

THE APPLICATION OF NON-LINEAR
FREQUENCY DOMAIN METHODS TO THE
EULER AND NAVIER-STOKES EQUATIONS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF AERONAUTICS AND ASTRONAUTICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Matthew Scott McMullen

March 2003

© Copyright by Matthew Scott McMullen 2003
All Rights Reserved

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Antony Jameson
(Principal Adviser)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Juan Alonso

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Robert MacCormack

Approved for the University Committee on Graduate Studies:

Preface

This research demonstrates the accuracy and efficiency of the Non-Linear Frequency Domain (NLFD) method in applications to unsteady flow calculations. The basis of the method is a pseudo-spectral approach to recast a non-linear unsteady system of equations in the temporal domain into a stationary system in the frequency domain. The NLFD method, in principle, provides the rapid convergence of a spectral method with increasing numbers of modes, and, in this sense, it is an optimal scheme for time-periodic problems. In practice it can also be effectively used as a reduced order method in which users deliberately choose not to resolve temporal modes in the solution.

The method is easily applied to problems where the time period of the unsteadiness is known *a priori*. A method is proposed that iteratively calculates the time period when it is not known *a priori*. Convergence acceleration techniques like local time-stepping, implicit residual averaging and multigrid are used in the solution of the frequency-domain equations. A new method, spectral viscosity is also introduced. In conjunction with modifications to the established techniques this produces convergence rates equivalent to state-of-the-art steady-flow solvers.

Two main test cases have been used to evaluate the NLFD method. The first is vortex shedding in low Reynolds number flows past cylinders. Numerical results demonstrate the efficiency of the NLFD method in representing complex flow field physics with a limited number of temporal modes. The shedding frequency is unknown *a priori*, which serves to test the application of the proposed variable-time-period method. The second problem is an airfoil undergoing a forced pitching motion in transonic flow. Comparisons with experimental results demonstrate that a limited

number of temporal modes can accurately represent a non-linear unsteady solution. Comparisons with time-accurate codes also demonstrate the efficiency gains realized by the NLFD method.

Acknowledgments

I am grateful to all the people who supported me during my graduate studies at Stanford. Without their help, this research would have been impossible.

First I would like to thank my adviser Antony Jameson. He has been open-minded to my ideas, patient with my progress, and consistent in his financial support of this research. I would also like to thank my co-adviser Juan Alonso whose mentorship prompted my initial investigations in this field. Not only did Juan and Antony significantly contribute to the scientific content of my research, but they made me feel like I was part of a family while at Stanford.

I would like to extend thanks to William Reynolds, Sanjiva Lele, and George (Bud) Homsy. Each of these individuals offered financial support during my graduate studies and for this I am duly grateful. I would also like to thank Robert MacCormack for his participation in my reading committee.

I owe a special debt of gratitude to a select group of individuals: Matthew Barone, Crisi Benford, Alison Bidwell, John Brennan, Kristin Kilbourne, Siva Nadarajah, James Pearce, Rebecca Thurston, and Stevan Vlaovic. These people reinforced the importance of friendship in my life and provided me with many fond memories of my time at Stanford

Finally, I would like to dedicate this work to my parents. I would be lost without their unconditional support.

Nomenclature

English Letters

A