ST scheme has off-diagonal components approximately 5.5 times larger than the GMC-PC scheme, which has similar magnitude terms for both CFL numbers. This corresponds to a correlation coefficient of 0.2 for the GMC-ST scheme compared to the 0.06 for the GMC-PC schemes, indicating a stronger preference for anisotropy with the GMC-ST scheme.






Figure 5.19: Z-vorticity (normal to the plane) in an isotropic decaying turbulence field along the center plane for the GMC-PC and GMC-ST schemes at an eddy-turnover time $\tau_t = 3.5$. Reprinted with permission from Gallagher et al. [117].



(a) GMC-PC scheme at $CFL_c = 1$

(b) GMC-PC scheme at $CFL_u = 1$



(c) GMC-ST scheme

Figure 5.20: Velocity correlations for the GMC-PC and GMC-ST schemes at an eddy-turnover time $\tau_t = 3.5$.

5.5 Summary of results

This chapter demonstrates the performance of the GMC-PC scheme vs. the GMC-ST scheme on a series of progressing to cases of interest for turbulent flows. The steady and unsteady performance for both inviscid and viscous configurations indicate that the GMC-PC scheme is more efficient than the GMC-ST scheme in all cases at low Mach numbers. Most importantly for the prediction of complex flow fields, the GMC-PC scheme has dissipative properties that are independent of Mach number. This is in contrast to the GMC-ST scheme, which shows severely degraded solution quality as the Mach number is reduced, and in some cases the scheme cannot predict the correct solution with a reasonable grid resolution.

Although the GMC-PC scheme exhibits less error than the GMC-ST scheme, it is more expensive per iteration due to the computational overhead of solving the linear system at each cell. This overhead causes each iteration in the GMC-PC scheme to be between 1.5 and 3 times more expensive than the comparable GMC-ST iteration. However, in some cases, the GMC-PC scheme requires substantially fewer iterations at low Mach numbers than the GMC-ST scheme and so it is not only more accurate, but also faster. In other cases at low Mach numbers, the GMC-PC scheme will take longer than the GMC-ST scheme on a given grid. When this occurs, the GMC-PC scheme will take more wall-time for the simulation on the same mesh but getting a comparable solution with the GMC-ST scheme may require a finer mesh and, consequently, more total resources.

With turbulent, reacting flows as the objective, the results from the decaying isotropic turbulence provide key insights. Notably, the results with the GMC-PC scheme are insensitive to the physical CFL number used. This is a useful property for LES for reasons that are not immediately obvious. As discussed in Chapter 1, implicitly filtered LES relies on the grid to act as a filter for the solution. This implies that the cutoff wave number in space is $k_{\text{cut}} \approx 1/\Delta x$. At the same time, the temporal integration introduces a cutoff frequency in physical time, $f_{\text{cut}} \approx 1/\Delta t$. In a turbulent flow, each eddy has an associated wave number

in the spatial domain, but it also has an associated frequency in the temporal domain. The spatial and temporal descriptions of an eddy are related, although the relationship is not always clear. However, under some circumstances, Taylor's hypothesis may be used to relate the temporal behavior of the flow directly to the spatial behavior of the flow [129].

If the hypothesis is valid, the frequency and the wave number are directly related:

$$f \propto \langle u \rangle k \quad (5.4)$$

where $\langle u \rangle$ is the mean flow velocity. More specifically, the spatial filter cutoff introduced by the grid can be used to define a frequency filter cutoff also introduced by the grid, $f_{\text{cut}} \propto \langle u \rangle k_{\text{cut}}$. From this, a cutoff time step can be defined:

$$\Delta t_{\text{cut}} = \frac{1}{\langle u \rangle k_{\text{cut}}} = \frac{\Delta x}{\langle u \rangle} \quad (5.5)$$

When the time step used in the integration is smaller than the cutoff time step just defined, the frequency content of the solution contains information corresponding to the sub-filter solution in space and this information is not physically useful. The cutoff time step can be used to define an ideal physical CFL number:

$$\text{CFL}_{\text{cut}} = 1 = \frac{\Delta t_{\text{cut}} \langle u \rangle}{\Delta x} \quad (5.6)$$

This ideal physical CFL number is simply $\text{CFL}_u = 1$ and the GMC-ST scheme is limited by stability constraints to a much smaller CFL number at low Mach numbers. Given the insensitivity to the physical CFL numbers with the GMC-PC scheme, the ideal physical time step for turbulent, non-reacting LES is $\text{CFL}_u = 1$ and this will serve as an upper limit for the reacting flow field tests.

CHAPTER 6

EXTENSIONS FOR STIFF CHEMISTRY

Careful consideration must be given to the treatment of the source terms that arise from chemical reactions. These source terms are often highly non-linear and form a system of stiff differential equations. In addition, the source terms result in a strong, localized heat release for many of the chemically reacting flows of interest. Therefore, numerical schemes must be designed to handle these unique challenges if they are to remain robust and accurate. This chapter begins with a description of how chemical source term integration is handled in single-time schemes using both fully-coupled and operator splitting approaches. Then, these concepts are extended to dual-time schemes and analyzed. This is an expanded presentation of the work in Gallagher and Menon [131].

To facilitate the presentation, the source term vector $\tilde{\mathbf{S}}$ is rearranged into stiff and non-stiff components, $\tilde{\mathbf{S}}_S$ and $\tilde{\mathbf{S}}_{NS}$ respectively. Unless otherwise needed, the tildes and bars are dropped from all terms and the equations apply to both LES and DNS configurations. The stiff vector is composed of the chemical reaction rates and the non-stiff vector is composed of the subgrid modeling terms, if present. Furthermore, the inviscid and viscous fluxes, as well as the non-stiff source term vector, are combined into a single residual vector \mathbf{R} , such that the dual-time and single-time governing equations are:

$$\begin{aligned} \Gamma \frac{\partial}{\partial \tau} \mathbf{Q} + \frac{\partial}{\partial t} \mathbf{W} &= -\mathbf{R} + \mathbf{S}_S \\ \frac{\partial}{\partial t} \mathbf{W} &= -\mathbf{R} + \mathbf{S}_S \end{aligned} \tag{6.1}$$

respectively.

6.1 Chemistry Integration in Single-Time Schemes

Broadly speaking, integration of chemical source terms for single-time schemes uses two different approaches. The first is a fully-coupled approach, where the chemical source terms are integrated alongside the spatial terms using the same temporal integrator. This may be done explicitly or implicitly depending on the formulation of the scheme. These schemes do not suffer from any splitting errors but rely on the temporal integrator to address the stiffness and time step limitations due to the source terms. For implicit temporal schemes, the stiffness is addressed automatically by nature of the scheme, but the presence of eigenvalues with large, positive components limits the time step selection. For explicit temporal schemes, the eigenvalues with positive real components do not impose stability constraints. However, the eigenvalues with negative real components do impose time step constraints for stability as the stiffness of the mechanism may demand time steps much smaller than the stability constraint to ensure accuracy. Both approaches can be successful if the limits on the time steps are taken into account or if techniques to reduce the stiffness are employed.

While the operator splitting approach segregates the workload into highly efficient, dedicated solvers, the approach introduces splitting errors from the decoupling of the source terms and the spatial fluxes [132]. Depending on the precise splitting scheme chosen, these errors have various leading order of truncation errors in physical time. Two first-order splitting methods are possible [132]. The first, Lie splitting, integrates the source terms over the time step and then integrates the fluxes based on the updated solution:

$$\frac{d\mathbf{W}^*}{dt} = \mathbf{S}_{S,0}^*, \quad \mathbf{W}^*(0) = \mathbf{W}_0, \quad t \in [0, \Delta t] \quad (6.2a)$$

$$\frac{d\mathbf{W}^{**}}{dt} = -\mathbf{R}_*^{**}, \quad \mathbf{W}^{**}(0) = \mathbf{W}^*(\Delta t), \quad t \in [0, \Delta t] \quad (6.2b)$$

An alternate splitting is also possible, where the fluxes are done first, followed by the sources:

$$\frac{d\mathbf{W}^*}{dt} = -\mathbf{R}_0^*, \quad \mathbf{W}^*(0) = \mathbf{W}_0, \quad t \in [0, \Delta t] \quad (6.3a)$$

$$\frac{d\mathbf{W}^{**}}{dt} = \mathbf{S}_{S,*}^{**}, \quad \mathbf{W}^{**}(0) = \mathbf{W}^*(\Delta t), \quad t \in [0, \Delta t] \quad (6.3b)$$

In these schemes, the temporal derivative may be solved using whatever techniques are suitable.

To analyze the errors, the right hand sides are linearized such that $\mathbf{S}_{S,m}^n = \mathbf{D}\mathbf{W}^m$ and $-\mathbf{R}_m^n = \mathbf{A}\mathbf{W}^m$ where m is the starting time index and n is the ending time index. When this is done, the truncation error for this scheme may be written as the difference between the solution using the split scheme and the exact solution:

$$\text{T.E.} = (\exp[\mathbf{D}\Delta t] \exp[\mathbf{A}\Delta t] - \exp[(\mathbf{A} + \mathbf{D})\Delta t]) \mathbf{W}_0 \quad (6.4)$$

for the first splitting and:

$$\text{T.E.} = (\exp[\mathbf{A}\Delta t] \exp[\mathbf{D}\Delta t] - \exp[(\mathbf{A} + \mathbf{D})\Delta t]) \mathbf{W}_0 \quad (6.5)$$

for the second. The matrix exponentials are expanded using Taylor series approximations and the terms are collected together:

$$\text{T.E.} = \frac{\mathbf{D}\mathbf{A} - \mathbf{A}\mathbf{D}}{2} \Delta t^2 + O(\Delta t^3) \quad (6.6)$$

for the first splitting and:

$$\text{T.E.} = \frac{\mathbf{A}\mathbf{D} - \mathbf{D}\mathbf{A}}{2} \Delta t^2 + O(\Delta t^3) \quad (6.7)$$

for the second. The global error is first order for both of the possible splittings, unless the operators commute, in which case the error terms cancel. Although interesting, this commutation rarely happens in anything other than model equations [132]. Additionally, despite similar formal truncation errors, the solution quality may suffer if the stiffest operator is not integrated last and so the two splittings are not equivalent in practice [132].

The other first order splitting method is called the *no-time splitting* method [132, 133]. In this scheme, the source terms are first evaluated over the time interval and then the source term is added to the residual and integrated together:

$$\frac{d\mathbf{W}^*}{dt} = \mathbf{S}_{S,0}^*, \quad \mathbf{W}^*(0) = \mathbf{W}_0, \quad t \in [0, \Delta t] \quad (6.8a)$$

$$\frac{d\mathbf{W}^{**}}{dt} = -\mathbf{R}_0^{**} + \frac{\mathbf{W}^*(\Delta t) - \mathbf{W}_0}{\Delta t}, \quad \mathbf{W}^{**}(0) = \mathbf{W}_0, \quad t \in [0, \Delta t] \quad (6.8b)$$

The solution using this scheme is [132]:

$$\mathbf{W}(\Delta t) = \exp[\mathbf{A}\Delta t] \left(\mathbf{I} + \frac{\exp[\mathbf{D}\Delta t] - \mathbf{I}}{\Delta t} \int_0^{\Delta t} \exp[-t\mathbf{A}] dt \right) \mathbf{W}_0 \quad (6.9)$$

which can again be expanded using Taylor series approximations for the matrix exponentials into:

$$\mathbf{W}(\Delta t) = \left[\mathbf{I} + (\mathbf{A} + \mathbf{D})\Delta t + \frac{\mathbf{A}^2 + \mathbf{D}^2 + \mathbf{AD}}{2}\Delta t^2 + O(\Delta t^3) \right] \mathbf{W}_0 \quad (6.10)$$

As before, the truncation error is found by subtracting the exact solution [132]:

$$\text{T.E.} = -\frac{\mathbf{DA}}{2}\Delta t^2\mathbf{W}_0 + O(\Delta t^3) \quad (6.11)$$

which is again a first-order accurate scheme. Unlike the Lie split scheme, however, this scheme always has errors present whether the operators commute or not.

Higher order splittings are possible. Strang splitting is a popular choice and these errors are second-order accurate in physical time [66, 132, 134]. This scheme integrates one operator over half of the time step, then uses that solution as the initial condition to integrate the other operator over the full time step, and finally uses that solution as the initial condition to integrate the first operator over the remaining half time step:

$$\frac{d\mathbf{W}^*}{dt} = \mathbf{S}_{S,0}^*, \quad \mathbf{W}^*(0) = \mathbf{W}_0, \quad t \in \left[0, \frac{\Delta t}{2}\right] \quad (6.12a)$$

$$\frac{d\mathbf{W}^{**}}{dt} = -\mathbf{R}_*^{**}, \quad \mathbf{W}^{**}(0) = \mathbf{W}^*\left(\frac{\Delta t}{2}\right), \quad t \in [0, \Delta t] \quad (6.12b)$$

$$\frac{d\mathbf{W}^{***}}{dt} = \mathbf{S}_{S,**}^{***}, \quad \mathbf{W}^{***}(0) = \mathbf{W}^{**}(\Delta t), \quad t \in \left[0, \frac{\Delta t}{2}\right] \quad (6.12c)$$

As before, the truncation error is:

$$\text{T.E.} = \left(\exp\left[\mathbf{D}\frac{\Delta t}{2}\right] \exp[\mathbf{A}\Delta t] \exp\left[\mathbf{D}\frac{\Delta t}{2}\right] - \exp[(\mathbf{A} + \mathbf{D})\Delta t] \right) \mathbf{W}_0 \quad (6.13)$$

which is again expanded using Taylor series approximations to yield:

$$\text{T.E.} = \left(\frac{1}{12} [\mathbf{A}\mathbf{A}\mathbf{D} - 2\mathbf{A}\mathbf{D}\mathbf{A} + \mathbf{D}\mathbf{A}\mathbf{A}] - \frac{1}{24} [\mathbf{D}\mathbf{D}\mathbf{A} - 2\mathbf{D}\mathbf{A}\mathbf{D} + \mathbf{A}\mathbf{D}\mathbf{D}] \right) \Delta t^3 \mathbf{W}_0 + O(\Delta t^4) \quad (6.14)$$

This scheme is globally second order accurate.

The formal order of accuracy is only one aspect of the actual error encountered using these schemes. In all cases, the truncation errors involve the reaction and flux operators, and they therefore also include the spectral content of those matrices. For many chemical mechanisms, the eigenvalues in the reaction operator may be large and although the formal order of accuracy may be acceptable, the magnitude of the errors may be significant. First order splitting schemes with time steps determined by the chemical eigenvalues will have errors that are $O(1)$, while second order schemes will have errors $O(\Delta t)$.

The way the schemes used to integrate the temporal derivative also influences the order of accuracy. For example, assuming the chemical source terms are integrated using a dedicated, stiff solver such as DVODE [47] and the flux terms are integrated using the MacCormack scheme, there are two possible ways to implement the splitting. To illustrate these approaches, the no-time splitting scheme in Equation (6.8) is used. The first way to implement the splitting is:

$$\frac{d\mathbf{W}^*}{dt} = \mathbf{S}_{S,0}^* \quad (\text{DVODE}) \quad (6.15a)$$

$$\mathbf{W}^{n+1/2} = \mathbf{W}^n - \Delta t \mathbf{R}^n + \Delta t \frac{\mathbf{W}^*(\Delta t) - \mathbf{W}_0}{\Delta t}, \quad (\text{MacCormack, Predictor}) \quad (6.15b)$$

$$\mathbf{W}^{n+1} = \frac{1}{2} \left[\mathbf{W}^{n+1/2} + \mathbf{W}^n - \Delta t \mathbf{R}^{n+1/2} + \Delta t \frac{\mathbf{W}^*(\Delta t) - \mathbf{W}_0}{\Delta t} \right], \quad (\text{MacCormack, Corrector}) \quad (6.15c)$$

This scheme computes the source terms and then holds them constant throughout the MacCormack integration step. When performed this way, the entire scheme is only first order accurate as the error analysis indicated. The second approach to splitting is to perform the operator splitting on each sub-step of the scheme:

MacCormack, Predictor:

$$\frac{d\mathbf{W}^*}{dt} = \mathbf{S}_{S,0}^* \quad (\text{DVODE}) \quad (6.16a)$$

$$\mathbf{W}^{n+1/2} = \mathbf{W}^n - \Delta t \mathbf{R}^n + \Delta t \frac{\mathbf{W}^*(\Delta t) - \mathbf{W}_0}{\Delta t}, \quad (\text{MacCormack Sub-Step}) \quad (6.16b)$$

MacCormack, Corrector:

$$\frac{d\mathbf{W}^{**}}{dt} = \mathbf{S}_{S,n+1/2}^{**} \quad (\text{DVODE}) \quad (6.16c)$$

$$\mathbf{W}^{n+1} = \frac{1}{2} \left[\mathbf{W}^{n+1/2} + \mathbf{W}^n - \Delta t \mathbf{R}^{n+1/2} + \Delta t \frac{\mathbf{W}^{**}(\Delta t) - \mathbf{W}^{n+1/2}}{\Delta t} \right], \quad (\text{MacCormack Sub-Step}) \quad (6.16d)$$

With this approach, each sub-step contains the first-order splitting error. When the sub-steps are combined together, the global order of accuracy improves to second-order. However, the impact of the chemistry stiffness on the solution accuracy remains.

6.2 Fully-Coupled Chemistry in Dual-Time Schemes

The chemical source terms may be integrated using the same time integration scheme as the the pseudo-time iterations in a fully-coupled fashion. This requires linearizing the implicit source terms, forming the Jacobian matrix $\mathbf{D} = \frac{\partial \mathbf{S}_S}{\partial \mathbf{Q}}$. The resulting scheme, performed using the Generalized MacCormack scheme using a two step, backward differencing formula for the physical time derivative is:

$$\begin{aligned} \left(\mathbf{\Gamma}^m + \frac{3}{2} \frac{\Delta\tau}{\Delta t} \frac{\partial \mathbf{W}}{\partial \mathbf{Q}} - \Delta\tau \mathbf{D}^m \right) \Delta Q^{m+1/2} = \\ - \Delta\tau \mathbf{R}^m - \frac{\Delta\tau}{2\Delta t} (3\mathbf{W}^m - 4\mathbf{W}^n + \mathbf{W}^{n-1}) + \Delta\tau \mathbf{S}_S^m \end{aligned} \quad (6.17a)$$

$$\begin{aligned} \left(\mathbf{\Gamma}^m + \frac{3}{4} \frac{\Delta\tau}{\Delta t} \frac{\partial \mathbf{W}}{\partial \mathbf{Q}} - \frac{1}{2} \Delta\tau \mathbf{D}^{m+1/2} \right) \Delta Q^{m+1} = \\ \frac{1}{2} \left[\mathbf{\Gamma}^m (\mathbf{Q}^m - \mathbf{Q}^{m+1/2}) - \Delta\tau \mathbf{R}^{m+1/2} \right. \\ \left. - \frac{\Delta\tau}{2\Delta t} (3\mathbf{W}^{m+1/2} - 4\mathbf{W}^n + \mathbf{W}^{n-1}) + \Delta\tau \mathbf{S}_S^{m+1/2} \right] \end{aligned} \quad (6.17b)$$

with the m index indicating the pseudo-time iteration and the n index indicating the physical-time iteration.

This scheme is only conditionally stable when the source term Jacobian \mathbf{D} has positive eigenvalues and the stable time step is proportional to the inverse of the largest, positive eigenvalue. To avoid the expense of computing the eigenvalues for large systems, the local time step is restricted to the minimum of the time step computed using Equation (4.9) and $(\max_k [|\dot{\omega}_k|/\rho])^{-1}$. Although this will limit the time step for both positive and negative eigenvalues and therefore give non-optimal time steps and convergence rates, it avoids the

need to compute the eigenvalues. The Jacobian matrix is approximated numerically using first order, finite differences in each of the primitive variables.

6.3 Operator Splitting in Dual-Time Schemes

The operator splitting approaches used for single-time schemes can be extended to dual-time schemes in a variety of ways. With a physical-time and pseudo-time derivative, the operator splitting can be applied in either time derivative. The choice of which derivative to split and how the splitting is done has implications for the accuracy of the scheme as will be shown in Chapter 7.

6.3.1 Physical-Time Splitting

The physical-time split scheme employs a second-order, symmetric Strang splitting in physical-time for the integration. This scheme is well-known and well-studied in the context of single-time schemes [66, 132, 135] and was given in Equation (6.12). This scheme integrates the chemistry source terms for a half time-step before and after the integration of the residual term. The DVODE solver is used to integrate the stiff source terms in the half-steps and the dual-time scheme is used to integrate the residual term in the middle step of the scheme:

DVODE:

$$\frac{d\mathbf{W}^*}{dt} = \mathbf{S}_{S,0}^*, \quad \mathbf{W}^*(0) = \mathbf{W}_0 \quad (6.18a)$$

MacCormack Scheme:

Iterate until dual-time converges:

$$\begin{aligned}
& \left(\mathbf{\Gamma}^m + \frac{\Delta\tau}{\Delta t} \frac{\partial \mathbf{W}}{\partial \mathbf{Q}} \right) \Delta Q^{m+1/2} = \\
& \quad - \Delta\tau \mathbf{R}^m - \frac{\Delta\tau}{\Delta t} \left(\mathbf{W}^m - \mathbf{W}^*(t + \frac{\Delta t}{2}) \right) \\
& \left(\mathbf{\Gamma}^m + \frac{1}{2} \frac{\Delta\tau}{\Delta t} \frac{\partial \mathbf{W}}{\partial \mathbf{Q}} \right) \Delta Q^{m+1} = \\
& \quad \frac{1}{2} \left[\mathbf{\Gamma}^m (\mathbf{Q}^m - \mathbf{Q}^{m+1/2}) - \Delta\tau \mathbf{R}^{m+1/2} \right. \\
& \quad \left. - \frac{\Delta\tau}{\Delta t} \left(\mathbf{W}^{m+1/2} - \mathbf{W}^*(t + \frac{\Delta t}{2}) \right) \right]
\end{aligned} \tag{6.18b}$$

DVODE:

$$\frac{d\mathbf{W}^{***}}{dt} = \mathbf{S}_{S,n+1}^{***}, \quad \mathbf{W}^{***}(0) = \mathbf{W}^{n+1} \tag{6.18c}$$

where now the physical-time derivative is discretized using the predicted value at the end of the first half time-step in a single-step, backward difference formula.

The Strang splitting requires the change in the physical time derivative relative to the other schemes. The second step, Equation (6.12c), must represent only the change due to the fluxes and non-stiff source terms and must not include changes due to the stiff source terms [136]. Multi-step physical time derivatives will include portions of the source term vector in the second step, violating this constraint. Therefore, multi-step schemes for the second step will have an error term of $O(1)$ regardless of the time step size [136].

An alternate scheme could be devised where the integration steps are reversed and the residual term is integrated over the half-steps. However, such a scheme is not optimal. The chemistry is, in general, the stiffest component of the system and the integration order in Equation (6.12) will ensure second-order accuracy is maintained for small-enough time steps [132]. Furthermore, the pseudo-time iteration is more expensive than the DVODE integration and performing the pseudo-iterations twice would increase the cost of the scheme.

As discussed previously, this scheme has second-order leading errors due to splitting in physical time [66, 132]. For large physical-time CFL numbers, this error may be unacceptable and lead to unphysical behavior [64, 135]. However, each physical time step requires just two stiff source term integration calls and this may drastically reduce the computational expense of the scheme for moderate CFL numbers.

6.3.2 Pseudo-Time Splitting

The splitting errors in the physical-time split scheme manifest in the physical solution. Splitting the scheme in pseudo-time instead may mitigate the effects of the splitting errors by keeping them at the pseudo-time level only. If traditional operator splitting is done on the pseudo-time level, steady-state splitting errors may persist and corrupt the physical-time solution. To demonstrate this, the no-time splitting approach in Equation (6.8) is performed at the pseudo-time level. This scheme is a direct, dual-time analogue to the existing, single-time scheme employed in the in-house solver LESLIE and uses DVODE to integrate the stiff chemical source terms. In this case, the source terms are coupled to the spatial fluxes using the no-time splitting approach in each sub-step. In each complete pseudo-time step, this scheme is second-order accurate in the pseudo-time step $\Delta\tau$. In particular, these errors are contained in the steady-state solution [65] and lead to a loss of time accuracy at the physical time level. This is demonstrated in Chapter 7.

The predictor and corrector steps of the Generalized MacCormack scheme are each split into two steps:

$$\begin{aligned}
 \frac{d}{d\tau} \mathbf{W}^* &= \mathbf{S}_S^{m+1/2}; & \mathbf{W}^*(\tau) &= \mathbf{W}^m \\
 \left(\mathbf{\Gamma}^m + \frac{3}{2} \frac{\Delta\tau}{\Delta t} \frac{\partial \mathbf{W}}{\partial \mathbf{Q}} \right) \Delta Q^{m+1/2} &= \\
 - \Delta\tau \mathbf{R}^m - \frac{\Delta\tau}{2\Delta t} (3\mathbf{W}^m - 4\mathbf{W}^n + \mathbf{W}^{n-1}) & \\
 + \Delta\tau \frac{\mathbf{W}^*(\tau + \Delta\tau) - \mathbf{W}^*(\tau)}{\Delta\tau} &
 \end{aligned} \tag{6.19a}$$

$$\begin{aligned}
\frac{d}{d\tau} \mathbf{W}^{**} &= \mathbf{S}_S^{m+1}; & \mathbf{W}^{**}(\tau) &= \mathbf{W}^{m+1/2} \\
\left(\Gamma^m + \frac{3}{4} \frac{\Delta\tau}{\Delta t} \frac{\partial \mathbf{W}}{\partial \mathbf{Q}} \right) \Delta \mathbf{Q}^{m+1} &= \\
\frac{1}{2} \left[\Gamma^m (\mathbf{Q}^m - \mathbf{Q}^{m+1/2}) - \Delta\tau \mathbf{R}^{m+1/2} \right. & \\
- \frac{\Delta\tau}{2\Delta t} (3\mathbf{W}^{m+1/2} - 4\mathbf{W}^n + \mathbf{W}^{n-1}) & \\
\left. + \Delta\tau \frac{\mathbf{W}^{**}(\tau + \Delta\tau) - \mathbf{W}^{**}(\tau)}{\Delta\tau} \right] &
\end{aligned} \tag{6.19b}$$

At the start of each predictor and corrector step, DVODE is used to integrate the stiff source term vector of the pseudo-time step $\Delta\tau$. This is then included as a source term in the integration of the fluxes and non-stiff source term vector. The physical time derivative is computed using a two-step, backward difference formula.

6.4 Consistent Splitting in Dual-Time Schemes

This scheme follows the single-time, consistent splitting approach from Schwer et al. [64] and uses DVODE in both the predictor and corrector stages to integrate the stiff source terms and the residual terms. The stiff source terms are evaluated implicitly within DVODE while the residual terms are computed explicitly using the spatial discretization of the Generalized MacCormack scheme. The predictor and corrector steps each take two steps, with the first step integrating the system of equations and the second step solving the linear system associated with preconditioning:

$$\begin{aligned}
\frac{d}{d\tau} \mathbf{W}^* &= \mathbf{S}_S^{m+1/2} - \mathbf{R}^m - \frac{1}{2\Delta t} (3\mathbf{W}^m - 4\mathbf{W}^n + \mathbf{W}^{n-1}) \\
\left(\Gamma^m + \frac{3}{2} \frac{\Delta\tau}{\Delta t} \frac{\partial \mathbf{W}}{\partial \mathbf{Q}} \right) \Delta \mathbf{Q}^{m+1/2} &= \Delta\tau \frac{\mathbf{W}^*(\tau + \Delta\tau) - \mathbf{W}^*(\tau)}{\Delta\tau}
\end{aligned} \tag{6.20a}$$

$$\begin{aligned}
\frac{d}{d\tau} \mathbf{W}^{**} &= \mathbf{S}_S^{m+1} - \mathbf{R}^{m+1/2} - \frac{1}{2\Delta t} (3\mathbf{W}^{m+1/2} - 4\mathbf{W}^n + \mathbf{W}^{n-1}) \\
\left(\mathbf{\Gamma}^m + \frac{3}{4} \frac{\Delta\tau}{\Delta t} \frac{\partial \mathbf{W}}{\partial \mathbf{Q}} \right) \Delta \mathbf{Q}^{m+1} &= \frac{1}{2} \left[\mathbf{\Gamma}^m (\mathbf{Q}^m - \mathbf{Q}^{m+1/2}) \right. \\
&\quad \left. + \Delta\tau \frac{\mathbf{W}^{**}(\tau + \Delta\tau) - \mathbf{W}^{**}(\tau)}{\Delta\tau} \right]
\end{aligned} \tag{6.20b}$$

Using DVODE to integrate both the stiff source terms and the residual terms confines the splitting errors to the left-hand side of the equation [64]. This may have effects on convergence rates as the chemical and flow time scales become more disparate, but it contains no errors in the steady state solution that arise from the operator splitting [64]. The physical time derivative uses a two-step, backward difference formula.

6.5 Summary

In this chapter, the GMC-PC scheme is extended to treat stiff chemical source terms that arise during combustion. This is done through the coupling of an external, dedicated stiff solver to the underlying Generalized MacCormack scheme. In this case, the popular DVODE solver is used, but any ODE solver may replace it as desired. In particular, the physical-time split scheme could be extended to use the preconditioned chemical solver from Sankaran and Oefelein [51] and Hansen and Sutherland [53] as those approaches mature. All of the schemes presented are second-order accurate in pseudo-time and have second-order accurate errors due to the operator splitting when that is used. However, the physical-time split scheme drops the order of accuracy in the physical-time derivative due to the limitations of multi-step schemes and operator split techniques [136]. Despite this reduction in order of accuracy, the number of DVODE calls is drastically reduced for this scheme relative to the others.

For clarity in the remaining chapters, the schemes will be referred to using a short-hand name. The consistent-split scheme in Section 6.4 is designated *Scheme A*. The physical-time split scheme in Section 6.3.1 is designated *Scheme B*. These two schemes are new and their

properties are explored in-depth in Chapter 7. The traditionally, fully-coupled scheme is called *Scheme C* and is used as a baseline, along with the GMC-ST scheme using the no-time splitting of Equation (6.16), to contrast the performance with Scheme A and Scheme B. The implementation of Scheme C does not contain a majority of the *ad hoc* corrections required to make the scheme successful across a wide range of conditions [3, 52]. The scheme in Section 6.3.2 using a traditional operator splitting approach in pseudo-time is included to demonstrate the pitfalls of naïvely using operator splitting approaches from single-time schemes in the dual-time framework. The steady-state errors from the operator splitting will be shown to destroy the physical time accuracy of the scheme. This scheme is called *Scheme D*.

CHAPTER 7

REACTING SCHEME VALIDATION

Cases focusing on the autoignition of lean hydrogen mixtures under elevated pressures are selected to test the critical behavior of the schemes. This configuration is chosen because it is stiff, contains eigenvalues with large, positive real components [137], it is of practical interest in homogeneous charge compression ignition (HCCI) engines, and it is frequently studied using direct numerical simulations (DNS). For all cases, the chemical mechanism used is the 9-species, 21-step mechanism of Li et al. [138].

The stiffness of the chemical mechanism and the spectral behavior of the Jacobian matrix for it are presented using a zero-dimensional analysis from Cantera [85] for a constant-volume reactor. The ignition delay, transient behavior of the solution, and the final conditions are calculated as a baseline for comparison. These baseline results are used to verify the split schemes are capable of accurately integrating the source terms in the absence of spatial derivatives and fluxes. One-dimensional configurations including temperature and composition inhomogeneities show the behavior of the schemes when spatial changes in the flow interact with the source terms. The compositional inhomogeneity includes the addition of the H radical. This causes the reaction front to separate from the diffusion manifold and has led to the failure of Strang split schemes previously [135]. Lastly, a large-eddy simulation (LES) of a three-dimensional, turbulent constant-volume HCCI case is considered. This case is based on multiple DNS studies done in 2D [87, 139–141] and 3D [142]. Using the DNS solutions as a baseline, the performance and the accuracy of the split schemes is established for a physically complex flow.




Figure 7.1: Evolution of key species with time from the Cantera constant volume simulation

7.1 Zero-Dimensional Ignition

The autoignition of a lean hydrogen mixture with air at a pressure of 41 atm and a temperature of 1100 K provides the opportunity to investigate the mathematical nature of the chosen chemical mechanism. The equivalence ratio is $\phi = 0.1$ and the Cantera package is used to generate baseline results for comparison with the splitting schemes. At these conditions, the ignition delay predicted by Cantera is $\tau_0 = 1.83$ ms and the final temperature of the system is $T_f = 1478.79$ K. The temporal evolution of the major species and the radicals is shown in Figure 7.1.

The eigenvalues λ_i of the Jacobian \mathbf{D} are needed to determine the stiffness of the system as well as to understand the stability characteristics of the integration scheme. As discussed previously, eigenvalues with large, positive, real components are destabilizing. These eigenvalues are frequently called the *explosive* eigenvalues [137, 143, 144]. Figure 7.2 follows the presentation of the eigenvalues by Gupta et al. [137] and shows the absolute




Figure 7.2: Eigenvalue and temperature evolution with time from the Cantera constant volume simulation. The magnitude of the real component of the eigenvalues is on the right axis where open circles represent negative real components and filled squares represent positive real components. Temperature is the dashed line on the right axis.

value of the real component of the eigenvalues, with negative real components shown as open circles and positive eigenvalues shown as solid squares. The positive eigenvalues are present during the pre-heat phase and they merge together at the time of ignition, where the temperature gradient is the largest. Post ignition, the eigenvalues are all initially negative, with a small, positive eigenvalue appearing at a late time. The positive eigenvalues are associated with the production of the OH radical [137].

While the positive eigenvalues pre-ignition are numerically destabilizing, the ratio of the largest to smallest eigenvalues indicates the stiffness of the system. Both pre- and post-ignition, there are several orders of magnitude separating the eigenvalues. Therefore, this chemical mechanism is stiff for all time. The stiffness requires special treatment to ensure robustness and accuracy and justifies the use of the DVODE solver to integrate the source terms in the split schemes.

All four schemes are run at four different physical time step sizes to assess the accuracy in the zero-dimensional reactor. The largest chemical time scale, proportional to the inverse of the largest eigenvalue, is on the order of $0.001 \mu\text{s}$. This is taken as the baseline physical time step Δt_0 and the time step increases by a factor of ten in four different time step sizes to a maximum of $1 \mu\text{s}$. The pseudo-time convergence is halted when the L_2 norm of the pseudo-time derivative vector drops below 1×10^{-12} , or the L2 norm decays four decades from its initial value, or when the L2 norm does not change by more than 1×10^{-14} for more than 10 pseudo-iterations.

Scheme C required an excessive number of iterations to converge when run at time step sizes $\Delta t \in \{100\Delta t_0, 1000\Delta t_0\}$ and the simulations were not run to completion. The scheme does not contain any techniques to adapt the pseudo-time step based on the stiffness of the Jacobian, leading to frequent over-shooting of the change in the solution, resulting in oscillations about the steady state in pseudo-time. While techniques are available for this scheme [3, 52], it is ancillary to the objectives of understanding the split schemes and improving Scheme C is not pursued. For the remaining cases, the number of pseudo-time iterations required at each physical time step for each scheme and Δt are shown in Figure 7.3, along with the normalized temperature $\theta = (T - T_i)/(T_f - T_i)$ where T is the temperature, T_i is the mean initial temperature and T_f is the final temperature. Scheme B is not included in Figure 7.3 because it does not require any pseudo-time iterations in this case. The number of iterations is averaged over bins $10 \mu\text{s}$ wide. In each scheme, increasing the time step by a factor of ten led to an increase in the number of pseudo-iterations smaller than ten, indicating the efficiency of the dual-time schemes at larger time steps. All four schemes yielded solutions matching the baseline Cantera solution at the time steps used.

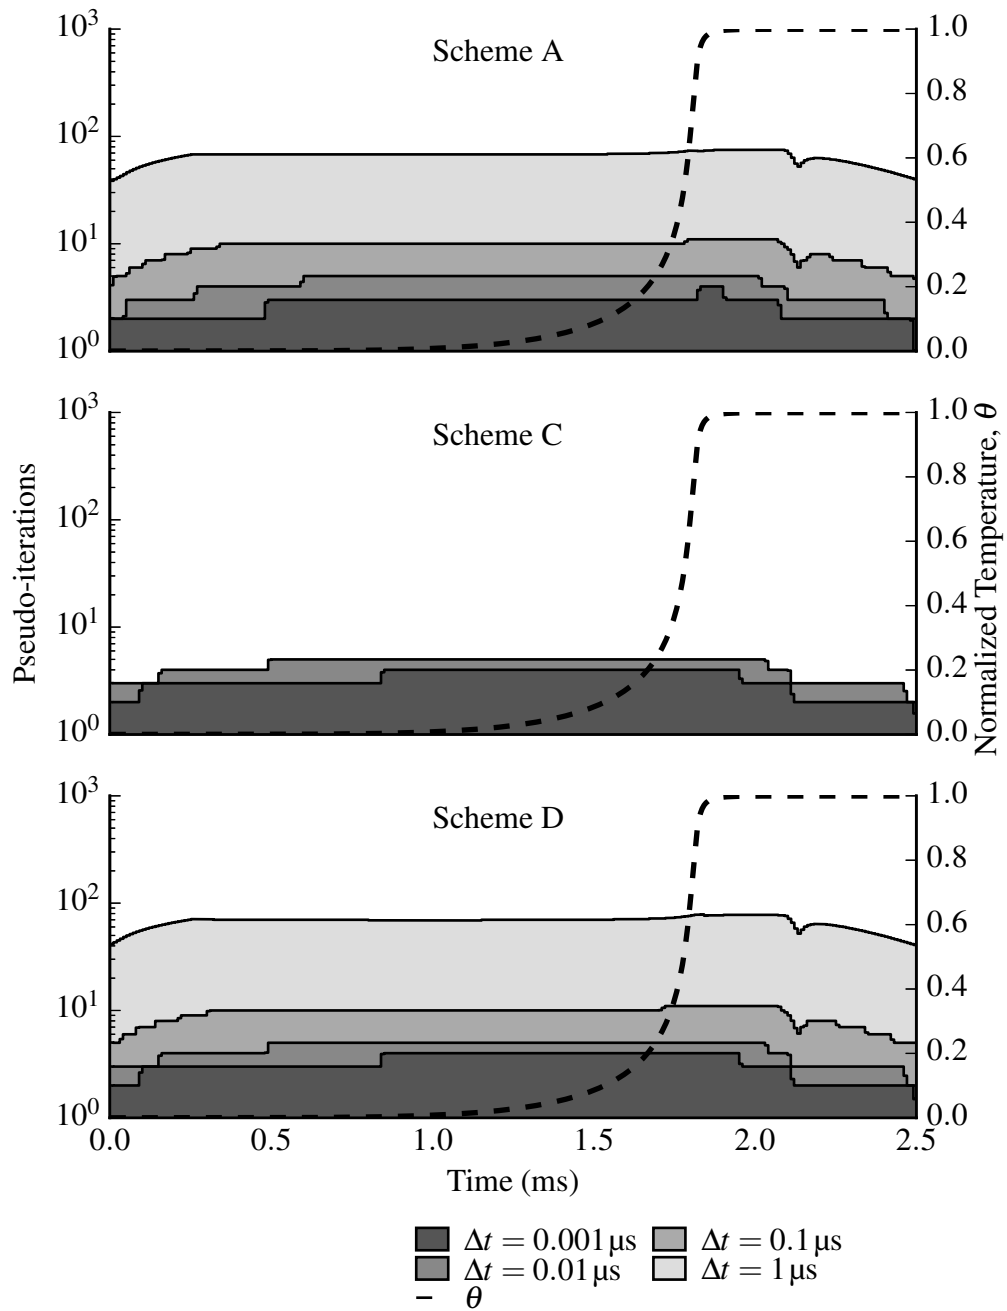



Figure 7.3: Number of pseudo-iterations per physical time step in the zero-dimensional reactor for each scheme, averaged over bins $10 \mu\text{s}$ wide. Scheme B does not require any pseudo-iterations for this configuration and is omitted. Scheme C failed to converge for the two largest physical time steps.


7.2 One-Dimensional Ignition

The decoupling of the spatial flux and the source terms leads to errors in Schemes A, B and D. Two simple tests demonstrate these errors using a one-dimensional, constant-volume, premixed hydrogen autoignition under similar conditions as the zero-dimensional case. Both cases use domains 1 mm in length and are discretized using 400 cells and periodic boundary conditions. The initial pressure is 41 atm and the mean temperature is 1095 K [145]. To ensure both fluxes and source terms are present, the flow is given a velocity of 1 m s^{-1} . This amounts to a reference frame transformation relative to the stationary case and does not alter the physics of the configuration. Case 1A introduces a sinusoidal temperature variation, as shown in Figure 7.4a, with amplitude 10 K and wavelength of 0.1 m, and the fuel is premixed hydrogen at an equivalence ratio of $\phi = 0.1$ [145]. Case 1B is a variation of Case 1A, where a compositional inhomogeneity is added. A small portion of the fuel is converted to atomic hydrogen, a total of 0.1% by mass. This is distributed through the domain using a sinusoidal function, shifted 90 degrees to the temperature inhomogeneity as illustrated in Figure 7.4b. The presence of atomic hydrogen has been shown to cause operator split schemes to fail in single-time schemes [135] and its effect on the split schemes in dual-time is assessed here. All four schemes are compared against the results using the GMC-ST scheme. The dual-time schemes are run at four different physical CFL numbers based on the speed of sound, $\text{CFL}_c \in \{1, 10, 100, 1000\}$. The convergence criteria are the same as those used for the zero-dimensional reactor.

Figure 7.5 compares the temporal evolution of the volume-averaged, normalized temperature for each scheme and CFL number for Case 1A. At $\text{CFL}_c = 1000$, both Scheme B and Scheme C become unstable and the simulations diverge. In the former case, there is a large, instantaneous release of heat during the source term step in the Strang splitting. This introduces significant volumetric expansion during the flux step of the Strang splitting, leading to large pressure waves that trigger numerical instabilities in the central scheme



(a) Initial temperature distribution for Case 1A.



(b) Initial temperature distribution and atomic H composition for Case 1B.

Figure 7.4: Initial fields for the one-dimensional test cases with a temperature inhomogeneity (Case 1A) and with both a temperature and compositional inhomogeneity through the introduction of atomic H (Case 1B).

Table 7.1: Errors measured by the L_2 norm of the volume-integrated temperature in Case 1A as a function of time for each scheme and CFL number, relative to the results for Scheme C at $CFL_c = 1$. The error for the explicit scheme relative to this case is 4.11×10^{-6} .

	$CFL_c = 1$	$CFL_c = 10$	$CFL_c = 100$	$CFL_c = 1000$
Scheme A	0.13×10^{-6}	0.71×10^{-6}	6.10×10^{-6}	184.68×10^{-6}
Scheme B	1.86×10^{-6}	3.71×10^{-6}	28.88×10^{-6}	—
Scheme C	—	1.57×10^{-6}	15.96×10^{-6}	—
Scheme D	0.28×10^{-6}	1.25×10^{-6}	15.70×10^{-6}	1478.35×10^{-6}

used. In the latter case, the divergence occurs during the point of peak heat release and the time step controls are not advanced enough to limit the time step at this phase. Except for Scheme D at $CFL_c = 1000$, the temperature evolution of all the remaining cases is indistinguishable with each other and with the baseline GMC-ST results. Table 7.1 reports the L_2 norm errors in the temperature profiles for this case, relative to the fully-coupled scheme at the smallest CFL number. Scheme A consistently has the lowest amount of error while Scheme B has errors that are of comparable order to the fully-coupled scheme. Scheme D at $CFL_c = 1000$ exhibits a delay in the ignition despite the convergence in pseudo-time, leading to a large value of the L_2 norm. The operator splitting is known to converge to the incorrect steady-state solution [65], causing operator splitting errors to manifest in physical-time. This demonstrates the need for consistent splitting as done in Scheme A.

More subtle differences between the behavior of the schemes for each CFL number can be seen in Figure 7.6 where the OH mass fraction as a function of space at multiple time instants is shown as a function of space. At early times, production of the OH radical occurs in the center of the domain where temperature is maximum. As time progresses, the amount of OH increases and by 1.9 ms, deflagration fronts begin to propagate outwards from the center. Scheme A shows minor variations in the profiles at the largest CFL number, despite

the temperature profiles all collapsing. These variations may be attributed to errors due to the increased numerical dissipation that occurs for large physical time step values. Variations in Scheme B begin to appear at a lower CFL number than in Scheme A. These errors in Scheme B are attributed to the splitting errors from the Strang splitting. Finally, the profiles for Scheme D show the formation of OH radical is significantly delayed at $CFL_c = 1000$ relative to the other CFL numbers and schemes. This delay is consistent with the delayed temporal variation in Figure 7.5 and again demonstrates the need for the consistent splitting approach in Scheme A to eliminate the steady-state errors due to operator splitting.

Case 1B is a more challenging case due to the importance of both the reactions and the radical mass diffusion. At early times, the atomic hydrogen participates in three reactions: recombination into H_2 , reaction with O_2 to form atomic oxygen and OH, and a three-body reaction with O_2 to form HO_2 [138]. The fundamental role of atomic hydrogen in the flame is the source of the errors when solved using an inconsistently split scheme. The large diffusion velocities and small mass of the atomic hydrogen lead to a rapid diffusion of the radical through the mixture where its role in the kinetics is to participate in the chain-branching reactions that drive the production of the OH radical. The formation of the OH radical is exothermic and begins driving the remainder of the reactions. Previous sensitivity studies conclude that the flame is most sensitive to variations in atomic hydrogen reaction rates and atomic hydrogen diffusion coefficients [146, 147].

This sensitivity to both reaction rates and diffusion coefficients suggests that decoupling the operators through a splitting scheme will generate large errors, a result borne out previously when investigating Strang splitting schemes in single-time schemes [135]. These errors may be compounded by additional sources of errors, such as numerical diffusion, that occur at large CFL numbers. Because of these effects, large variations in the solution exist between the schemes as shown in Figure 7.7a. For this case, Scheme C is unstable for all CFL numbers and is omitted. All three schemes match the explicit simulation at a physical $CFL_c = 1$ but begin to diverge from the baseline solution as the CFL number increases. In

the case of both Scheme B and Scheme D, increasing the CFL number delays the ignition while increasing the CFL number for Scheme A has the opposite effect.

Inspection of the OH mass fraction evolution suggests the source of errors is different between the consistent splitting in Scheme A and the inconsistent splitting in Scheme B and Scheme D. Figure 7.7b tracks the evolution of the volume-averaged OH mass fraction through the ignition process. The inset in Figure 7.7b shows the differences between the schemes appears very early in the evolution. In Scheme B and Scheme D, the early-time peak of OH mass fractions is delayed in time by the same amount of the increase in CFL number, whereas Scheme A predicts values close to the baseline, GMC-ST solution during this phase. In addition to the delay, the peak values are significantly higher than expected for these two schemes.

At late times, both schemes also show a delay in OH mass fraction production as the CFL number increases. On the other hand, Scheme A shows production occurring earlier for each CFL number as well as a reduction in the values of the peak OH mass fraction. As the CFL number increases, numerical dissipation increases and reduces the inhomogeneities in the mixture. This, in turn, leads to more rapid ignition.

From these two cases, we make several conclusions.

- The traditional approach to integrating chemistry in dual-time, Scheme C, exhibits robustness issues and instabilities when the time steps are large or the chemistry is very fast.
- Scheme D, which uses a classic operator splitting in pseudo-time, exhibits errors in the solution for Case 1A for the largest CFL number, but exhibits large departures from the solution for multiple CFL numbers in Case 1B. The errors that arise in Scheme D are due to splitting errors that persist at steady-state and these errors are more dominant in Case 1B by design.

- Scheme A is designed with a consistent splitting approach that eliminates these problems and its results are superior to Scheme D.
- The physical-time split scheme, Scheme B, exhibits robustness issues at large time steps and shows significant errors when the reaction and diffusion operators are decoupled as in Case 1B. Despite these drawbacks, it does result in a considerable cost reduction relative to Scheme A.

Based on these conclusions, Scheme C will not be considered further. When considering Scheme B, the number of pseudo-iterations per physical time step is similar for both Scheme A and Scheme B, but Scheme A uses the expensive DVODE call twice on each pseudo-iteration while Scheme B needs only two DVODE calls per physical time step. This can amount to approximately two orders of magnitude fewer DVODE calls in the Scheme B at moderate CFL numbers. Because of this savings, it is worth evaluating further.




Figure 7.5: Volume-averaged temperature profiles for each scheme and CFL number for Case 1A. Scheme B at $CFL_c = 1000$ is unstable due to the instantaneous volumetric expansion at the large, physical time step while Scheme C at the same CFL number is unstable at the point of maximum heat release. Both are omitted.

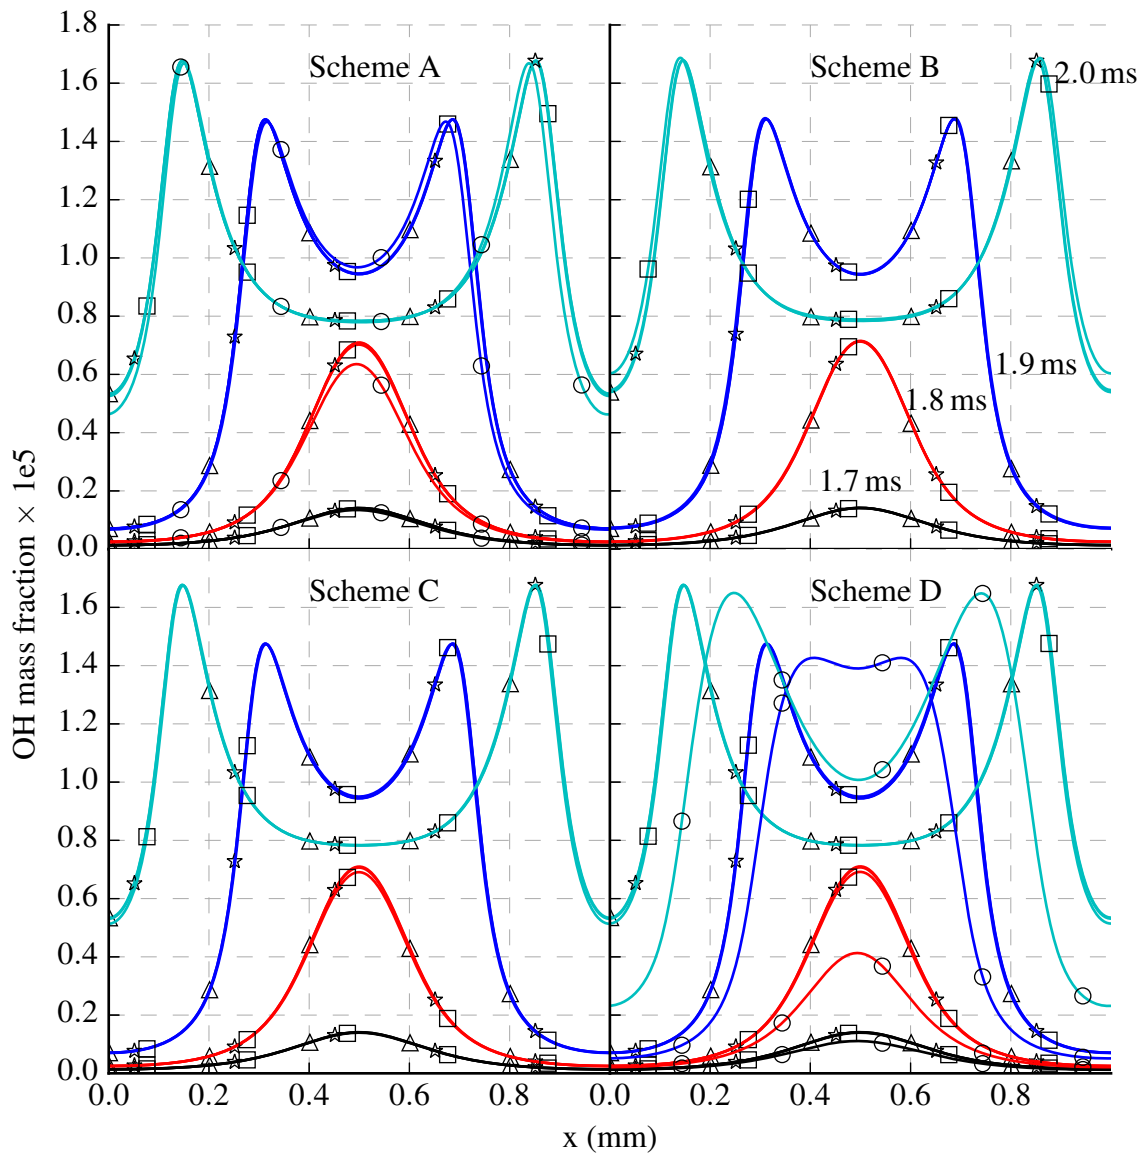
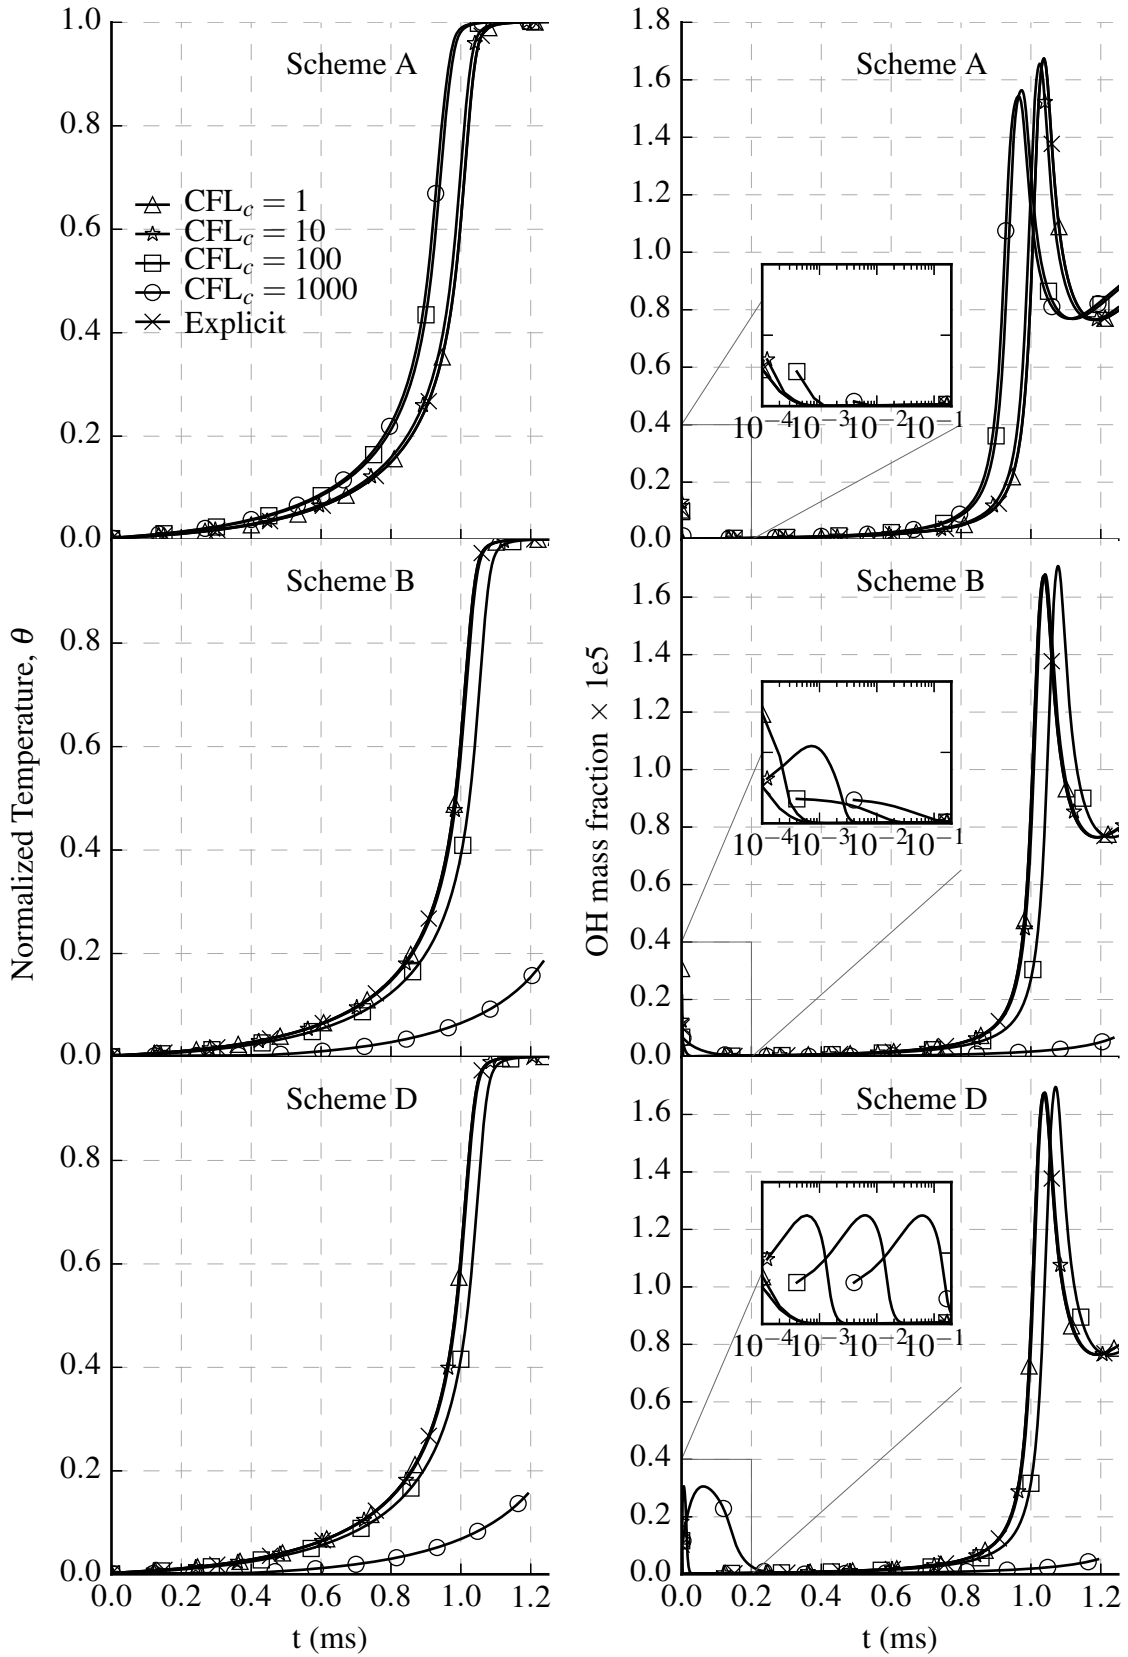


Figure 7.6: Instantaneous OH mass fraction at various instants in time through the ignition process for Case 1A. Symbols are the same as in Figure 7.5



(a) Volume-averaged temperature.

(b) Volume-averaged OH mass fraction.

Figure 7.7: Temperature and OH profiles for each scheme and CFL number for Case 1B. Scheme C is unstable for all physical CFL numbers tested and is omitted.

7.3 Three-Dimensional Ignition

As a final assessment of the schemes, we consider the autoignition of a homogeneous, premixed hydrogen mixture in a turbulent, constant volume system designed to mimic the conditions in an HCCI engine [87, 140–142]. The system is a periodic box with sides measuring 4.1 mm and contains a premixed hydrogen mixture with $\phi = 0.1$. The mean initial temperature is $\tilde{T} = 1070$ K at a pressure of $\bar{p} = 41$ atm. Superimposed on the mean fields are isotropic velocity and temperature perturbations with the RMS of the velocity fluctuations $u' = 0.5$ m s⁻¹ and the RMS of the temperature fluctuations $T' = 15$ K. The velocity and temperature fluctuations are generated independently and are therefore uncorrelated with each other. Both fluctuations are generated using the model spectrum [129]:

$$E(k) = 16 (\psi')^2 \sqrt{\frac{2}{\pi}} \left(\frac{k^4}{k_0^5} \right) \exp \left(-2 \frac{k^2}{k_0^2} \right) \quad (7.1)$$


where ψ is the variable whose spectrum is being computed, k is the wavenumber and k_0 is a user-specified parameter designating the peak wavenumber. These spectra are generated such that the most energetic length scale for the velocity fluctuations is 1 mm and for the temperature fluctuations is 1.32 mm [139]. All simulations start from identical realizations of the initial conditions to eliminate variations due to differences in random seeds. The domain is discretized with 64 cells in each direction, giving an initial resolution of $\Delta x/\eta = 4$ and $\Delta x/\delta_F^0 = 0.25$, where η is the Kolmogorov length scale and $\delta_F^0 = 0.257$ mm is the unstrained thermal thickness of the flame for the mean conditions [87]. Also of interest, the laminar flame speed is $S_L = 50$ cm s⁻¹ and the ignition delay time at the mean conditions is $\tau_0 = 3.3$ ms. With a $u'/S_L = 1$ and the laminar flame front resolved over four cells on the original mesh, the quasi-laminar chemistry assumption is used for the chemistry-turbulence interaction.

Investigations of non-reacting isotropic turbulence in Chapter 5 indicate the turbulent structures are independent of the physical time step selection with the preconditioned scheme.


With the introduction of chemical time scales and the coupling between the fluid mechanics and the chemistry, the dependence on physical time step size is evaluated here. Both Scheme A and Scheme B use physical time step sizes based on four CFL numbers. The turbulent Mach number based on the mean conditions and the RMS velocity is $M_t = 0.0012$ and the physical CFL numbers selected are $CFL_c = 1$, $CFL_c = 10$, $CFL_{u'} = 0.1$ ($CFL_c \approx 100$) and $CFL_{u'} = 1$ ($CFL_c \approx 1000$). These give physical time steps that are approximately 1×10^{-7} s, 1×10^{-6} s, 1×10^{-5} s and 1×10^{-4} s respectively, and are therefore 100 to 1×10^5 times larger than the chemical time scales.

As with the one dimensional simulations, Scheme B is numerically unstable when run at $CFL_{u'} = 1$ at the point of peak heat release rate and it is omitted from discussion. Figure 7.8 shows the surface area of the first propagating front (defined as the isosurface of $\tilde{Y}_{H_2} = 8.5 \times 10^{-4}$) [142] for the two schemes. At the large CFL numbers, $CFL_{u'} = 0.1$ and 1, the Scheme A results indicate a significant delay in the ignition compared to the results at $CFL_c = 1$ and 10. Additionally, the pressure rise is entirely absent at $CFL_{u'} = 1$, and a slight pressure rise is present at $CFL_{u'} = 0.1$. However, the results at $CFL_c = 1$ and 10 match the results predicted by Scheme B for all stable CFL numbers.

In both schemes, the ignition is delayed relative to the DNS data from Yu and Bai [142] and appears close to the ignition time of the homogeneous mixture. Additionally, at small CFL numbers with Scheme A and all stable CFL numbers with Scheme B, the flame area and steady state are insensitive to the choice of CFL number. Both schemes show the ignition front surface appearing later and reaching a lower peak value than the DNS data. These differences in ignition delay and flame surface area are due to the coarser grid resolutions used in this LES study. In particular, the coarse resolution enhances the dissipation of the temperature fluctuations at early times, prior to any appreciable chemical reactions. This delays the onset of ignition as the mixture is more homogeneous. Likewise, the coarser resolution will have less wrinkling and fine-scale curvature than is present in the DNS results, reducing the peak front area.



(a) Scheme A



(b) Scheme B

Figure 7.8: Front surface area (the isosurface of $\tilde{Y}_{H_2} = 8.5 \times 10^{-4}$) and pressure throughout the ignition process for the two schemes at each CFL number, compared to the 3D DNS data from Yu and Bai [142].

This loss of accuracy with the Scheme A at large CFL numbers is attributed to the behavior of the DVODE solver under these conditions. The DVODE solver is provided a set of ODEs in time where the right-hand side of the system represents the spatially discretized residual. This ODE system is integrated using an implicit, backward differencing formulation for stiff equations and the system Jacobian is computed numerically. In Scheme A, only the chemical source terms are treated implicitly in this nature and the spatial residual is held constant through the integration step. To properly treat the chemical source terms implicitly, the system of ODEs is augmented with an equation for the evolution of temperature. This temperature equation may be solved assuming constant enthalpy (constant pressure) or constant energy (constant volume). Both have been explored for this case and the loss of accuracy manifests for both choices, so this assumption is not the root cause.

To treat the stiffness in the system, DVODE computes the Jacobian numerically and uses the spectral information contained within it to detect stiffness and to determine time step sizes. In the case of Scheme A for $CFL_{u'} = 0.1$ and $CFL_{u'} = 1$, the loss of accuracy is attributed to the loss of numerical precision in the system when both the fluxes and source terms are added together. For these CFL numbers, the unsteady preconditioner is less active than at lower CFL numbers and the pseudo-time step sizes become large relative to the time step sizes when the unsteady preconditioner is fully activated. This, in turn, creates spatial residual terms that are large enough relative to the chemical source terms that double-precision floating point values do not contain enough resolution to represent the sum as different from the largest term. The loss of precision impacts the formation of radical species, including the atomic hydrogen radical, which is key to the production of OH and the release of heat to initiate the reactions. It also implies the numerical Jacobian is misrepresented within DVODE. This, in turn, causes DVODE to take larger time steps than should be used based on the stiffness of the system and the species production rates are predicted inaccurately. Species and reactions more prevalent in the flame may have production rates large enough relative to the fluxes to be captured when added together, so

some amount of heat is released and an equilibrium is eventually reached. This equilibrium is not the expected state, with H_2O over represented and all of the radicals under represented.

On the other hand, Scheme B at $\text{CFL}_{u'} = 0.1$ predicts the initial evolution accurately relative to the smaller CFL numbers until the point the primary ignition front disappears. The formation of the ignition kernels is predominantly attributed to the reaction rates, with convection and mixing a small component to their initial formation. Once the kernels grow and become deflagration fronts, the coupling between the reaction rates and the flow-field becomes more important to the evolution of the system. At $\text{CFL}_{u'} = 0.1$, the splitting errors due to the decoupling of the reactions from the flow become large enough to cause a difference in the system evolution. After this point, the rest of the fuel burns more slowly at this CFL number, taking approximately 20% longer to consume the remaining fuel relative to the other CFL numbers.

The mass fraction contours of HO_2 in Figure 7.9 at $t = 0.9\tau_0$ show minor differences as the CFL number increases for Scheme B and the ignition front indicated by the black lines (defined to be where $\widetilde{Y}_{\text{H}_2} = 8.5 \times 10^{-4}$) suggest the flow dynamics until this stage are independent of physical time step size, as observed in the non-reacting turbulence case [117]. Likewise, for $\text{CFL}_c = 1$ and 10, the results with Scheme A show the initiation and growth of the ignition kernels in the same locations as with Scheme B, with minor differences as the CFL number increases. The differences in the front contours can be attributed to differences in the front displacement speed. This displacement speed of the reaction front is defined as [87, 139, 140, 142]:

$$s_d = \frac{\frac{D\widetilde{Y}_{\text{H}_2}}{Dt}}{|\nabla\widetilde{Y}_{\text{H}_2}|} \quad (7.2)$$


evaluated on a constant mass fraction surface. The material derivative of the species mass fraction is:

$$\bar{\rho} \frac{D\widetilde{Y}_{\text{H}_2}}{Dt} = \nabla \cdot \left(\bar{\rho} D_{\text{H}_2} \nabla \widetilde{Y}_{\text{H}_2} \right) + \bar{\omega} \quad (7.3)$$

where D_{H_2} is the diffusion coefficient of hydrogen into the mixture. The displacement speed may be density weighted:


$$s_d^* = \left(\frac{\bar{\rho}}{\bar{\rho}_0} \right) s_d \quad (7.4)$$

Using this measure, spontaneous ignition fronts are defined as $s_d^* > 1.1S_L$ and the slowly moving deflagration fronts as $s_d^* < 1.1S_L$ [139]. The three-dimensional isosurfaces of the primary ignition front in Figure 7.10 show that both slow, deflagration fronts (colored blue) and rapid, spontaneous ignition fronts (colored yellow) exist in approximately the same locations for Scheme A and Scheme B, consistent with expectations [142]. This insensitivity to CFL number in the flow topology is similar to the preconditioned scheme behavior in non-reacting flows [117], although it is only true over a small range of CFL numbers in the reacting flows.




(a) Scheme A at $CFL_c = 1$

(b) Scheme B at $CFL_c = 1$



(c) Scheme A at $CFL_c = 10$

(d) Scheme B at $CFL_c = 10$



(e) Scheme A at $CFL_{u'} = 0.1$

(f) Scheme B at $CFL_{u'} = 0.1$





Figure 7.9: Contours of HO_2 for both schemes and each CFL number along the center plane at $t = 0.9\tau_0$ with the ignition front defined as the surface of $\tilde{Y}_{H_2} = 8.5 \times 10^{-4}$ indicated as black lines.



(a) Scheme A

(b) Scheme B

Figure 7.10: Isosurfaces of the ignition front defined as the surface of $\tilde{Y}_{\text{H}_2} = 8.5 \times 10^{-4}$ at $t = 0.9\tau$ for both schemes at $\text{CFL}_c = 1$, where the blue regions indicate the slow, deflagration front propagation ($s_d^* < 1.1S_L$) and the yellow regions indicate the rapid, spontaneous ignition front propagation ($s_d^* > 1.1S_L$).

7.4 Summary of results

This chapter investigates the efficacy of four different schemes to integrate chemical source terms in the autoignition of premixed hydrogen at elevated temperatures and pressures. The tests are chosen because they are stiff, the fluxes and source terms are highly coupled, and small discrepancies in accuracy manifest in large errors in the final solution. From these tests, we can conclude the following:

1. The traditional approach to chemistry integration in dual-time schemes, Scheme C, is unstable for very stiff problems such as Case 1B, where the presence of atomic hydrogen triggers rapid initial reactions. Even for simpler problems, the scheme exhibits reduced convergence rates. While *ad hoc* fixes are available for this scheme to work, such as slowly introducing the chemical source terms over the pseudo-iterations,

these techniques require user intervention and are not based in the physics of the problem.

2. Using an external solver, in this case DVODE, eliminates all of the user intervention and contains no *ad hoc* corrections to either the time step or the chemical source terms. When the solver is coupled using a consistent splitting technique as in Scheme A, the scheme is stable for very large physical CFL numbers in the zero- and one-dimensional cases, including Case 1B. However, the three-dimensional isotropic turbulence case revealed numerical precision errors and DVODE is unable to properly evaluate the Jacobian to determine proper step sizes.
3. The importance of properly splitting the operators is shown through the contrast with Scheme D, which employs a traditional operator splitting approach that is popular in single-time schemes. This scheme exhibits large errors as the physical CFL number increases. These errors are attributed to the convergence to an incorrect steady-state in the pseudo-time iteration that is an artifact of the traditional operator splitting techniques.
4. Finally, the scheme employing Strang splitting in physical time, Scheme B, performs surprisingly well in most cases. For moderate physical CFL numbers, the results in all cases are accurately predicted, including the turbulent autoignition simulations where the results were the same for time steps up to 1000 times larger than the chemical time scales. However, for large physical CFL numbers, the decoupling of the reactions and the convective-diffusive fluxes results in errors. While relatively minor in the turbulent autoignition simulation, the errors were severe in Case 1B, which is designed to trigger the deficiencies of Scheme B. Lastly, for the largest CFL numbers, robustness is often an issue as the decoupling of the source terms and the fluxes causes a large, local heat release that triggers pressure waves large enough to cause numerical instabilities.

Despite the drawbacks and limitations at large, physical CFL numbers, both the consistent split scheme and physical-time split scheme are superior to the traditional approach in terms of robustness and convergence rates. These schemes have the additional benefit of requiring no user intervention or *ad hoc* corrections without a basis in the physics of the problem. Furthermore, the techniques are not confined to using the DVODE solver and any solver designed for stiff chemical kinetics may be substituted, offering more flexibility than the traditional schemes. Because both Scheme A and Scheme B are limited to similar CFL numbers, the former due to numerical precision losses and the latter due to the coupling between sources and fluxes, only Scheme B is considered in the next chapter. This scheme also provides the benefit of orders of magnitude fewer DVODE calls and therefore offers the best potential improvement in performance.

CHAPTER 8


APPLICATIONS

The previous chapters provide the details and validation for the non-reacting and reacting forms of the Generalized MacCormack scheme with preconditioning. Based on the successes and limitations of these schemes, this chapter explores two turbulent, chemically reacting configurations. The first is the Volvo flygmotor, a bluff-body stabilized, premixed combustor that serves as a canonical test case for turbulent combustion. This configuration is the target of numerous experimental and computational studies [68–78]. The results in this chapter are a comparison between an existing set of results using an explicit scheme presented by Sankaran and Gallagher [80] and the GMC-PC scheme using the physical time split chemistry integration. The GMC-PC scheme is performed only on the coarsest grid used in Sankaran and Gallagher [80] and differences are apparent in the structure of the flow behind the bluff body.


The second configuration is the Cambridge slot burner that has both premixed and stratified configurations. Experimental data is available for comparison [81, 82, 84], along with the only other known computational study, performed using a two-dimensional RANS approach [83]. The GMC-PC scheme is run with the physical time split chemistry integration using a reduced, stiff chemical mechanism. In addition to the comparisons with experimental data, the topological structure of the flow throughout the flame is presented to identify key differences between the premixed and stratified flows.

8.1 Volvo Flygmotor

The computational configuration for the Volvo flygmotor matches the requirements of the Model Validation for Propulsion workshop [79] and the domain is detailed in Figure 8.1. The domain represents a small section down the center of the experimental rig, with periodic



(a) Cross-section of the Volvo flymotor computational domain.



(b) Side view of the Volvo flymotor computational domain.

Figure 8.1: Schematic of the computational domain used for the premixed, bluff-body stabilized flame simulations using the specifications of the Model Validation for Propulsion workshop [79].

boundaries in the spanwise direction. The inflow is a constant-mass, partially-reflective characteristic boundary with the massflow rate set to 0.2083 kg s^{-1} at $\tilde{T} = 288 \text{ K}$ and an equivalence ratio of $\phi = 0.65$. The fuel is propane and the oxidizer is air. The outflow is a constant pressure outflow with a back pressure of 100 kPa and a convective sponge boundary is applied to eliminate acoustic waves. All remaining boundaries are adiabatic, no-slip walls. A two-step, global mechanism is used [148] along with Sutherland's law for the transport properties and thermally-perfect assumptions for the equation of state. The DTF model is used for turbulence-chemistry interaction.




Figure 8.2: Grid clustering around the bluff body and shear layers.

The baseline results for comparison use the hybrid MacCormack-MUSCL scheme [149], where the MUSCL scheme provides numerical stabilization through the flame front. These results were presented previously as part of the MVP workshop [80]. Grid independence was assessed using three, progressively finer grids designated *coarse*, *medium*, and *fine*. The grids contain 0.6 million, 1.2 million, and 3.7 million cells respectively. In all cases, the grid stretching is confined to under 10% in order to ensure scheme accuracy. These baseline results are relatively insensitive to the grid for the dimensions chosen [80]. Therefore, only the coarse grid is used to with the GMC-PC scheme and Figure 8.2 shows the grid in the region of the bluff body for this setup. The GMC-PC scheme is run at a physical time step of 1×10^{-6} s, approximately 100 times larger than the explicit time step. Stability through the gradients in the flame front is provided by JST artificial dissipation. For this case, the convergence criteria is set to 4 decades residual decay, requiring approximately 50 pseudo-iterations per physical time step on average. This is the minimum number of decades that ensured stability in the physical time solution over long simulation times.

Figure 8.3 and Figure 8.4 show the instantaneous temperature and spanwise vorticity contours along the center plane for the GMC-PC scheme and the explicit scheme. The

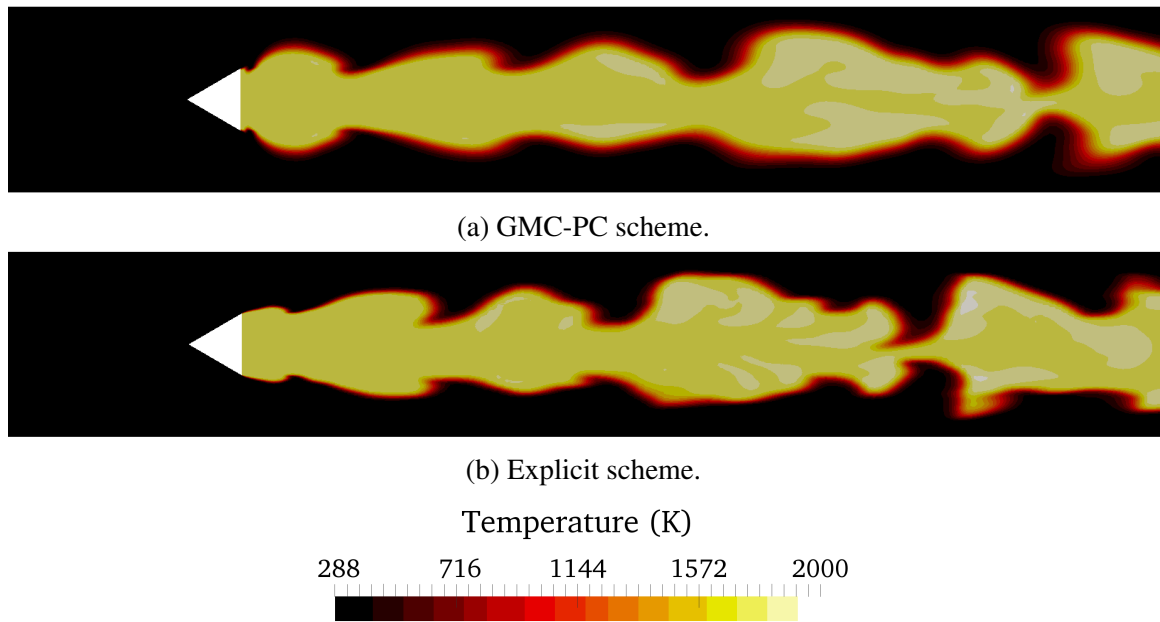


Figure 8.3: Temperature contours along the centerplane of the Volvo flygmotor using the GMC-PC scheme and the explicit scheme. Images are not at the same time instant.

time instants are different between the two schemes, however it is clear that both schemes predict a symmetric flame shape, which is an important metric for success in this case. The GMC-PC scheme shows less wrinkling along the edges of the flame in Figure 8.3, which is attributed to the JST dissipation providing more numerical diffusion than the MUSCL scheme in the explicit results. Likewise, the vortical structures also appear more diffuse with the GMC-PC scheme in Figure 8.4. However, both schemes capture the change in sign of the vorticity downstream of the bluff-body, which is again another key metric for measuring the success of the scheme. The flow field immediately behind the bluff body is different between the GMC-PC and explicit schemes. The latter shows little vorticity until approximately one bluff-body height downstream, whereas the former shows vortical structures much closer to the bluff body.

This difference in vortical structures near the body is more clear in the time-averaged axial velocity along the centerline in Figure 8.5a. The GMC-PC scheme shows a more pronounced recirculation region, coming closer to the experimental data, than the explicit

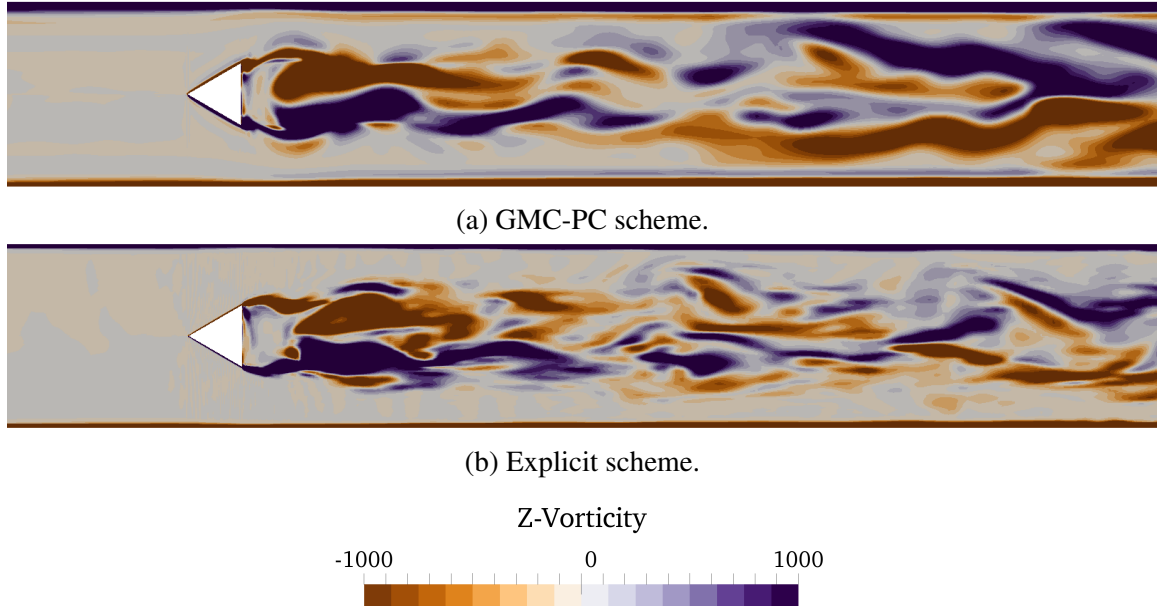


Figure 8.4: Spanwise vorticity contours along the centerplane of the Volvo flymotor using the GMC-PC scheme and the explicit scheme. Images are not at the same time instant.

scheme on any of the grids tested. With most of the flow Mach number behind the bluff body over $M = 0.1$, the region behind the bluff-body is the only location where a low Mach number flow exists and it is not surprising the GMC-PC scheme exhibits the largest differences in this location. The mean fluctuation level, defined as $\sqrt{u'^2 + v'^2}/U_{\text{inlet}}$, in Figure 8.5b shows the GMC-PC scheme predicts increased levels of fluctuation along the entire centerline, but most noticeably near the bluff body. Although all simulations were time-averaged for similar lengths of time, the fluctuation level in the GMC-PC case shows signs of not being fully converged. The lower numerical dissipation levels in the GMC-PC scheme leads to longer averaging times required for second order statistics to converge due to increased level of variation in the solution, and so additional averaging time may be needed. The centerline anisotropy levels in Figure 8.5c, defined as v'/u' , shows all schemes and grids produce similar levels of anisotropy with axial velocity fluctuations dominating the transverse velocity fluctuations in contrast with the experimental data.

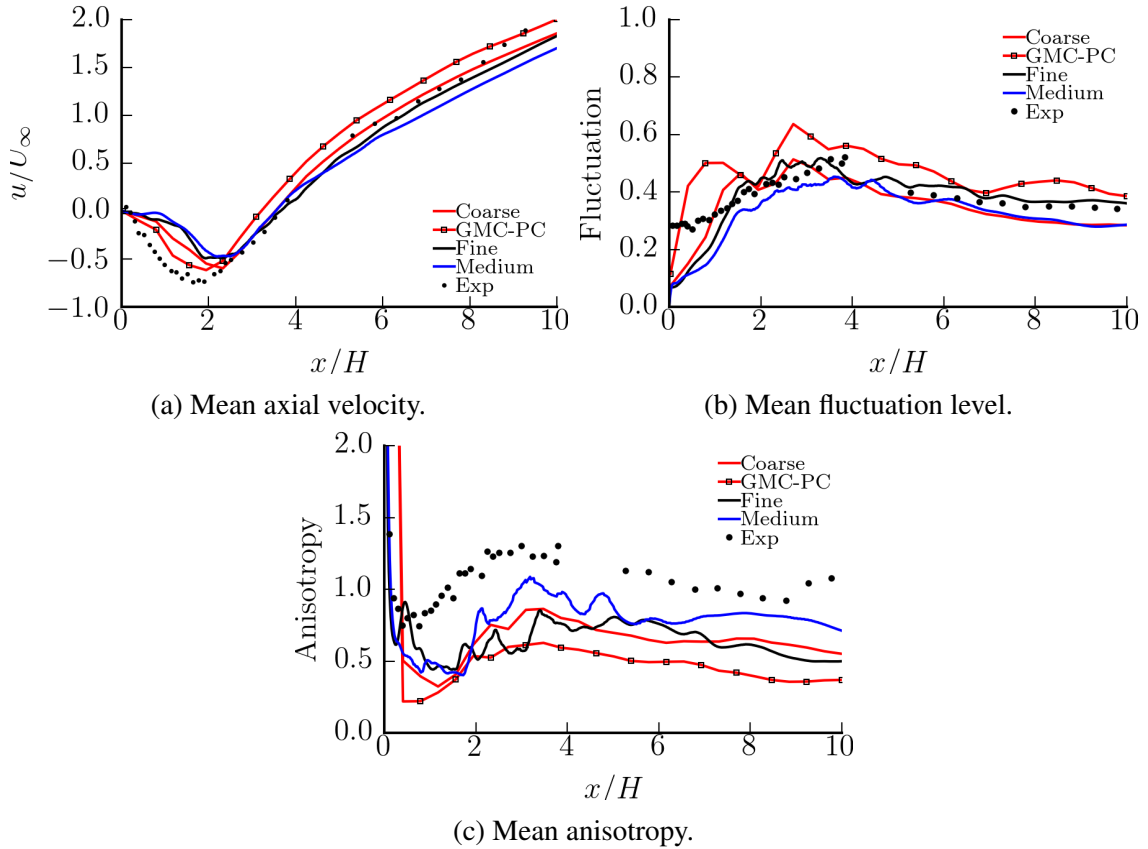


Figure 8.5: Profiles along the centerline behind the bluff-body for the GMC-PC scheme compared to the explicit scheme. Results for the explicit scheme on several consecutively finer meshes provide context for the GMC-PC scheme’s results.

Transverse velocity profiles at multiple locations downstream of the bluff body in Figure 8.6 show similar trends as the centerline values. The mean velocity values for all schemes and grids match each other and the experimental data well in Figure 8.6a and Figure 8.6b. The RMS velocities in Figure 8.6c and Figure 8.6d show more variation, particularly relative to the GMC-PC scheme. At all locations, the GMC-PC scheme again shows signs of requiring additional convergence. The GMC-PC scheme captures the peaks in the RMS velocities better, particularly relative to the explicit scheme on the coarse grid, but the values near the centerline are generally overpredicted.

The simulations with the GMC-PC scheme require approximately 10 times as much wall time as the explicit scheme on the same grid to reach the same physical time. The

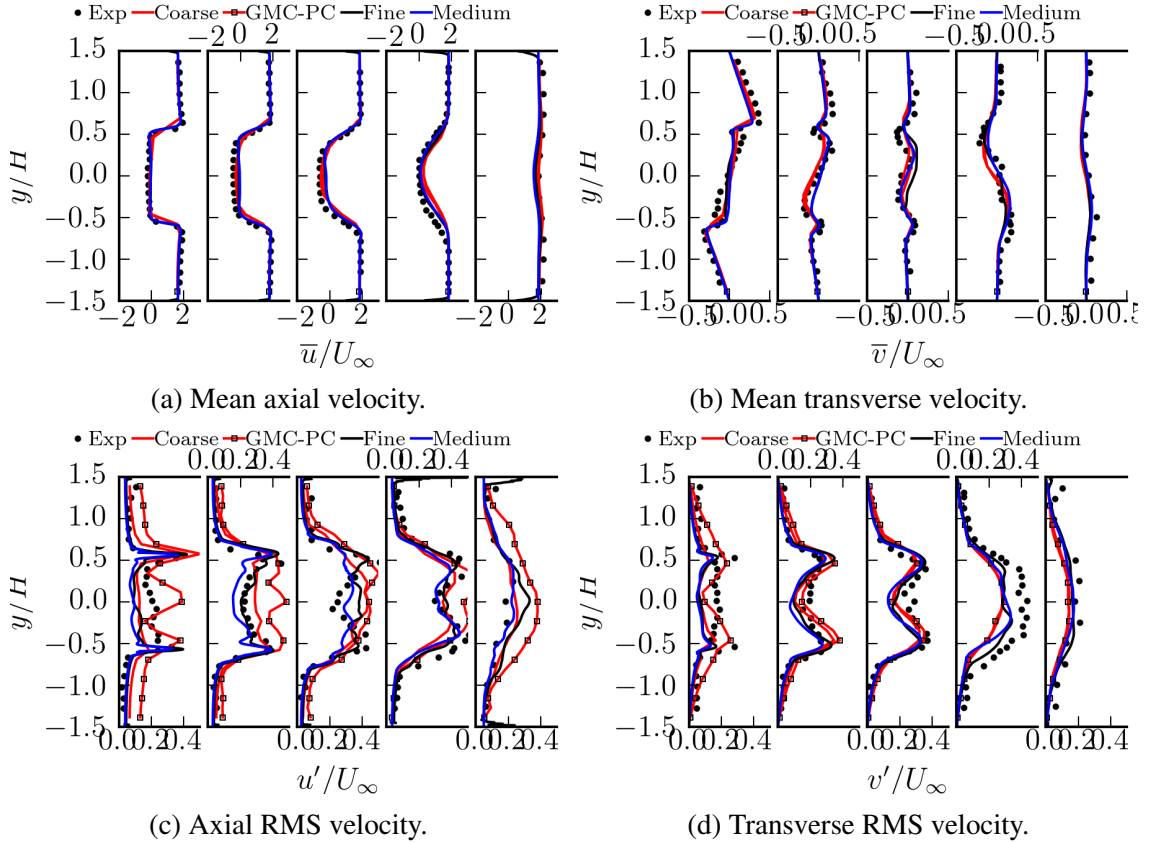


Figure 8.6: Transverse profiles at the downstream locations $x/H \in \{0.375, 0.95, 1.53, 3.75, 9.4\}$ for the GMC-PC scheme, the explicit scheme, and experimental measurements. Results for the explicit scheme are on several consecutively finer meshes.

explicit scheme integrates the chemical source terms explicitly; had the explicit runs used DVODE for chemical integration, the GMC-PC scheme would be between 1 and 1.5 times more expensive than the explicit scheme. This is because the time step is approximately 100 times larger than the explicit scheme, each physical iteration requires on average 50 pseudo-iterations, and each pseudo-iteration is between two and three factors more time than the explicit scheme iterations when using DVODE. However, the solution quality is similar throughout the majority of the flow field using the GMC-PC scheme, with the largest changes occurring in the region behind the bluff body. This region is the only portion of the flow that is at a low Mach number, and the GMC-PC scheme predicts the proper recirculation region whereas the explicit scheme on all grids predicts a wake-like structure. As suggested

by the studies in previous chapters, the explicit scheme would require a much finer grid in the low Mach number region to achieve similar results. In this sense, the GMC-PC scheme is more efficient and provides higher quality solutions on a given grid.

8.2 Cambridge Slot Burner

The Cambridge slot burner is a series of experiments operating at a range of conditions from perfectly premixed through a range of stratification ratios. The experimental burner consists of two pairs of slots, one lean pair and one rich pair, surrounded by single slots of coflowing air as shown in Figure 8.7. Each slot is 5 mm wide and 50 mm deep with walls between each slot. A rod with a diameter of 1.5 mm located 10 mm above the slot exit and 2 mm towards the rich side serves as the anchor for a V-flame. The slots have a square mesh turbulence generator at the exit.

Two configurations are tested here. The premixed configuration corresponds to the *fsh1* case [81–84] with volumetric flow rates of methane through both pairs of slots at 8.82 L/min and air flow rates of 116 L min⁻¹. The equivalence ratio for both slots is $\phi = 0.73$. The stratified configuration corresponds to the highest stratification ratio case available, *fsh6* [81–84], where the volumetric flow rate of methane through the rich slots is 13.2 L min⁻¹ and 4.41 L min⁻¹ through the lean slots. The coflowing air remains unchanged. The lean equivalence ratio is $\phi = 0.37$ and the rich equivalence ratio is $\phi = 1.10$.

The computational domain starts at the exit of the slots and extends upwards by 150 mm. To keep the number of grid points reasonable, the walls between the slots are not resolved and there are discrete jumps in the inflow boundary conditions between the slots. The domain is 132 mm wide, with the extra space on the sides of the air slots also providing coflowing air to avoid entrainment issues and recirculation regions outside of the air slots [83]. Only the central 20 mm of the depth direction is simulated with periodic boundary conditions in that direction. The left and right boundaries are slip walls, and the top boundary is a partially reflecting characteristic outflow imposing atmospheric pressure. The inflow

is a partially reflecting, characteristic inflow with the coflowing air at 2.88 m s^{-1} and the methane slots at 4.386 m s^{-1} . Isotropic turbulence is imposed on the inflow velocities with a turbulence intensity of 10% [83].

The grid contains approximately 2.6 million cells. A large region around the rod is discretized with approximately uniform cells in all directions, with gentle stretching from the inflow to the uniform region, as shown in Figure 8.8a. Outside this region, the grid is stretched with under a 10% stretching ratio towards the edge walls and the outflow to reduce the number of grid points. The rod is resolved with enough cells sufficient to place two points within the viscous sub-layer and stretched at less than 10% to the uniform region as shown in Figure 8.8b.

The simulation is run using the GMC-PC scheme with the physical time split chemical integration and the DTF model. The grid is fine enough in the region of interest that the DTF model only thickens the flame by under 1% to ensure it is resolved over 5 cells. JST artificial dissipation is used above $y = 32 \text{ mm}$, where the grid progressively becomes too coarse to resolve the flame and remain stable. However, no additional artificial dissipation is required within the region of interest. The time step using the explicit scheme averages $1 \times 10^{-9} \text{ s}$ due to the small cells at the rod surface. The GMC-PC scheme is run at a time step of $5 \times 10^{-7} \text{ s}$, which corresponds to a physical CFL number of approximately 10 away from the rod surface where the grid is uniform. This time step is chosen based on the limits based on solution accuracy observed during the validation studies. A time step of $5 \times 10^{-8} \text{ s}$ was also run for a brief period and showed little change in the instantaneous structures, therefore extended time-averaging using this time step was not performed. The pseudo-time is converged such that the residual decays at least 2 decades, and for both premixed and stratified configurations this requires an average of 100 pseudo-iterations. The chemistry is modeled using the 16 species, 25 step Smooke-Giovangigli reduced mechanism [86]. The transport properties are evaluated using constant, non-unity Lewis numbers and the state

Table 8.1: Species and Lewis numbers used with the Smooke-Giovangigli mechanism.

Species	CH ₄	O ₂	N ₂	CH ₂ O	H ₂ O	HO ₂	H ₂ O ₂	HCO
Lewis	0.97	1.11	1.0	1.28	0.83	1.1	1.12	1.27
	H ₂	H	O	OH	CO ₂	CO	CH ₃	CH ₃ O
	0.3	0.18	0.7	0.73	1.39	0.3	1.0	1.3

equation uses the thermally perfect assumptions. The species and the corresponding Lewis numbers are listed in Table 8.1.

The mean temperature profiles at two locations, 20 mm and 30 mm, above the slot exit in Figure 8.9 show the results with the GMC-PC scheme compared to the experimental data and RANS data [83]. Error measurements for the experimental data are not provided, so agreement may only be determined qualitatively. In the premixed configuration shown in Figure 8.9a, the flame predicted by the GMC-PC scheme is slightly wider than the experimental data, whereas the RANS is slightly narrower. Further downstream, shown in Figure 8.9b, the width of the flame is over-predicted by both the GMC-PC scheme and the RANS results. At 20 mm downstream from the burner in the stratified configuration, the experimental data indicates a flame of similar width to the premixed case, but the temperature peak is shifted to the right of the centerline of the rod as shown in Figure 8.9c. Both the RANS data [83] and the GMC-PC data indicate a flame shifted to the left of the centerline, with the GMC-PC scheme producing a wider flame than the RANS results. The peak temperature from the GMC-PC scheme matches the peak temperature from the experiments well, while the RANS data over-predicts the peak temperature. This distinction is likely due to differences in the chemical modeling. The RANS data uses a flamelet approach and the product temperature in the premixed case is accurate but in the stratified case is over-predicted. With the reduced kinetic mechanism used here, the final product temperature

is predicted better. Downstream, the GMC-PC scheme gives a flame position consistent with the experimental data and superior to the RANS data as shown in Figure 8.9d.

Figure 8.10 shows the vorticity contours on the center plane for both the premixed and stratified configurations. Immediately evident is the transition from a symmetric flow in the premixed case to an asymmetric flow in the stratified case, where the stratified flow deflects further into the lean side and deflects slightly further towards the coflowing air on the other side. In the near-field behind the rod, the flow fields are different; the premixed case indicates a separated flow similar to a laminar flow over a cylinder, while the stratified flow is more indicative of a turbulent flow that remains attached longer. More pronounced, however, is the vorticity field further downstream. The premixed case maintains the sign of the vorticity away from the rod. On the other hand, the stratified case indicates a change in sign of the streamwise vorticity downstream of the rod. This change in sign is similar to what is observed in the Volvo case. Additionally, the premixed vorticity contours show a relatively large vortex sheet on the left side of the V-flame while the stratified case indicates smaller vortices within a sheet. These observations are more evident in the species mass fractions.

Two species are used to indicate the leading and trailing edge of the flame brush. The leading edge is tracked with the CH_2O mass fraction, while the trailing edge is tracked with the OH mass fraction. In addition to the flame asymmetry, stratification increases the amount of wrinkling on the flame front [83]. Inspection of the CH_2O mass fraction on the leading edge of the flame shows a larger number of smaller spatial features in Figure 8.11c relative to the larger structures in Figure 8.11a. The right side of the flame in the premixed case indicated a symmetric shedding pattern, while the stratified case shows asymmetry in the size and location of the vortices. Consistent with these leading edge vortices, the OH contours on the inner edge of the flame indicate the amount of wrinkling that persists through the flame front. For the premixed case in Figure 8.11b, the contours indicate minor wrinkling caused by regular, large-scale structures. In the stratified case in Figure 8.11d,

the left edge of the contours suggest a number of small scale structures that persist through the front. The right side of the stratified flame shows the impact of the relatively larger structures. However, as indicated by the OH contours, the flame thickness is considerably larger in the stratified case and the wrinkling induced by the stratification does not persist deep into the flame.

The types of structures in the flame may be determined through a quadrant analysis of the invariants of the velocity gradient tensor. This analysis has been performed on incompressible wall-bounded flows and jets [150–153], compressible decaying isotropic turbulence [154], compressible boundary layers [155], and premixed-flame turbulence interaction [156] as well as a premixed turbulent jet flame [157]. To start, the velocity gradient tensor is $\mathbf{A} = \frac{\partial u_i}{\partial x_j}$ at each point in space and the tensor invariants are defined as the coefficients in the characteristic equation [156]:

$$\lambda^3 + P\lambda^2 + Q\lambda + R = 0 \quad (8.1)$$

The invariants are:

$$\begin{aligned} P &= -\text{Tr}(\mathbf{A}) = -A_{ii} \\ Q &= \frac{1}{2} (\text{Tr}(\mathbf{A})^2 - \text{Tr}(\mathbf{A}^2)) = \frac{1}{2} A_{ij}A_{ji} \\ R &= -\det(\mathbf{A}) = \frac{1}{3} A_{ij}A_{jk}A_{ki} \end{aligned} \quad (8.2)$$

The invariant P is the volumetric expansion of the fluid, where $P < 0$ indicates expansion, $P = 0$ is dilation free, and $P > 0$ indicates compression. Non-dimensionalized versions of all three invariants may be used:

$$\begin{aligned} \hat{P} &= \frac{P}{\|\mathbf{A}\|} \\ \hat{Q} &= \frac{Q}{\|\mathbf{A}\|^2} \\ \hat{R} &= \frac{R}{\|\mathbf{A}\|^3} \end{aligned} \quad (8.3)$$

where $\|\mathbf{A}\| = (A_{ij}A_{ij})^{1/2}$.

The three-dimensional $(\hat{P}, \hat{Q}, \hat{R})$ space can be divided into sectors divided by surfaces where the discriminant \hat{D} is zero [156]:

$$\hat{D} = \frac{1}{108} \left(27\hat{R}^2 + \left(4\hat{P}^3 - 18\hat{P}\hat{Q} \right) \hat{R} + 4\hat{Q}^3 - \hat{P}^2\hat{Q}^2 \right) \quad (8.4)$$

Two surfaces divide the space, above which the points are part of focal topologies and below which the points are part of nodal topologies:

$$\begin{aligned} \frac{1}{3}\hat{P} \left(\hat{Q} - \frac{2}{9}\hat{P}^2 \right) - \frac{2}{27} \left(-3\hat{Q} + \hat{P}^2 \right)^{3/2} - \hat{R} &= 0 \\ \frac{1}{3}\hat{P} \left(\hat{Q} - \frac{2}{9}\hat{P}^2 \right) + \frac{2}{27} \left(-3\hat{Q} + \hat{P}^2 \right)^{3/2} - \hat{R} &= 0 \end{aligned} \quad (8.5)$$

A third surface also exists where the eigenvalues of \mathbf{A} are purely imaginary:

$$\hat{P}\hat{Q} - \hat{R} = 0 \quad (8.6)$$

These dividing surfaces break the space into eight sectors as shown in Figure 8.12. Each sector is defined in Table 8.2 based on the critical points associated with that topology. With the velocity gradient tensor, the sectors also have a physical interpretation relating to the type of flow that occurs at each point [153]. These physical interpretations are also included in Table 8.2.

To identify the topology throughout the flame surface, three different isosurfaces are extracted from the dataset. These surfaces represent regions of constant temperature, $\tilde{T} = 800$ K, 1300 K, and 1800 K. This corresponds to roughly the leading edge, center, and trailing edge of the flame brush for both the premixed and stratified configurations. Because each value of \hat{P} represents a (\hat{Q}, \hat{R}) plane, the values of \hat{P} are grouped into four bins so they may be projected onto a plane. The resulting correlations are in Figure 8.13 and Figure 8.14 for the premixed and stratified cases, respectively.

Table 8.2: Sectors of the $(\hat{P}, \hat{Q}, \hat{R})$ space divided by the discriminant surfaces along with their topological characteristics [156] and physical interpretations.

Sector	Topology	Interpretation
S1	Unstable focus/compressing	Vortex Compression
S2	Unstable node/saddle/saddle	Vortex Sheets
S3	Stable node/saddle/saddle	Vortex Tubes
S4	Stable focus/stretching	Vortex Stretching
S5	Stable focus/compressing	Vortex Compression
S6	Stable node/node/node	
S7	Unstable focus/stretching	Vortex Stretching
S8	Unstable node/node/node	

The top row of Figure 8.13 and Figure 8.14 are points that experience the largest amounts of volumetric expansion. In the premixed case, the points are predominantly in the S3 region indicating the presence of vortex tubes on the surfaces. Moving from left to right, the density of points decreases as the temperature increases. For the premixed flame, the high temperature isosurface is in the product side of the flame where less volumetric expansion occurs and the majority of the points move to lower values of \hat{P} . The stratified case also shows a predominance of S3 structures (vortex tubes) through all three isosurfaces, with the tubes persisting through towards the product side. In contrast with the premixed case, there are also a significant number of points in the S1 sector indicating vortex compression that is not present in the premixed case. By the $\tilde{T} = 1800$ K surface, the vortex compression is no longer present.

At lower levels of volumetric expansion in the second row of figures, S3 structures are again the dominant features in both the premixed and stratified cases. The premixed case also shows a large number of points in the S4 and S5 region indicating both vortex stretching and vortex compression occur along the surfaces, with it being more frequent

at the lower temperatures The stratified case shows less stretching and more compression, persisting through all three isosurfaces. The third row corresponds to the dilatation free regime and both configurations show a majority of the points lie in the vortex stretching and compression regions on the lower temperature isosurfaces. At the high temperature surface, the premixed case shows both S2 and S3 regions are populated, indicating the presence of stable vortex tubes and unstable vortex sheets. On the other hand, the stratified case shows points primarily in the S3 region indicating stable vortex tubes are more frequent than unstable vortex sheets.

The structures identified from the tensor analysis are consistent with the visual observations in Figure 8.10 and Figure 8.11. Both configurations have a large number of vortex tubes throughout the reaction front, with the stratified case dominated by tubes while the premixed case also contains regions of vortex sheets. In both cases, vortex compression is common, but the stratified case shows an increased number of regions undergoing stretching. The stratified case indicates the vorticity changes sign downstream of the rod suggesting a strong baroclinic torque component exists. This baroclinic torque may be the cause of the vortex stretching observed in the invariants and suggests a fundamentally different flow topology exists in the presence of stratification. For both the premixed and stratified cases, neither shows the classic teardrop shape in the invariants that is associated with non-reacting turbulent flows, a trend that has been observed in other premixed reacting systems [156, 157].






Figure 8.7: Schematic of the Cambridge slot burner.



(a) Grid in the measurement window of the experiment.



(b) Grid around the rod.

Figure 8.8: Grid within the important regions of the domain, including the near-wall location where the flame will anchor.

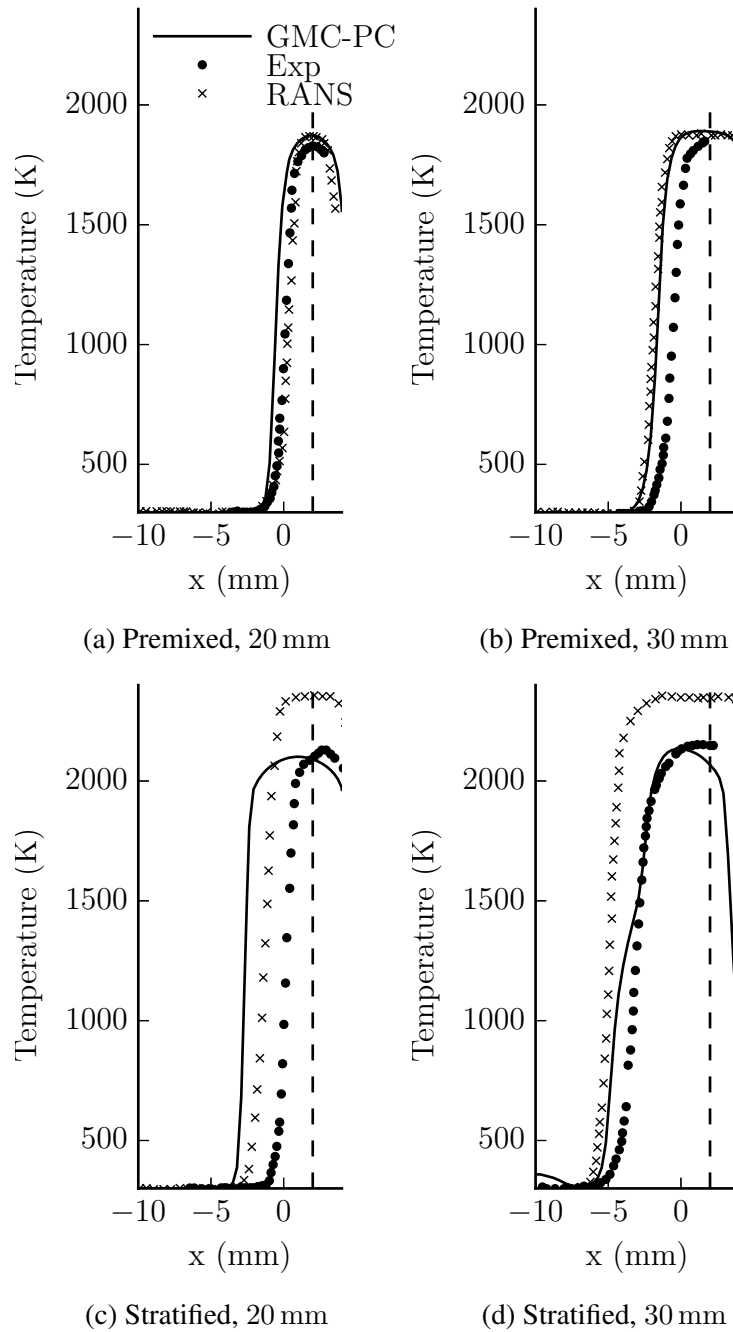


Figure 8.9: Mean temperature profiles at two streamwise locations downstream from the burner exit for both the premixed and stratified conditions. Experimental and RANS data digitized from Darbyshire et al. [83]. The vertical dashed line marks the centerline of the rod.



(a) Premixed



(b) Stratified

Z-Vorticity




Figure 8.10: Instantaneous spanwise vorticity along the center plane of the burner for both the premixed and stratified configurations.




Figure 8.11: Instantaneous contours of reactant-side and product-side radical species indicative of the leading and trailing edges of the flame brush for both configurations along the center plane.

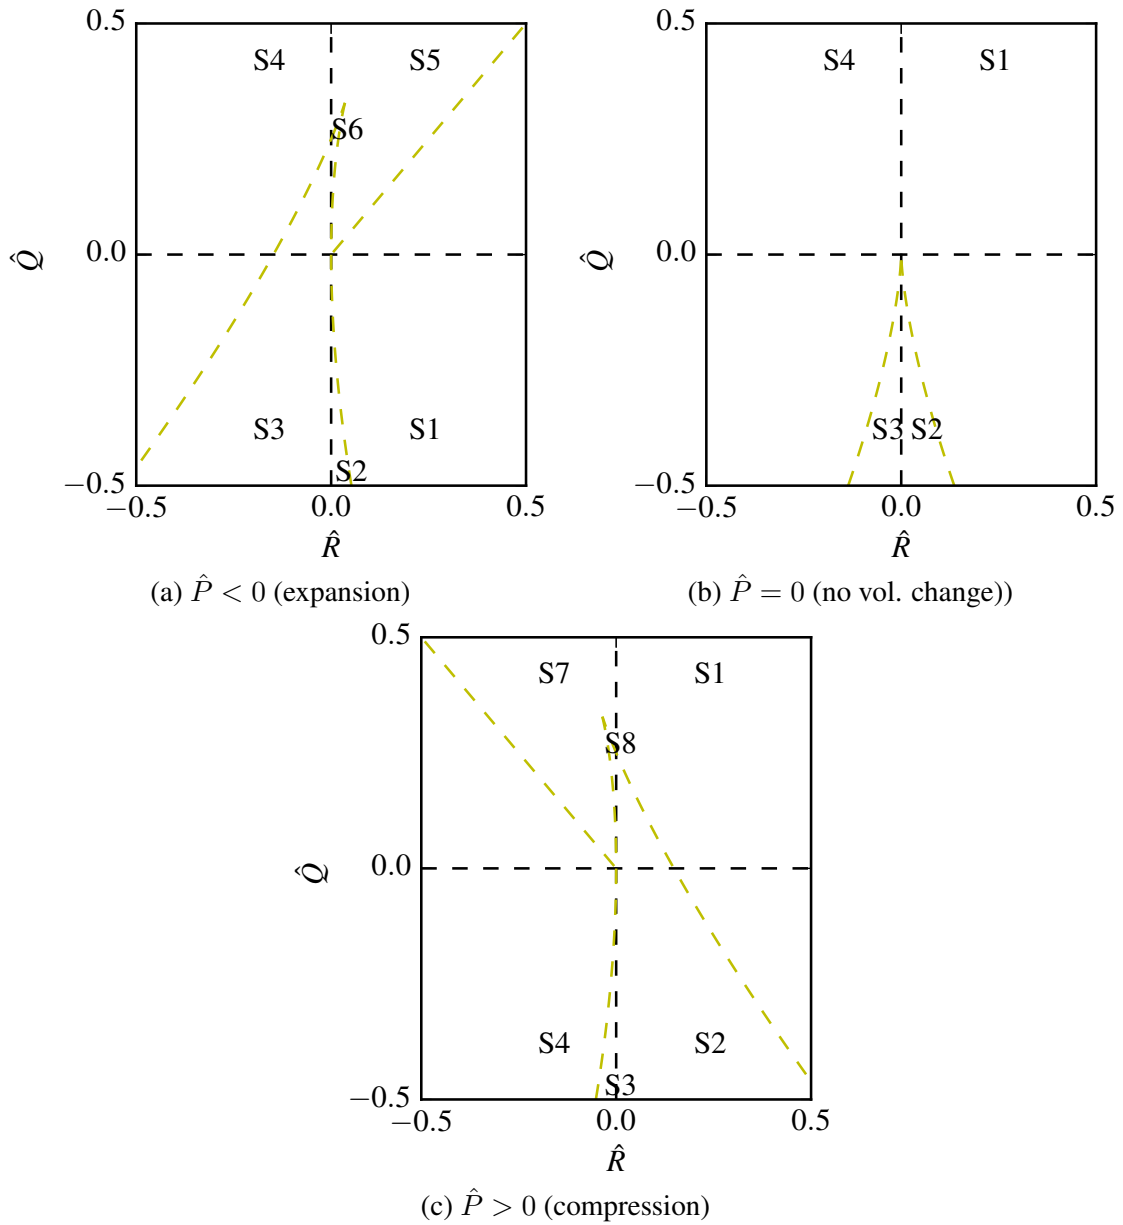


Figure 8.12: Sectors representing different critical points in the topological analysis of the invariants of the velocity gradient tensor. Sectors are defined in Table 8.2. Yellow dashed lines represent the surfaces where the discriminant is zero.

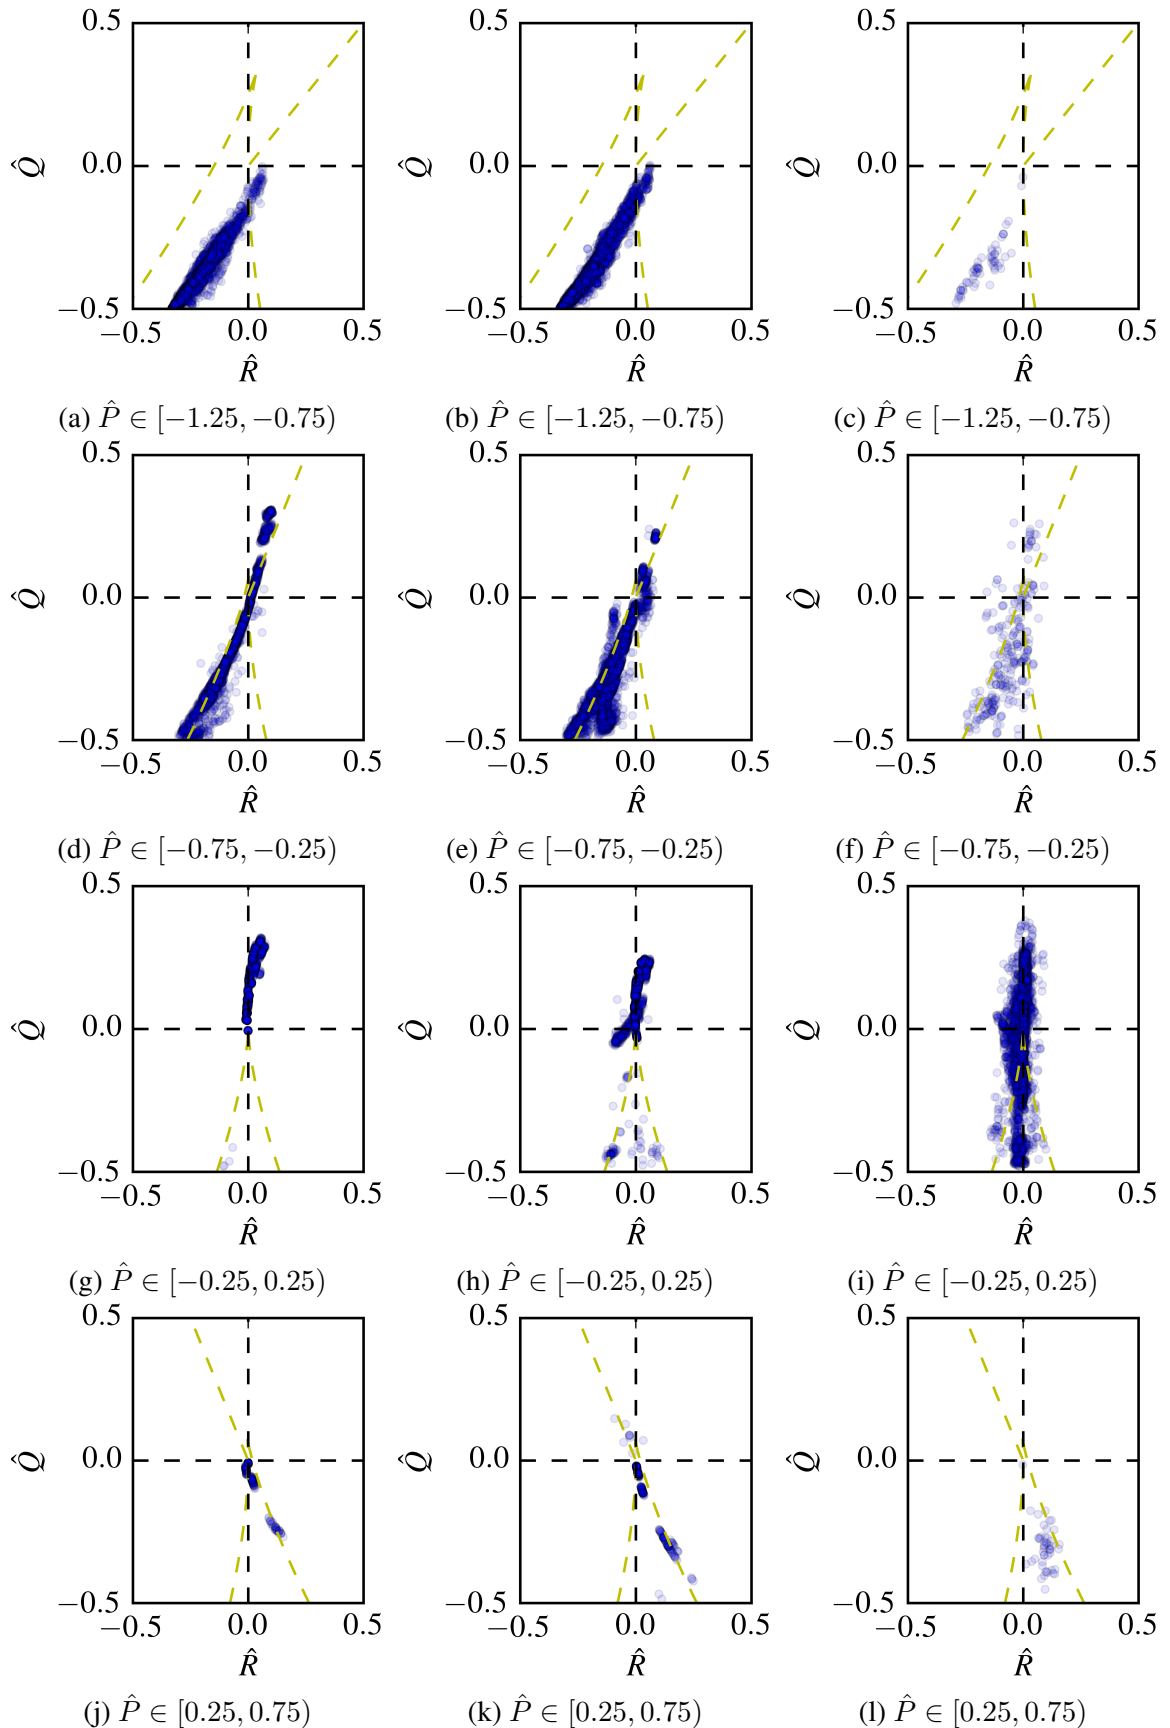


Figure 8.13: Topological analysis of the velocity gradient tensor invariants for the premixed case conditioned on $\tilde{T} = 800\text{ K}$, 1300 K , 1800 K from left to right.

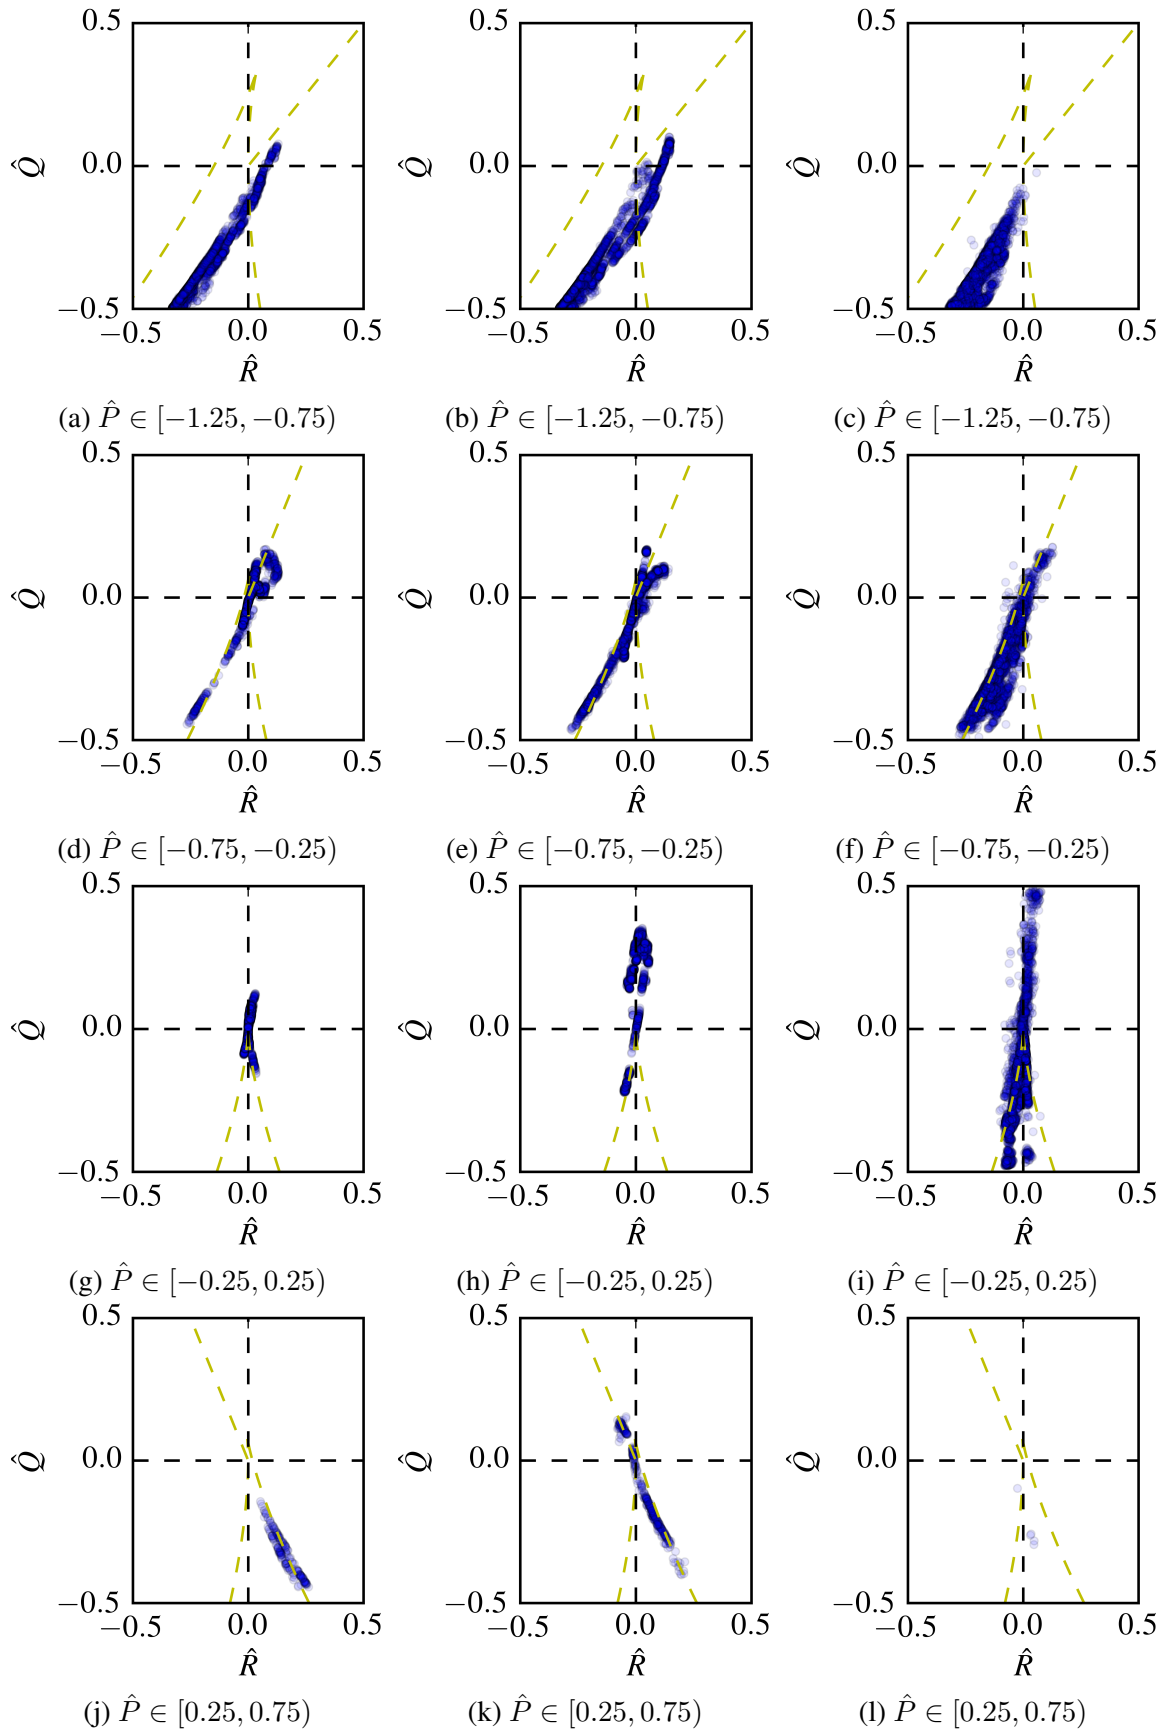


Figure 8.14: Topological analysis of the velocity gradient tensor invariants for the stratified case conditioned on $\tilde{T} = 800$ K, 1300 K, 1800 K from left to right.

CHAPTER 9

CONCLUSIONS AND FUTURE WORK

The objective of this work is to extend an existing compressible solver so that it is capable of resolving low Mach number, turbulent, reacting flows using the large-eddy simulation (LES) technique. The reliance on turbulence models in LES to provide the proper amount of dissipation representing the missing scales demands numerical schemes whose dissipative properties are understood, controlled, and minimized. The application of compressible solvers to the low Mach number regime often results in excess numerical dissipation [2, 3] that may overwhelm the contributions of the turbulence models. These low Mach number flows may be present globally or locally, such as behind bluff bodies or in recirculation zones. Therefore, this work focused on the development of a numerical scheme using a preconditioning approach that provides Mach number independent dissipative properties and the ability to use dedicated, stiff solvers to treat the chemical kinetics. A summary of the new contributions of this work follows, along with an assessment of what tasks remain and what new areas of research are now possible.

9.1 Summary

The popular MacCormack scheme, a member of the Lax-Wendroff family of schemes, serves as the baseline numerical approach for this work. It is attractive because it is a central scheme that is inherently stable and uses a predictor-corrector scheme with first order derivatives in each step. Because of this, it requires a small computational stencil, low memory footprint, is simple to implement, and it requires no user specified parameters or additional artificial dissipation. Despite its popularity and the large amount of literature regarding the scheme, its properties at low Mach numbers have not been reported.

Chapter 3 presents the original, 1971 form of the MacCormack scheme and analyzes its dissipative properties at low Mach numbers. The inherent dissipation in the scheme is found to suffer from the same degradation as artificial dissipation schemes in the low Mach number limit. This degradation is caused by the improper scaling of the eigenvalues multiplying the dissipation terms. The introduction of a preconditioning matrix to rescale the eigenvalues and to eliminate stiffness in the system alleviates the excessive dissipation problem at low Mach numbers when the preconditioning parameter equalizes the eigenvalues. The preconditioning parameter depends on the local flow conditions and, when used in a dual-time scheme for time accurate simulations, it also depends on the physical time step. As a result, for unsteady flow configurations, the preconditioning is effectively disabled and the modified equation analysis suggests the 1971 MacCormack scheme with preconditioning will behave the same as the scheme without preconditioning. The 1971 form of the scheme with preconditioning is called MC-PC.

Chapter 4 introduces a new form of the scheme, the Generalized MacCormack scheme, designed to overcome the limitations of the original scheme. This scheme is presented in both single-time form, suitable for compressible simulations, and a dual-time form with preconditioning, which provides Mach number independent behavior. The single-time form is similar to the original scheme; the only difference is an improvement to a non-linearity correction from the original scheme. That correction is found to produce large-scale pressure oscillations in otherwise smooth flows, while the new correction proposed here maintains stability in complex flows without introducing the pressure oscillations. The dual-time form of the scheme includes preconditioning at the pseudo-time level and splits the flux terms into convective and acoustic components. The acoustic component is treated using central differences for the interface pressure and an artificial dissipation using the properly scaled eigenvalues is added to the pressure field. This ensures stability as well as proper scaling of the numerical dissipation at low Mach numbers. The single-time and dual-time schemes are called GMC-ST and GMC-PC respectively.

A technique to develop robust preconditioner definitions based on a Fourier analysis of perturbation responses is presented in addition to the new Generalized MacCormack schemes. This technique is used to develop new preconditioning definitions designed to limit the amount of preconditioning applied when the flow contains strong non-linearities. These new definitions are more robust than existing preconditioner definitions and enhance the stability of the scheme when faced with strong pressure perturbations by many orders of magnitude. Additional analysis of the local time stepping methodology suggests that the time steps must be smooth to ensure stability, and a local smoothing operation is applied to the time steps. Lastly, a well-known technique for limiting the preconditioner definition in the neighborhood of stagnation points is analyzed and shown to eliminate the need for global preconditioner definitions.

Results for several non-reacting flows in Chapter 5 verifies the analysis of the MC-PC, GMC-ST and GMC-PC schemes. For large physical time steps, such that the unsteady preconditioner is not active and the preconditioning parameter equalizes the wave speeds, the MC-PC scheme and the GMC-PC scheme exhibit substantially less numerical dissipation than the GMC-ST in the convection of an inviscid vortex at a low Mach number. However, at small physical time steps where the preconditioning is effectively disabled, the MC-PC scheme solution matches that of the GMC-ST scheme, while the GMC-PC scheme exhibits superior solution quality as predicted by the modified equation analysis. Meanwhile, the convection of the same vortex at a moderate Mach number demonstrates the GMC-ST and GMC-PC scheme both predict accurate solutions with minimal numerical dissipation.

The same conclusions are found through investigations of a steady, viscous flow. The lid-driven cavity case at $M = 0.3$ using the GMC-ST and GMC-PC schemes matches the reference data well for both a low and a high Reynolds number. However, at $M = 0.001$, the GMC-ST scheme is unable to capture the correct solution on any of the grids tested, while the GMC-PC scheme performs as well as it did at $M = 0.3$. Furthermore, at both Mach numbers and Reynolds numbers, the GMC-PC converges at second-order rates with grid

refinement as expected. This, combined with the vortex case results, demonstrate the Mach number independence and minimized dissipative properties of the new GMC-PC scheme.

These cases also reveal the efficiency of the GMC-PC scheme. For the vortex convection at $M = 0.5$, the GMC-PC scheme and GMC-ST scheme predict solutions that are negligibly different, although the GMC-PC scheme requires more pseudo-time iterations per physical time step and the cost per step is increased due to matrix inversions. In this instance, the GMC-PC is not an effective scheme under these conditions and the GMC-ST scheme is preferred. However, at the low Mach number, $M = 0.001$, the GMC-PC scheme requires more iterations per physical time step at an increased cost on the same grid relative to the GMC-ST scheme, but the solution quality is superior. Tests with the GMC-ST scheme indicate the same solution quality using that scheme requires approximately 30 times the computational workload. Therefore, depending on the physical CFL number chosen, the GMC-PC scheme is at least 3 times more efficient and possibly up to an order of magnitude more efficient than the GMC-ST scheme. The GMC-PC scheme is even more efficient in the low Mach number lid-driven cavity case, where convergence required approximately 30 times fewer iterations than the GMC-ST scheme, while also providing a vastly improved solution quality. These performance and accuracy improvements are most noticeable at low Mach numbers.

Isotropic decaying turbulence at a low turbulent Mach number serves as a final verification of the improved dissipative properties of the GMC-PC scheme. The GMC-PC scheme at a physical CFL number based on the speed of sound, $CFL_c = 1$, and at a physical CFL number based on the initial RMS velocity, $CFL_{u'} = 1$, shows the solution is relatively insensitive to the physical CFL number chosen. CFL numbers larger than $CFL_{u'} = 1$ are not tested, as this exceeds the ideal CFL number for LES discussed in Section 5.5. The temporal evolution of both the large and small scales in the flow exhibit less numerical dissipation with the GMC-PC scheme relative to the GMC-ST scheme, and the turbulent kinetic energy

spectra contain a richer range of scales with the GMC-PC scheme. The GMC-ST scheme also indicates a tendency towards anisotropy not seen in the GMC-PC scheme.

Before moving to the schemes for reacting flows, these results should be put into a broader context. Most central and upwind schemes have low Mach number variants, but the Lax-Wendroff family of schemes is noticeably absent in that regime. The non-reacting GMC-PC scheme extends the MacCormack scheme, but the analysis of the MacCormack scheme also holds for other members of the Lax-Wendroff family. Therefore, the the inherent dissipation in all members of the family of schemes is expected to be overly dissipative at low Mach numbers without treatment, and the application of preconditioning to the scheme directly is unlikely to improve these properties for unsteady flows. Modification of the pressure terms combined with preconditioning improved the performance, and it is likely that all members of this family of schemes would benefit from a similar change. This opens up an entire family of second-order accurate, single-step or two-step explicit methods for application to low Mach number flows, and existing solvers that employ schemes from this family may be extended with minimal changes to the scheme.

Chapter 6 outlines extensions to the GMC-PC scheme to account for stiff chemical kinetics. These extensions are designed to remove the limitations of the traditional approach to integrating chemistry in dual-time schemes. This traditional approach linearizes the source term vector and uses the same temporal integration scheme as the fluxes in a fully-coupled approach. The robustness and accuracy of this scheme rely on *ad hoc* techniques to ensure stability, such as limiting the pseudo-time step based on the chemical kinetics or limiting the introduction of the source terms over several steps. This, in turn, reduces the effectiveness of preconditioning and is purely numerical in nature without any physical underpinning. The extensions proposed in Chapter 6 remove these *ad hoc* corrections by coupling the dual-time scheme to an external, dedicated ODE solver.

Two schemes are proposed based on techniques commonly used in single-time schemes. The first, a consistent splitting approach, uses DVODE to integrate both the fluxes and the

source terms together, with the fluxes treated explicitly using the Generalized MacCormack spatial differencing, and the source terms treated implicitly. This addresses the stiffness of the chemical source terms and ensures there are no steady-state errors due to the operator splitting. The second scheme is an extension of the common Strang operator splitting to dual-time schemes with the splitting performed in physical time. Although this extension is seemingly straight forward, the physical time derivative in the dual-time scheme must be altered. Specifically, the physical time derivative cannot be multi-step as the flux integration term must not contain any artifacts of the source term integration. Therefore, the physical time accuracy of this scheme is limited to first order.

A third scheme is also detailed to highlight the importance of the consistently split scheme. This scheme is the application to dual-time schemes of traditional operator splitting techniques used in single-time schemes. These operator splitting techniques are known to converge to incorrect steady state solutions due to splitting errors. Because the solution at each physical time level is the steady state solution in pseudo-time, any steady-state errors in the pseudo-time iterations will manifest as physical-time errors and degrade the quality of the solution.

A progressive series of cases in Chapter 7 based on a premixed, lean hydrogen mixture at elevated temperatures and pressures consistent with the conditions faced in Homogeneous Charge-Compression Ignition (HCCI) engines tests the accuracy and limitations of the proposed schemes. A zero-dimensional, constant volume reactor verifies that all schemes can predict the same solution as the stand-alone Cantera package. This is done with a range of time steps, ranging from a step size on the order of the chemical time scale, to a step size approximately 1000 times larger. The traditional, fully-coupled scheme failed to converge for large time steps, demonstrating the lack of efficiency in the traditional approach. The other three schemes all matched the Cantera solution for all time steps. The zero-dimensional reactor contained no flux terms, and therefore there are no errors due to splitting flux and source term operators. The physical split scheme is vastly more efficient in

this case, requiring only two DVODE calls per physical time step, as opposed to the schemes using DVODE at the pseudo-time level that required orders of magnitude more DVODE calls.

One-dimensional simulations with temperature and compositional inhomogeneities on a background flow demonstrate the behavior of the schemes when fluxes and source terms are present. The first case, with only a temperature variation, shows the success of the two new schemes. The consistently split scheme is both stable and accurate for all physical CFL numbers tested, while the scheme using operator splitting in pseudo-time deviates from the expected solution as the CFL number increases. Although the errors are minor in the volume averaged data for this case, it is indicative of the importance of eliminating steady-state errors and demonstrates the fallacy of directly applying single-time techniques to dual-time schemes. More significant errors are evident in the instantaneous, spatial fields, where this scheme shows a large temporal delay in the transition from ignition kernel to propagating deflagration waves. The physical split scheme also performs well in this case, providing accurate solutions for all but the largest time step size. At the largest time step, the source term integration operator introduces a large, localized heat release and this triggers sharp changes in thermodynamic variables. This leads to numerical instabilities generated by insufficient numerical dissipation in the central scheme around sharp interfaces. The traditional, fully-coupled scheme also fails at the largest time step size at the point of peak heat release rates.

The same case is repeated with the addition of a compositional inhomogeneity, atomic hydrogen, offset from the temperature inhomogeneity. The presence of atomic hydrogen begins the chain-branching reactions and establishes gradients in the radical species opposed to the temperature gradients. Under these conditions, the physical split scheme using Strang splitting is expected to exhibit large errors in the ignition behavior. Consistent with these expectations, the physical split scheme shows substantial deviations with the expected solution at larger physical CFL numbers, but performs unexpectedly well at smaller CFL

numbers. The pseudo-time scheme with operator splitting performs similarly, and the errors in both schemes are attributed to splitting the source terms and the diffusion operator.

The traditional, fully-coupled scheme exhibited numerical instabilities at all physical CFL numbers and is unable to provide a solution for this flow. On the other hand, the consistently split scheme is both robust and accurate. Errors in the solution do increase as the physical CFL number increases, with the ignition occurring earlier than expected. This is attributed to the diffusion of the temperature wave, leading to a more uniform temperature throughout the domain and an earlier ignition. The diffusion of the temperature wave is expected and due to the increased numerical dissipation at large physical CFL numbers.

The autoignition in a turbulent field with temperature fluctuations serves as a final measure of the two, new schemes. The results revealed weaknesses in both schemes. In the case of the consistently split scheme, the results for the two smaller CFL numbers tested demonstrate the same insensitivity as the non-reacting turbulence results. However, at larger CFL numbers, the scheme predicts markedly different solutions that miss both the transient and the steady-state expectations. This failure is attributed to the loss of numerical precision within the DVODE solver when it is calculating the Jacobian to assess stiffness in the system. The source terms are so much smaller than the flux terms that finite precision math is unable to generate their sum as different than the flux term itself, leading to DVODE incorrectly predicting a constant system. This, in turn, leads to inappropriately large time step values within DVODE and the solution is incorrect as a result. At the lower physical time step sizes, the pseudo-time step is of a similar order of magnitude as the chemical time scales and the loss of precision has less impact due to the small time steps DVODE is using. However, at large physical CFL numbers, preconditioning is more active and the pseudo-time step is several orders of magnitude larger than the chemical time scale and the loss of precision has more of an impact.

The physical split scheme also breaks down at large physical CFL numbers. In this instance, as with the previous cases, the scheme is numerically unstable at the largest CFL

number due to the large heat release in the chemical integration step. For large time steps where it is stable, the scheme predicts solutions that are insensitive of the time step until the point where the coupling between diffusion, convection, and reaction rates is dominant. When this transition occurs, the Strang splitting is unable to predict the correct solution due to operator splitting errors. However, for small enough time steps, the solution is insensitive to the physical CFL number used and the results are strikingly similar between the consistently split and physical split schemes. Even at these modest CFL numbers, the number of DVODE calls relative to the GMC-ST and consistently split schemes is reduced by at least one order of magnitude, representing a significant reduction in cost.

Two practical configurations are the final demonstration of the new scheme for reacting flows in Chapter 8. Due to the numerical limitations of the consistently split scheme, only the physical split scheme is employed. The first is a premixed, bluff-body stabilized flame frequently studied both numerically and experimentally. The case is not typically considered low Mach number, but it features a prominent recirculation zone behind the bluff-body that provides the stabilization for the flame. This region, by its nature, contains low velocities and is therefore important to resolve correctly. Previous studies using an explicit scheme and multiple grids are used as a baseline for comparison with the GMC-PC scheme. These previous studies failed to capture the recirculation region behind the bluff-body, even as the grid is refined, and instead predict a wake-like structure. The GMC-PC scheme, on the other hand, correctly captures the recirculation zone on the coarsest grid used for the baseline solutions. On this grid, the GMC-PC scheme is approximately 10 times slower than the explicit scheme used in the other study. The finest grid used in the previous study failed to capture the recirculation region, suggesting that the trends identified in the simpler configurations hold for complex cases as well. Specifically, although the solution on a given grid may be more expensive with the GMC-PC scheme, the solution at low Mach numbers is superior and resolving the same quality solution with the explicit scheme is more expensive.

The second configuration is a rod-stabilized V-flame over a two-dimensional slot burner. This case has both premixed and stratified configurations with experimental data available. This work is the first known LES of this configuration, and with inflow velocities of approximately 5 m s^{-1} , the flow is at a low Mach number. Temperature profiles throughout the flame match experimental data and RANS simulation data well for both the premixed and stratified cases. A new topological analysis of the results indicates predominantly tube-like structures dominate the flame fronts. The topology of the stratified case shows an increased amount of vortex stretching and the vorticity fields show a switch in the sign of the vorticity downstream of the rod. This suggests a strong baroclinic torque exists in the flame and the topology of the stratified flame contains fundamental differences relative to the premixed case.

9.2 Future Work

The dual-time iterative procedure offers numerous opportunities for future work and research regarding optimization and acceleration techniques. At each sub-iteration in the pseudo-time step operation, a linear system is solved at each cell. Although the preconditioning matrix is held fixed over all of the sub-iterations, the conservative to primitive Jacobian matrix added to the preconditioning matrix must be updated using the latest values. This prevents the inverse of the matrix from being stored and reused between sub-iterations and represents a majority of the increased cost per step when solving the preconditioned system. Unfortunately, for even the largest chemical mechanisms, these matrices are still dense and have no unique structural advantages to reduce the cost of solving the system. Because the matrices are unique at each cell and change each step, offloading them to dedicated acceleration hardware like GPUs or coprocessors likely will not be advantageous due to the memory transfer overhead. Therefore, techniques to update the matrix inverse directly without needing to recompute it could lead to a large reduction in cost. This may be most useful towards the end of the pseudo-time convergence, when the changes in the solution

are small and may be confined to one or two entries in the matrix inverse, at which point techniques for performing rank-one updates directly to the matrix inverse may be possible.

Along similar lines, it is possible to reformulate the consistently split scheme so that the physical time derivative remains implicit in pseudo-time alongside the implicit source terms. This way, the Jacobian does not appear on the left hand side of the system and only the preconditioning matrix is required. Because this matrix is constant over the sub-iterations, its inverse can be computed once and stored for all sub-iterations. This inverse can be computed numerically, or the inverse preconditioning matrix may be written directly. Directly writing this matrix will make the solution of the linear system simply a matrix-vector multiplication operation where the matrix may be evaluated and stored once per pseudo-iteration. This option was not explored in this work because of the numerical precision difficulties already present in the scheme.

The behavior of both the consistently split scheme and the physical-time split scheme is promising, but unsatisfying due to the numerical challenges encountered. The physical split scheme is limited to moderate physical CFL numbers to ensure accuracy in physical time. This is not unexpected given the nature of the splitting errors in physical time, and the scheme does result in significant cost savings at low Mach numbers due to the reduced number of DVODE calls. However, the time step in this case is dictated by the chemical dynamics and the allowable time steps are far from the ideal step sizes for LES. The consistently split scheme has the potential to allow physical time steps that are based on the fluid mechanics and not the flame dynamics, but the numerical loss of precision stymied that effort. For this scheme, the flame dynamics and fluid mechanics remain coupled through the pseudo-time iteration and conceptually it is possible the physical time step may be chosen much larger and still provide accurate solutions. Reformulating the consistently split scheme to eliminate or mitigate this loss of precision is the highest priority in continuing this work.

An outstanding question is how the solution accuracy for chemically reacting flows is affected by the physical time step size. This work showed that non-reacting turbulence pro-

perties are relatively insensitive to the physical CFL number, and purely reacting problems are likewise insensitive. But the coupling between the source terms and the fluxes may lead to more stringent requirements on the physical time step selection than either component does alone.

Finally, evaluation of features and improvements to the underlying scheme may extend the range of conditions the scheme can solve. For instance, the artificial dissipation is not TVD and extending the range of Mach numbers into supersonic or hypersonic regimes would require modifications. These modifications may come through reformulating an upwind scheme such as MUSCL or WENO to account for the preconditioned wave speeds, and combining the GMC-PC scheme with the preconditioned upwind scheme to create a hybrid scheme. Alternatively, a TVD version of the MacCormack scheme can be implemented and extended for low Mach numbers. Further verifications of the other types of MacCormack discretizations are also needed, specifically the higher order versions of the scheme.

Appendices

APPENDIX A

NONLINEAR INSTABILITIES IN THE MACCORMACK SCHEME

The generalized MacCormack scheme for single-time integration (GMC-ST) introduced in Chapter 4 contains an update to the non-linearity switch first proposed in the 1971 version of the scheme [39]. It was noted in Chapter 5 that the original switch generates large scale pressure fluctuations in smooth flows whereas the new switch eliminates a majority of the oscillations. In this appendix, the single-time behavior of the non-linear switch is assessed through model equations designed to trigger the instabilities. Further discussion of the oscillations in the inviscid vortex flow are also presented. These sections are extracted from Gallagher et al. [117].

A.1 Model Equations

Model equations provide a test bed to assess the non-linear switch in its original and proposed forms relative to the scheme without the switch. This is done using the linear wave equation and the non-linear Burgers' equation. The linear wave equation given by:

$$\frac{\partial u}{\partial t} + c \frac{\partial u}{\partial x} = 0 \quad (\text{A.1})$$

where the wave speed c is a constant in time but varies in space and u is the unknown. The initial conditions and spatial variation in the wave speed are designed to trigger the conditions that turn on the non-linearity correction:

$$c(x) = \begin{cases} 1 & \text{if } x > 0 \\ -1 & \text{if } x \leq 0 \end{cases} \quad (\text{A.2})$$

and the grid is created such that the $i + 1/2$ face of the cell centered on $x = 0$ has $c_i < 0$ and $c_{i+1} > 0$. The initial condition on the scalar is a Gaussian pulse.

The scheme without a switch is discontinuous at the cell face, with a positive value on the left side of the origin and negative value on the right as shown in Figure A.1a. This solution is consistent with the observations of MacCormack [39], where the sum of the solution across the cell face with divergent velocities is zero, even though each side may grow unbounded. Although in this case the solution remains bounded, this jump represents an infinite source of the scalar at the origin. On the other hand, both the 1971 form of the non-linearity switch as well as the new form proposed here are continuous and zero through the origin. Each half of the Gaussian wave convects away from the interface as expected. Oscillations in all solutions occur due to the Gibbs phenomenon at the discontinuities.

The switch is based on the need to control non-linearities, so the second test uses the inviscid Burgers' equation:


$$\frac{\partial u}{\partial t} + \frac{1}{2} \frac{\partial uu}{\partial x} = 0 \quad (\text{A.3})$$

with initial conditions again designed to trigger the switch:

$$u(x, t = 0) = \begin{cases} -\frac{x}{x_c - x_{0.25}} - 1.0 & \text{if } x_{0.25} \leq x < x_c \\ -\frac{x}{x_{0.75} - x_c} + 1.0 & \text{if } x_c \leq x < x_{0.75} \\ 0 & \text{else} \end{cases} \quad (\text{A.4})$$

where x_c is the center of the domain, $x_{0.25}$ and $x_{0.75}$ are the coordinates at 25% and 75% of the domain respectively.

The expected solution is two ramps that move away from the origin. Results in Figure A.1b show the same trends as with the linear equation. Specifically, both the 1971 switch and the new switch produce identical solutions. However, the scheme without the switch again produces an infinite source of the scalar at the origin where the discontinuity is located. Despite the non-linear equation, the instabilities seen in the 1971 switch in the



(a) Linear wave equation

(b) Inviscid Burgers' equation

Figure A.1: Comparison of the schemes without the nonlinearity switch, with the original 1971 nonlinearity switch, and with the new nonlinearity switch given by Equation (4.1) on the linear and nonlinear model equations. Reprinted with permission from Gallagher et al.[117].


Euler and Navier-Stokes equations do not appear. This is suggestive that the instabilities are not simply non-linear. They may be due to the coupling between the pressure and velocity in the full Euler equation set, but a stability analysis of a coupled set of simple non-linear equations may reveal the extent to which the instabilities are specific to the Euler equations. Despite the lack of instabilities in these simple tests, the model equations confirm the new switch produces the same solution as the original switch.

A.2 Vortex Convection


Based on the results of the model equations, the full Euler equation set is required to reproduce the instabilities. The vortex test case described in Section 5.2 at $M = 0.3$ demonstrates the severe degradation in the solution when the original 1971 form of the switch is used. This is shown in Figure A.2 and compared to the new form of the switch and the solution without the switch. The original 1971 switch generates large-scale pressure

oscillations in the solution. For this flow, the simulation remains stable when the switch is disabled and the resulting pressure field is well behaved. This is not a general solution, however, because the switch is required for stability in more complex flows.

The new switch proposed in this work leads to a pressure solution similar to the case without the switch. All three pressure solutions are shown in Figure A.2a. The corresponding velocity profiles are shown in Figure A.2b. The velocity profiles are unaffected by the pressure oscillations and remain well-behaved. The local minima and maxima in the velocity profiles are larger in magnitude with the original switch compared to the new switch and the case without the switch. The pressure oscillations are due to insufficient dissipation with the 1971 switch, which also explains the differences in the velocity predictions.



(a) Pressure profiles across the vortex.



(b) V-velocity profiles across the vortex.

Figure A.2: Velocity and pressure profiles across the vortex using the original, 1971 nonlinearity switch, the new switch proposed in this work (Equation (4.1)), and no nonlinearity switch for one turn around the domain. Reprinted with permission from Gallagher et al. [117].

APPENDIX B

COMPUTATIONAL SETTINGS AND PERFORMANCE

The dual-time scheme with preconditioning in this work introduces a number of new settings that control the behavior the scheme. These settings will affect the robustness, accuracy, and performance of the scheme. Therefore, all of the settings used for each of the cases presented in this work are summarized below and may be used as a guide for setting up similar problems or as a starting point for new problems. Additionally, the computational cost of the simulations are also provided where available, again to serve as a baseline for establishing expectations. Where appropriate, notes are added to each case to describe how the settings were chosen and what effect might happen based on changes to them, if known.

Flow over Bump

This case is described in Section 5.1.

Table B.1: Preconditioning settings for the flow over a bump case.

Case	Pressure PC	Viscous PC	Unsteady PC, l	Max Steps	Max Norm	Flatness Steps	Flatness Measure	Convergence Decades	Local Timestep Smoothing	Smoothing Steps	Smoothing Weight	JST $k^{(2)}$	JST $k^{(4)}$
	x	x	2.0	—	1×10^{-9}	10	1×10^{-12}	4	x	—	—	—	—

Notes

- Computational performance was not measured for this case.
- Turning on the local time step smoothing causes the case to take excessively long to converge.
- Turning on the pressure preconditioner causes the case to take excessively long to converge.

Inviscid Vortex Convection

This case is described in Section 5.2.

Table B.2: Preconditioning settings for the inviscid vortex case.

Case	Pressure PC	Viscous PC	Unsteady PC, l	Max Steps	Max Norm	Flatness Steps	Flatness Measure	Convergence Decades	Local Timestep Smoothing	Smoothing Steps	Smoothing Weight	JST $\kappa^{(2)}$	JST $\kappa^{(4)}$
$CFL_c = 1$	✓	✗	0.3112	10000	1×10^{-9}	10	1×10^{-12}	—	✗	—	—	—	—
$CFL_u = 1$	✓	✗	0.3112	10000	1×10^{-9}	10	1×10^{-12}	2–4	✗	—	—	—	—

Notes

- Computational performance (wall time) was not recorded for this case, but iterative efficiency and solution quality metrics indicate the preconditioned scheme is at least 3 times more efficient than the non-preconditioned scheme; see Section 5.2 for details.
- Provided the number of decades the residual converges is greater than or equal to 2, the results are insensitive to the convergence criteria.

- Less than 2 decades of convergence leads to numerical divergence.

Lid-Driven Cavity

This case is described in Section 5.3.

Table B.3: Preconditioning settings for the lid-driven cavity case.

Case	Pressure PC	Viscous PC	Unsteady PC, l	Max Steps	Max Norm	Flatness Steps	Flatness Measure	Convergence Decades	Local Timestep Smoothing	Smoothing Steps	Smoothing Weight	JST $\kappa^{(2)}$	JST $\kappa^{(4)}$
All	x	✓	—	300000	1×10^{-9}	—	—	—	x	—	—	—	—

Table B.4: Relative wall times for the lid-driven cavity case. All timing metrics are relative to the GMC-ST simulation on the same grid.

Case	16×16	32×32	64×64	128×128	256×256
$Re = 100, M = 0.3$	227.91	115.02	5.28	1.09	—
$Re = 100, M = 0.001$	0.011	0.030	0.026	0.121	—
$Re = 1000, M = 0.3$	0.216	0.149	0.259	3.13	1.11
$Re = 1000, M = 0.001$	0.019	0.028	0.222	0.188	0.167

Notes

- The relative wall times indicate the GMC-PC scheme is significantly faster than the GMC-ST scheme for all low Mach number cases. It is important to remember that the GMC-ST scheme is unable to accurately capture the results under these conditions.

- The GMC-PC scheme exhibited stalled convergence in the $M = 0.3$ cases, but the flatness convergence criteria was not used to detect this. The timing studies are therefore skewed in favor of the GMC-ST scheme.
- Turning on the local time step smoothing prevents proper convergence in this case.

Decaying Isotropic Turbulence

This case is described in Section 5.4.

Table B.5: Preconditioning settings for the non-reacting isotropic turbulence case.

Case	Pressure PC	Viscous PC	Unsteady PC, l	Max Steps	Max Norm	Flatness Steps	Flatness Measure	Convergence Decades	Local Timestep Smoothing	Smoothing Steps	Smoothing Weight	JST $\kappa^{(2)}$	JST $\kappa^{(4)}$
$CFL_c = 1$	✓	✓	0.7	100	1×10^{-9}	10	1×10^{-12}	10	✗	—	—	—	—
$CFL_u = 1$	✓	✓	0.7	100000	1×10^{-9}	10	1×10^{-12}	10	✗	—	—	—	—

Notes

- Computational performance was not measured for this case.
- Turning on the local time step smoothing causes the case to take excessively long to converge.

Zero-Dimensional Ignition

This case is described in Section 7.1.

Table B.6: Preconditioning settings for the zero-dimensional ignition cases.

Case	Pressure PC	Viscous PC	Unsteady PC, l	Max Steps	Max Norm	Flatness Steps	Flatness Measure	Convergence Decades	Local Timestep Smoothing	Smoothing Steps	Smoothing Weight	JST $\kappa^{(2)}$	JST $\kappa^{(4)}$
Scheme A, all CFLs	✓	✓	0.01	20000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme B, all CFLs	✓	✓	0.01	—	1×10^{-12}	—	—	—	✓	10	0.05	—	—
Scheme C, all CFLs	✓	✓	0.01	20000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme D, all CFLs	✓	✓	0.01	20000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—

Table B.7: Relative wall times for the zero-dimensional ignition case. All timing metrics are relative to the GMC-ST simulation.

Scheme	$CFL_c = 1$	$CFL_c = 10$	$CFL_c = 100$	$CFL_c = 1000$
Scheme A	3.07	0.473	0.104	0.073
Scheme B	1.40	0.140	0.014	0.002
Scheme C	3.09	0.522	—	—
Scheme D	3.38	0.593	0.125	0.091

Notes

- For this case, Scheme B did not require any pseudo-iterations because the flux terms are not computed; therefore, preconditioning settings are not relevant.
- Scheme C took excessively long to converge at larger CFL numbers.
- Local time step smoothing was used in this case, but it is not relevant because the cells all contain the same solution values. It does, however, affect performance and so it is reported for completeness.

- Scheme C took excessively long times to converge at the two largest CFL numbers, so the simulation times are not recorded.

One-Dimensional Ignition

These cases are described in Section 7.2

Table B.8: Preconditioning settings for the one-dimensional ignition cases.

Scheme	Pressure PC	Viscous PC	Unsteady PC, l	Max Steps	Max Norm	Flatness Steps	Flatness Measure	Convergence Decades	Local Timestep Smoothing	Smoothing Steps	Smoothing Weight	JST $\kappa^{(2)}$	JST $\kappa^{(4)}$
Case 1A													
Scheme A, CFL _c = 1	✓	✓	0.01	5000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme A, CFL _c = 10	✓	✓	0.01	10000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme A, CFL _c = 100	✓	✓	0.01	20000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme A, CFL _c = 1000	✓	✓	0.01	50000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme B, CFL _c = 1	✓	✓	0.01	5000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme B, CFL _c = 10	✓	✓	0.01	5000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme B, CFL _c = 100	✓	✓	0.01	10000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme C, CFL _c = 1	✓	✓	0.01	10000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme C, CFL _c = 10	✓	✓	0.01	10000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme C, CFL _c = 100	✓	✓	0.01	20000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme D, CFL _c = 1	✓	✓	0.01	5000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme D, CFL _c = 10	✓	✓	0.01	10000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme D, CFL _c = 100	✓	✓	0.01	20000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme D, CFL _c = 1000	✓	✓	0.01	50000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Case 1B													
Scheme A, CFL _c = 1	✓	✓	0.01	5000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme A, CFL _c = 10	✓	✓	0.01	10000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme A, CFL _c = 100	✓	✓	0.01	20000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme A, CFL _c = 1000	✓	✓	0.01	50000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme B, CFL _c = 1	✓	✓	0.01	5000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme B, CFL _c = 10	✓	✓	0.01	5000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme B, CFL _c = 100	✓	✓	0.01	10000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme B, CFL _c = 1000	✓	✓	0.01	50000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme D, CFL _c = 1	✓	✓	0.01	5000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme D, CFL _c = 10	✓	✓	0.01	10000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme D, CFL _c = 100	✓	✓	0.01	20000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme D, CFL _c = 1000	✓	✓	0.01	50000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—

Three-Dimensional Ignition

This case is described in Section 7.3.

Table B.9: Preconditioning settings for the three-dimensional ignition case.

Case	Pressure PC	Viscous PC	Unsteady PC, l	Max Steps	Max Norm	Flatness Steps	Flatness Measure	Convergence Decades	Local Timestep Smoothing	Smoothing Steps	Smoothing Weight	JST $\kappa^{(2)}$	JST $\kappa^{(4)}$
Scheme A, all CFLs	✓	✓	0.0041	20000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—
Scheme B, all CFLs	✓	✓	0.0041	20000	1×10^{-12}	10	1×10^{-14}	4	✓	10	0.05	—	—

Volvo Flymotor

This case is described in Section 8.1.

Table B.10: Preconditioning settings for the Volvo case.

Case	Pressure PC	Viscous PC	Unsteady PC, l	Max Steps	Max Norm	Flatness Steps	Flatness Measure	Convergence Decades	Local Timestep Smoothing	Smoothing Steps	Smoothing Weight	JST $\kappa^{(2)}$	JST $\kappa^{(4)}$
	✓	✓	0.04	1000	1×10^{-6}	10	1×10^{-12}	4	✓	10	0.05	0.25	0.04

Notes

- Computational timing was not performed for this case.

- Less than 4 decades convergence leads to an eventual divergence in the code.
- The flame creates gradients large enough to cause numerical instabilities unless JST dissipation is used.

Cambridge Slot Burner

This case is described in Section 8.2.

Table B.11: Preconditioning settings for the Cambridge slot burner cases.

Case	Pressure PC	Viscous PC	Unsteady PC, τ	Max Steps	Max Norm	Flatness Steps	Flatness Measure	Convergence Decades	Local Timestep Smoothing	Smoothing Steps	Smoothing Weight	JST $\kappa^{(2)}$	JST $\kappa^{(4)}$
All	✓	✓	0.03	500	1×10^{-9}	10	1×10^{-12}	2	✗	—	—	0.25	0.007

Notes

- Computational timing was not performed for this case.
- The code is modified such that JST dissipation is only active beyond 30 mm above the burner. Within the primary measurement window, the grid is fine enough to resolve the flame. However, numerical instabilities exist as the grid coarsens towards the outflow, so the artificial dissipation is used for stability there.

REFERENCES

- [1] F. M. Denaro, “What does finite volume-based implicit filtering really resolve in Large-Eddy Simulations?” *Journal of Computational Physics*, vol. 230, no. 10, pp. 3849–3883, 2011.
- [2] E. Turkel, “Preconditioning techniques in computational fluid dynamics,” *Annual Review of Fluid Mechanics*, vol. 31, no. 1, pp. 385–416, 1999.
- [3] V Sankaran and C. Merkle, “Analysis of preconditioning methods for Euler and Navier-Stokes computations,” *Von Karman Institute Lecture Series on Computational Fluid Dynamics*, 1999.
- [4] R. H. Pletcher, J. C. Tannehill, and D. Anderson, *Computational fluid mechanics and heat transfer*. CRC Press, 2012.
- [5] J. M. Weiss and W. A. Smith, “Preconditioning applied to variable and constant density flows,” *AIAA journal*, vol. 33, no. 11, pp. 2050–2057, 1995.
- [6] A. J. Chorin, “A numerical method for solving incompressible viscous flow problems,” *Journal of Computational Physics*, vol. 135, no. 2, pp. 118–125, 1997.
- [7] S. Patankar, *Numerical heat transfer and fluid flow*. CRC press, 1980.
- [8] S. Acharya, B. Baliga, K. Karki, J. Murthy, C. Prakash, and S. Vanka, “Pressure-based finite-volume methods in computational fluid dynamics,” *Journal of Heat Transfer*, vol. 129, no. 4, pp. 407–424, 2007.
- [9] K. Karki and S. Patankar, “Pressure based calculation procedure for viscous flows at all speeds in arbitrary configurations,” *AIAA Journal*, vol. 27, no. 9, pp. 1167–1174, 1989.
- [10] C.-D. Munz, S. Roller, R. Klein, and K. J. Geratz, “The extension of incompressible flow solvers to the weakly compressible regime,” *Computers & Fluids*, vol. 32, no. 2, pp. 173–196, 2003.
- [11] R. Issa, “Solution of the implicitly discretised fluid flow equations by operator-splitting,” *Journal of Computational Physics*, vol. 62, no. 1, pp. 40–65, 1986.
- [12] R. Issa, B. Ahmadi-Befrui, K. Beshay, and A. Gosman, “Solution of the implicitly discretised reacting flow equations by operator-splitting,” *Journal of Computational Physics*, vol. 93, no. 2, pp. 388–410, 1991.

- [13] N. Bressloff, “A parallel pressure implicit splitting of operators algorithm applied to flows at all speeds,” *International Journal for Numerical Methods in Fluids*, vol. 36, no. 5, pp. 497–518, 2001.
- [14] C. Rhie and W. Chow, “Numerical study of the turbulent flow past an airfoil with trailing edge separation,” *AIAA Journal*, vol. 21, no. 11, pp. 1525–1532, 1983.
- [15] A. J. Chorin, “A numerical method for solving incompressible viscous flow problems,” *Journal of Computational Physics*, vol. 2, no. 1, pp. 12–26, 1967.
- [16] J. H. Ferziger and M. Peric, *Computational methods for fluid dynamics*. Springer Science & Business Media, 2002.
- [17] A. Tomboulides, J. Lee, and S. Orszag, “Numerical simulation of low Mach number reactive flows,” *Journal of Scientific Computing*, vol. 12, no. 2, pp. 139–167, 1997.
- [18] A. G. Tomboulides and S. A. Orzag, “A quasi-two-dimensional benchmark problem for low Mach number compressible codes,” *Journal of Computational Physics*, vol. 146, no. 2, pp. 691–706, 1998.
- [19] A. Meister, “Asymptotic single and multiple scale expansions in the low Mach number limit,” *SIAM Journal on Applied Mathematics*, vol. 60, no. 1, pp. 256–271, 1999.
- [20] Y.-H. Choi and C. L. Merkle, “The application of preconditioning in viscous flows,” *Journal of Computational Physics*, vol. 105, no. 2, pp. 207–223, 1993.
- [21] D. Choi and C. L. Merkle, “Application of time-iterative schemes to incompressible flow,” *AIAA Journal*, vol. 23, no. 10, pp. 1518–1524, 1985.
- [22] E. Turkel, “Preconditioned methods for solving the incompressible and low speed compressible equations,” *Journal of Computational Physics*, vol. 72, no. 2, pp. 277–298, 1987.
- [23] F. N. Felten and T. S. Lund, “Kinetic energy conservation issues associated with the collocated mesh scheme for incompressible flow,” *Journal of Computational Physics*, vol. 215, no. 2, pp. 465–484, 2006.
- [24] A. Jameson, W. Schmidt, and E. Turkel, “Numerical solutions of the Euler equations by finite volume methods using Runge-Kutta time-stepping schemes,” in *AIAA 14th Fluid and Plasma Dynamics Conference*, AIAA Paper 81-1259.
- [25] D. V. Gaitonde, J. S. Shang., and J. L. Young, “Practical aspects of high-order accurate finite-volume schemes for electromagnetics,” in *35th AIAA AIAA Aerospace Sciences Meeting*, AIAA Paper 97-0363, Reno, Nevada, 1997.

- [26] D. V. Gaitonde and M. R. Visbal, “High-order schemes for Navier-Stokes equations: Algorithm and implementation into FDL3DI,” Air Force Research Laboratory, Wright-Patterson Air Force Base, Ohio, Tech. Rep. AFRL-VA-WP-TR-1998-3060, 1998.
- [27] ———, “Further development of a Navier-Stokes solution procedure based on higher-order formulas,” in *37th AIAA AIAA Aerospace Sciences Meeting*, AIAA Paper 99-0557, Reno, Nevada, 1999.
- [28] D. V. Gaitonde, J. S. Shang., and J. L. Young, “Practical aspects of higher-order numerical schemes for wave propagation phenomena,” *International Journal for Numerical Methods in Engineering*, vol. 45, pp. 1849–1869, 1999.
- [29] D. V. Gaitonde and M. R. Visbal, “Padé-type higher-order boundary filters for the Navier–Stokes equations,” *AIAA Journal*, vol. 38, no. 11, pp. 2103–2112, 2000.
- [30] A. Jameson, “Analysis and design of numerical schemes for gas dynamics, 1: Artificial diffusion, upwind biasing, limiters and their effect on accuracy and multigrid convergence,” *International Journal of Computational Fluid Dynamics*, vol. 4, no. 3-4, pp. 171–218, 1995.
- [31] ———, “Analysis and design of numerical schemes for gas dynamics, 2: Artificial diffusion and discrete shock structure,” *International Journal of Computational Fluid Dynamics*, vol. 5, no. 1-2, pp. 1–38, 1995.
- [32] M.-S. Liou and C. J. Steffen, “A new flux splitting scheme,” *Journal of Computational physics*, vol. 107, no. 1, pp. 23–39, 1993.
- [33] R. V. Chima and M.-S. Liou, “Comparison of the AUSM+ and H-CUSP schemes for turbomachinery applications,” in *16th AIAA Computational Fluid Dynamics Conference*, AIAA paper 2003-4120, 2003.
- [34] S. Venkateswaran and C. Merkle, “Efficiency and accuracy issues in contemporary CFD algorithms,” in *Fluids 2000*, AIAA Paper 2000-2251, 2000.
- [35] M.-S. Liou, “A sequel to AUSM, part II: AUSM+-up for all speeds,” *Journal of Computational Physics*, vol. 214, no. 1, pp. 137–170, 2006.
- [36] J. Sachdev, A. Hosangadi, and V. Sankaran, “Improved flux formulations for unsteady low Mach number flows,” in *42nd AIAA Fluid Dynamics Conference and Exhibit*, AIAA Paper 2012-3067, 2012.
- [37] D. Folkner, A. Katz, and V. Sankaran, “An unsteady preconditioning scheme based on convective-upwind split-pressure (CUSP) artificial dissipation,” *International Journal for Numerical Methods in Fluids*, vol. 78, pp. 1–16, 2014.

- [38] R. MacCormack, “The effect of viscosity in hypervelocity impact cratering,” *Frontiers of Computational Fluid Dynamics*, pp. 27–44, 1969.
- [39] R. MacCormack, “Numerical solution of the interaction of a shock wave with a laminar boundary layer,” in *Proceedings of the Second International Conference on Numerical Methods in Fluid Dynamics*, ser. Lecture Notes in Physics, M. Holt, Ed., vol. 8, Springer Berlin Heidelberg, 1971, pp. 151–163.
- [40] C. Merkle, “Time-accurate unsteady incompressible flow algorithms based on artificial compressibility,” in *Fluid Dynamics and Co-located Conferences*, AIAA Paper 87-1137, 1987.
- [41] J. P. Withington, V. Yang, and J. Shuen, “A time-accurate implicit method for chemically reacting flows at all mach numbers,” in *29th AIAA Aerospace Sciences Meeting*, AIAA Paper 91-0581, 1991.
- [42] S. Venkateswaran, J. Weiss, C. Merkle, and Y.-H. Choi, “Propulsion-related flowfields using the preconditioned Navier-Stokes equations,” in *Joint Propulsion Conference and Exhibit*, AIAA Paper 92-3437, 1992.
- [43] J. M. Weiss and J. Murtry, “Computation of reacting flowfields using unstructured adaptive meshes,” in *33rd AIAA Aerospace Sciences Meeting*, AIAA Paper 95-0870, 1995.
- [44] A. Jameson, “Time dependent calculations using multigrid, with applications to unsteady flows past airfoils and wings,” in *AIAA 10th Computational Fluid Dynamics Conference*, AIAA Paper 91-1596, 1991.
- [45] T. Poinsot and D. Veynante, *Theoretical and Numerical Combustion*. Edwards, 2001.
- [46] C. Lian, G. Xia, and C. L. Merkle, “Impact of source terms on reliability of CFD algorithms,” *Computers & Fluids*, vol. 39, no. 10, pp. 1909–1922, 2010.
- [47] A. C. Hindmarsh, “ODEPACK, a systematized collection of ODE solvers,” in *Scientific Computing, Applications of Mathematics and Computing to the Physical Sciences, Volume I*, Stepleman, Carver, Peskin, Ames, and Vichnevetsky, Eds., vol. 1, IMACS World Congress, 1983, pp. 55–64.
- [48] Y. Gao, Y. Liu, Z. Ren, and T. Lu, “A dynamic adaptive method for hybrid integration of stiff chemistry,” *Combustion and Flame*, vol. 162, no. 2, pp. 287–295, 2015.
- [49] C. Xu, Y. Gao, Z. Ren, and T. Lu, “A sparse stiff chemistry solver based on dynamic adaptive hybrid integration for efficient combustion simulations,” in *9th U.S. National Combustion Meeting*, 2015.

- [50] Y. Gao, Y. Liu, Z. Ren, and T. Lu, “A dynamic adaptive method for hybrid integration of stiff chemistry,” *Combustion and Flame*, vol. 162, no. 2, pp. 287–295, 2015.
- [51] V. Sankaran and J. Oefelein, “Advanced preconditioning strategies for chemically reacting flows,” in *45th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA Paper 2007-1432, 2007.
- [52] S Venkateswaran, M Deshpande, and C. Merkle, “The application of preconditioning to reacting flow computations,” in *12th Computational Fluid Dynamics Conference*, AIAA Paper 95-1673, 1995.
- [53] M. A. Hansen and J. C. Sutherland, “Dual timestepping methods for detailed combustion chemistry,” *Combustion Theory and Modelling*, pp. 1–17, 2016.
- [54] H. N. Najm, P. S. Wyckoff, and O. M. Knio, “A semi-implicit numerical scheme for reacting flow: I. stiff chemistry,” *Journal of Computational Physics*, vol. 143, no. 2, pp. 381–402, 1998.
- [55] H. N. Najm, O. M. Knio, P. H. Paul, and P. S. Wyckoff, “Response of stoichiometric and rich premixed methane-air flames to unsteady strain rate and curvature,” *Combustion Theory and Modelling*, vol. 3, no. 4, pp. 709–726, 1999.
- [56] O. M. Knio, H. N. Najm, and P. S. Wyckoff, “A semi-implicit numerical scheme for reacting flow: II. stiff, operator-split formulation,” *Journal of Computational Physics*, vol. 154, no. 2, pp. 428–467, 1999.
- [57] O. M. Knio and H. N. Najm, “Effect of stoichiometry and strain rate on transient flame response,” *Proceedings of the Combustion Institute*, vol. 28, no. 2, pp. 1851–1857, 2000.
- [58] V Sankaran and S Menon, “Subgrid combustion modeling of 3-D premixed flames in the thin-reaction-zone regime,” *Proceedings of the Combustion Institute*, vol. 30, no. 1, pp. 575–582, 2005.
- [59] R. Yu, J. Yu, and X.-S. Bai, “An improved high-order scheme for DNS of low Mach number turbulent reacting flows based on stiff chemistry solver,” *Journal of Computational Physics*, vol. 231, no. 16, pp. 5504–5521, 2012.
- [60] C. Lian, G. Xia, and C. L. Merkle, “Solution-limited time stepping to enhance reliability in CFD applications,” *Journal of Computational Physics*, vol. 228, no. 13, pp. 4836–4857, 2009.
- [61] J. M. Weiss, “Calculation of reacting flowfields involving stiff chemistry,” in *14th AIAA Computational Fluid Dynamics Conference*, AIAA Paper 99-3369, 1999.

- [62] F. Génin, “Study of compressible turbulent flows in supersonic environment by large-eddy simulation,” PhD thesis, Georgia Institute of Technology, 2009.
- [63] M. M. Masquelet, “Large-eddy simulations of high-pressure shear coaxial flows relevant for H_2/O_2 rocket engines,” PhD thesis, Georgia Institute of Technology, 2013.
- [64] D. A. Schwer, P. Lu, W. H. Green, and V. Semiao, “A consistent-splitting approach to computing stiff steady-state reacting flows with adaptive chemistry,” *Combustion Theory and Modelling*, vol. 7, no. 2, pp. 383–399, 2003.
- [65] R. L. Speth, W. H. Green, S. MacNamara, and G. Strang, “Balanced splitting and rebalanced splitting,” *SIAM Journal on Numerical Analysis*, vol. 51, no. 6, pp. 3084–3105, 2013.
- [66] G. Strang, “On the construction and comparison of difference schemes,” *SIAM Journal on Numerical Analysis*, vol. 5, no. 3, pp. 506–517, 1968.
- [67] G. Wang, M. Boileau, and D. Veynante, “Implementation of a dynamic thickened flame model for large eddy simulations of turbulent premixed combustion,” *Combustion and Flame*, vol. 158, no. 11, pp. 2199–2213, 2011.
- [68] A Sjunnesson, S Olovsson, and B Sjoblom, “Validation rig- a tool for flame studies,” in *International Symposium on Air Breathing Engines, 10 th, Nottingham, England*, 1991, pp. 385–393.
- [69] A Sjunnesson, C Nelsson, and E Max, “LDA measurements of velocities and turbulence in a bluff body stabilized flame,” *Laser Anemometry*, vol. 3, pp. 83–90, 1991.
- [70] P Nilsson and X. Bai, “Level-set flamelet library approach for premixed turbulent combustion,” *Experimental Thermal and Fluid Science*, vol. 21, no. 13, pp. 87–98, 2000.
- [71] E. Giacomazzi, V. Battaglia, and C. Bruno, “The coupling of turbulence and chemistry in a premixed bluff-body flame as studied by LES,” *Combustion and Flame*, vol. 138, no. 4, pp. 320–335, 2004.
- [72] C. Fureby, “A fractal flame-wrinkling large eddy simulation model for premixed turbulent combustion,” *Proceedings of the Combustion Institute*, vol. 30, no. 1, pp. 593–601, 2005.
- [73] I. Porumbel and S. Menon, “Large eddy simulation of bluff body stabilized premixed flame,” in *44th AIAA Aerospace Sciences Meeting and Exhibit*, AIAA Paper 2006-152, 2006.

- [74] P. A. Cocks, V. Sankaran, and M. C. Soteriou, "Is LES of reacting flows predictive? Part 1: Impact of numerics," in *51st AIAA Aerospace Sciences Meeting, Grapevine TX, January*, AIAA Paper 2013-0170, 2013.
- [75] P. A. Cocks, M. C. Soteriou, and V. Sankaran, "Impact of numerics on the predictive capabilities of reacting flow LES," *Combustion and Flame*, vol. 162, no. 9, pp. 3394–3411, 2015.
- [76] H.-G. Li, P. Khare, H.-G. Sung, and V. Yang, "A large-eddy-simulation study of combustion dynamics of bluff-body stabilized flames," *Combustion Science and Technology*, vol. 188, no. 6, pp. 924–952, 2016.
- [77] N Zettervall, K Nordin-Bates, E. Nilsson, and C Fureby, "Large eddy simulation of a premixed bluff body stabilized flame using global and skeletal reaction mechanisms," *Combustion and Flame*, vol. 179, pp. 1–22, 2017.
- [78] C. Y. Lee and S. Cant, "Large-eddy simulation of a bluff-body stabilized turbulent premixed flame using the transported flame surface density approach," *Combustion Theory and Modelling*, pp. 1–27, 2017.
- [79] A. Comer, M. Ihme, C. Li, S. Menon, J. Oefeline, B. Rankin, V. Sankaran, and V. Sankaran, Eds., *First Model Validation for Propulsion Workshop*, Accessed online 4/25/2017, <https://community.apan.org/wg/afrlccg/mvpws/p/proceedings>, 2017.
- [80] V. Sankaran and T. Gallagher, "Grid convergence in les of bluff body stabilized flames," in *55th AIAA Aerospace Sciences Meeting*, AIAA 2017-1791, 2017.
- [81] P Anselmo-Filho, S Hochgreb, R. Barlow, and R. Cant, "Experimental measurements of geometric properties of turbulent stratified flames," *Proceedings of the Combustion Institute*, vol. 32, no. 2, pp. 1763–1770, 2009.
- [82] R. S. Barlow, G.-H. Wang, P. Anselmo-Filho, M. Sweeney, and S. Hochgreb, "Application of Raman/Rayleigh/LIF diagnostics in turbulent stratified flames," *Proceedings of the Combustion Institute*, vol. 32, no. 1, pp. 945–953, 2009.
- [83] O. Darbyshire, N Swaminathan, and S Hochgreb, "The effects of small-scale mixing models on the prediction of turbulent premixed and stratified combustion," *Combustion Science and Technology*, vol. 182, no. 9, pp. 1141–1170, 2010.
- [84] M. Sweeney, S Hochgreb, and R. Barlow, "The structure of premixed and stratified low turbulence flames," *Combustion and Flame*, vol. 158, no. 5, pp. 935–948, 2011.

- [85] D. G. Goodwin, H. K. Moffat, and R. L. Speth, *Cantera: An object-oriented software toolkit for chemical kinetics, thermodynamics, and transport processes*, <http://www.cantera.org>, Version 1.8.0, 2011.
- [86] M. D. Smooke and V. Giovangigli, “Premixed and nonpremixed test problem results,” in *Reduced kinetic mechanisms and asymptotic approximations for methane-air flames*, Springer, 1991, pp. 29–47.
- [87] F. Bisetti, J.-Y. Chen, J. H. Chen, and E. R. Hawkes, “Differential diffusion effects during the ignition of a thermally stratified premixed hydrogen–air mixture subject to turbulence,” *Proceedings of the Combustion Institute*, vol. 32, no. 1, pp. 1465–1472, 2009.
- [88] P. Sagaut, *Large eddy simulation for incompressible flows: an introduction*. Springer Science & Business Media, 2006.
- [89] F. Génin and S. Menon, “Studies of shock/turbulent shear layer interaction using large-eddy simulation,” *Computers & Fluids*, vol. 39, pp. 800–819, 2010.
- [90] M. Masquelet and S. Menon, “Large-eddy simulation of flame-turbulence interactions in a shear coaxial injector,” *Journal of Propulsion and Power*, vol. 26, no. 5, pp. 924–935, 2010.
- [91] O Colin, F Ducros, D Veynante, and T. Poinso, “A thickened flame model for large eddy simulations of turbulent premixed combustion,” *Physics of Fluids*, vol. 12, no. 7, pp. 1843–1863, 2000.
- [92] F. Charlette, C. Meneveau, and D. Veynante, “A power-law flame wrinkling model for LES of premixed turbulent combustion Part I: Non-dynamic formulation and initial tests,” *Combustion and Flame*, vol. 131, no. 1, pp. 159–180, 2002.
- [93] ———, “A power-law flame wrinkling model for LES of premixed turbulent combustion Part II: Dynamic formulation,” *Combustion and Flame*, vol. 131, no. 1, pp. 181–197, 2002.
- [94] P. A. Strakey and G. Eggenpieler, “Development and validation of a thickened flame modeling approach for large eddy simulation of premixed combustion,” *Journal of Engineering for Gas Turbines and Power*, vol. 132, no. 7, 2010.
- [95] B. G. Franzelli, “Impact of the chemical description on direct numerical simulations and large eddy simulations of turbulent combustion in industrial aero-engines,” PhD thesis, Institut de Mécanique des Fluides de Toulouse, 2011.

- [96] B. Franzelli, E. Riber, and B. Cuenot, “Impact of the chemical description on a Large Eddy Simulation of a lean partially premixed swirled flame,” *Comptes Rendus Mécanique*, vol. 341, no. 1-2, pp. 247–256, 2013.
- [97] M Berglund, E Fedina, C Fureby, J Tegnér, and V Sabel’Nikov, “Finite rate chemistry large-eddy simulation of self-ignition in supersonic combustion ramjet,” *AIAA Journal*, vol. 48, no. 3, pp. 540–550, 2010.
- [98] C. Duwig, K.-J. Nogenmyr, C.-k. Chan, and M. J. Dunn, “Large eddy simulations of a piloted lean premix jet flame using finite-rate chemistry,” *Combustion Theory and Modelling*, vol. 15, no. 4, pp. 537–568, 2011.
- [99] A. G. Smith, “Simulations of vitiated bluff body stabilized flames,” PhD thesis, Georgia Institute of Technology, 2016.
- [100] K. W. Thompson, “Time dependent boundary conditions for hyperbolic systems,” *Journal of Computational Physics*, vol. 68, no. 1, pp. 1–24, 1987.
- [101] K. W. Thompson, “Time dependent boundary conditions for hyperbolic systems, II,” *Journal of Computational Physics*, vol. 89, pp. 439–461, 1990.
- [102] C. Yoo and H. Im, “Characteristic boundary conditions for simulations of compressible reacting flows with multi-dimensional, viscous and reaction effects,” *Combustion Theory and Modelling*, vol. 11, no. 2, pp. 259–286, 2007.
- [103] D. L. Darmofal, P. Moinier, and M. B. Giles, “Eigenmode analysis of boundary conditions for the one-dimensional preconditioned Euler equations,” *Journal of Computational Physics*, vol. 160, no. 1, pp. 369–384, 2000.
- [104] V. Sankaran, C. L. Merkle, X. Zeng, and D. Li, “Influence of large-scale pressure changes on preconditioned solutions at low speeds,” *AIAA Journal*, vol. 42, no. 12, pp. 2490–2498, 2004.
- [105] D. Darmofal and K. Siu, “A robust multigrid algorithm for the Euler equations with local preconditioning and semi-coarsening,” *Journal of Computational Physics*, vol. 151, no. 2, pp. 728–756, 1999.
- [106] D. Darmofal and P.-I. Schmid, “The importance of eigenvectors for local preconditioners of the Euler equations,” in *12th Computational Fluid Dynamics Conference*, AIAA 1996-1655, 1996.
- [107] D. Darmofal and B. van Leer, “Local preconditioning: Manipulating mother nature to fool father time,” in *Computing the Future II: Advances and Prospects in Computational Aerodynamics*, John Wiley and Sons, 1998.

- [108] ———, “Local preconditioning of the Euler equations: A characteristic interpretation,” *Von Karman Institute Lecture Series on Computational Fluid Dynamics*, 1999.
- [109] H.-G. Li, N. Zong, X.-Y. Lu, and V. Yang, “A consistent characteristic boundary condition for general fluid mixture and its implementation in a preconditioning scheme,” *Advances in Applied Mathematics and Mechanics*, vol. 4, no. 01, pp. 72–92, 2012.
- [110] K. Hejranfar and R. Kamali-Moghadam, “Assessment of three preconditioning schemes for solution of the two-dimensional Euler equations at low Mach number flows,” *International Journal for Numerical Methods in Engineering*, vol. 89, no. 1, pp. 20–52, 2012.
- [111] ———, “Preconditioned characteristic boundary conditions for solution of the preconditioned Euler equations at low Mach number flows,” *Journal of Computational Physics*, vol. 231, no. 12, pp. 4384–4402, 2012.
- [112] R. W. MacCormack, “A numerical method for solving the equations of compressible viscous flow,” *AIAA Journal*, vol. 20, no. 9, pp. 1275–1281, 1982.
- [113] P Garcia-Navarro, F Alcrudo, and J. Saviron, “1-D open-channel flow simulation using TVD-MacCormack scheme,” *Journal of Hydraulic Engineering*, vol. 118, no. 10, pp. 1359–1372, 1992.
- [114] C. G. Mingham, D. M. Causon, and D. M. Ingram, “A TVD MacCormack scheme for transcritical flow,” in *Proceedings of the Institution of Civil Engineers-Water and Maritime Engineering*, Thomas Telford Ltd, vol. 148, 2001, pp. 167–175.
- [115] D. Liang, B. Lin, and R. A. Falconer, “Simulation of rapidly varying flow using an efficient TVD–MacCormack scheme,” *International Journal for Numerical Methods in Fluids*, vol. 53, no. 5, pp. 811–826, 2007.
- [116] L. R. Mendez-Nunez and J. J. Carroll, “Comparison of leapfrog, Smolarkiewicz, and MacCormack schemes applied to nonlinear equations,” *Monthly weather review*, vol. 121, no. 2, pp. 565–578, 1993.
- [117] T. P. Gallagher, M. Akiki, S. Menon, and V. Sankaran V.Sankaran, “Development of the generalized MacCormack scheme and its extension to low Mach number flows,” *International Journal for Numerical Methods in Fluids*, n/a–n/a, 2017.
- [118] R. W. Newsome, “Euler and Navier-Stokes solutions for flow over a conical delta wing,” *AIAA Journal*, vol. 24, no. 4, pp. 552–561, 1986.

- [119] R. Warming and B. Hyett, “The modified equation approach to the stability and accuracy analysis of finite-difference methods,” *Journal of computational physics*, vol. 14, no. 2, pp. 159–179, 1974.
- [120] A Fiterman, E Turkel, and V Vatsa, “Pressure updating methods for the steady-state fluid equations,” in *Proceedings of the 12th AIAA Computational Fluid Dynamics Conference*, AIAA Paper 95-36508, 1995.
- [121] M. A. Potsdam, V. Sankaran, and S. A. Pandya, “Unsteady low Mach preconditioning with application to rotorcraft flows,” in *18th AIAA Computational Fluid Dynamics conference*, AIAA Paper, AIAA paper 2007-4473, 2007.
- [122] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, *LAPACK Users’ Guide*, Third. Philadelphia, PA: Society for Industrial and Applied Mathematics, 1999, ISBN: 0-89871-447-8 (paperback).
- [123] U. Ghia, K. N. Ghia, and C. Shin, “High-Re solutions for incompressible flow using the Navier-Stokes equations and a multigrid method,” *Journal of Computational Physics*, vol. 48, no. 3, pp. 387–411, 1982.
- [124] I Demirdžić, Ž Lilek, and M Perić, “A collocated finite volume method for predicting flows at all speeds,” *International Journal for Numerical Methods in Fluids*, vol. 16, no. 12, pp. 1029–1050, 1993.
- [125] S. M. H. Karimian and G. Schneider, “Pressure-based control-volume finite element method for flow at all speeds,” *AIAA Journal*, vol. 33, no. 9, pp. 1611–1618, 1995.
- [126] J. R. Edwards and M.-S. Liou, “Low-diffusion flux-splitting methods for flows at all speeds,” *AIAA Journal*, vol. 36, no. 9, pp. 1610–1617, 1998.
- [127] R. Prosser, “Improved boundary conditions for the direct numerical simulation of turbulent subsonic flows. I. inviscid flows,” *Journal of Computational Physics*, vol. 207, pp. 736–768, 2005.
- [128] S. Huang and Q. Li, “A new dynamic one-equation subgrid-scale model for large eddy simulations,” *International Journal for Numerical Methods in Engineering*, vol. 81, no. 7, pp. 835–865, 2010.
- [129] S. B. Pope, *Turbulent Flows*. Cambridge University Press, 2000.
- [130] G. Xia, V. Sankaran, D. Li, and C. L. Merkle, “Modeling of turbulent mixing layer dynamics in ultra-high pressure flows,” in *36th AIAA Fluid Dynamics Conference and Exhibit*, AIAA Paper 2006-3729, 2006.

- [131] T. Gallagher and S. Menon, “Stiff source term integration in dual-time schemes,” n/a–n/a, In preparation.
- [132] B. Sportisse, “An analysis of operator splitting techniques in the stiff case,” *Journal of Computational Physics*, vol. 161, no. 1, pp. 140–168, 2000.
- [133] P. Sun, “A pseudo-non-time-splitting method in air quality modeling,” *Journal of Computational Physics*, vol. 127, no. 1, pp. 152–157, 1996.
- [134] M. Singer, S. Pope, and H. Najm, “Operator-splitting with ISAT to model reacting flow with detailed chemistry,” *Combustion Theory and Modelling*, vol. 10, no. 2, pp. 199–217, 2006.
- [135] Z. Lu, H. Zhou, S. Li, Z. Ren, T. Lu, and C. K. Law, “Analysis of operator splitting errors for near-limit flame simulations,” *Journal of Computational Physics*, vol. 335, pp. 578–591, 2017.
- [136] W. Hundsdorfer and J. Verwer, “A note on splitting errors for advection-reaction equations,” *Applied Numerical Mathematics*, vol. 18, no. 1-3, pp. 191–199, 1995.
- [137] S. Gupta, H. G. Im, and M. Valorani, “Classification of ignition regimes in HCCI combustion using computational singular perturbation,” *Proceedings of the Combustion Institute*, vol. 33, no. 2, pp. 2991–2999, 2011.
- [138] J. Li, Z. Zhao, A. Kazakov, and F. L. Dryer, “An updated comprehensive kinetic model of hydrogen combustion,” *International Journal of Chemical Kinetics*, vol. 36, no. 10, pp. 566–575, 2004.
- [139] J. H. Chen, E. R. Hawkes, R. Sankaran, S. D. Mason, and H. G. Im, “Direct numerical simulation of ignition front propagation in a constant volume with temperature inhomogeneities: I. fundamental analysis and diagnostics,” *Combustion and Flame*, vol. 145, no. 1, pp. 128–144, 2006.
- [140] E. R. Hawkes, R. Sankaran, P. P. Pébay, and J. H. Chen, “Direct numerical simulation of ignition front propagation in a constant volume with temperature inhomogeneities: II. parametric study,” *Combustion and Flame*, vol. 145, no. 1, pp. 145–159, 2006.
- [141] G. Bansal and H. G. Im, “Autoignition and front propagation in low temperature combustion engine environments,” *Combustion and Flame*, vol. 158, no. 11, pp. 2105–2112, 2011.
- [142] R. Yu and X.-S. Bai, “Direct numerical simulation of lean hydrogen/air auto-ignition in a constant volume enclosure,” *Combustion and Flame*, vol. 160, no. 9, pp. 1706–1716, 2013.

- [143] T. Lu, C. S. Yoo, J. Chen, and C. K. Law, "Three-dimensional direct numerical simulation of a turbulent lifted hydrogen jet flame in heated coflow: A chemical explosive mode analysis," *Journal of Fluid Mechanics*, vol. 652, pp. 45–64, 2010.
- [144] R. Shan, C. S. Yoo, J. H. Chen, and T. Lu, "Computational diagnostics for n-heptane flames with chemical explosive mode analysis," *Combustion and Flame*, vol. 159, no. 10, pp. 3119–3127, 2012.
- [145] G. Bansal and H. Im, "Autoignition of a hydrogen-air mixture with temperature and composition inhomogeneities," in *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*, AIAA Paper 2009-1559, 2009.
- [146] T. P. Coffee and J. M. Heimerl, "Sensitivity analysis for premixed, laminar, steady state flames," *Combustion and Flame*, vol. 50, pp. 323–340, 1983.
- [147] M. Smooke, H Rabitz, Y Reuven, and F. Dryer, "Application of sensitivity analysis to premixed hydrogen-air flames," *Combustion Science and Technology*, vol. 59, no. 4-6, pp. 295–319, 1988.
- [148] C. Westbrook and F. Dryer, "Simplified reaction mechanisms for the oxidation of hydrocarbon fuels in flames," vol. 27, pp. 31–43, 1981.
- [149] F. Génin and S. Menon, "Dynamics of sonic jet injection into supersonic crossflow," *Journal of Turbulence*, vol. 11, no. 4, pp. 1–30, 2010.
- [150] A. Perry and M. Chong, "Topology of flow patterns in vortex motions and turbulence," *Applied Scientific Research*, vol. 53, no. 3, pp. 357–374, 1994.
- [151] H. M. Blackburn, N. N. Mansour, and B. J. Cantwell, "Topology of fine-scale motions in turbulent channel flow," *Journal of Fluid Mechanics*, vol. 310, pp. 269–292, 1996.
- [152] M. Chong, J Soria, A. Perry, J Chacin, B. Cantwell, and Y Na, "Turbulence structures of wall-bounded shear flows found using DNS data," *Journal of Fluid Mechanics*, vol. 357, pp. 225–247, 1998.
- [153] C. B. da Silva and J. C. Pereira, "Invariants of the velocity-gradient, rate-of-strain, and rate-of-rotation tensors across the turbulent/nonturbulent interface in jets," *Physics of Fluids*, vol. 20, no. 5, 2008.
- [154] S. Suman and S. S. Girimaji, "Velocity gradient invariants and local flow-field topology in compressible turbulence," *Journal of Turbulence*, no. 11, N2, 2010.
- [155] L. Wang and X.-Y. Lu, "Flow topology in compressible turbulent boundary layer," *Journal of Fluid Mechanics*, vol. 703, p. 255, 2012.

- [156] L. Cifuentes, C. Dopazo, J. Martin, and C. Jimenez, “Local flow topologies and scalar structures in a turbulent premixed flame,” *Physics of Fluids*, vol. 26, no. 6, p. 065 108, 2014.
- [157] L. Cifuentes, C. Dopazo, J. Martin, P. Domingo, and L. Vervisch, “Effects of the local flow topologies upon the structure of a premixed methane-air turbulent jet flame,” *Flow, Turbulence and Combustion*, vol. 96, no. 2, pp. 535–546, 2016.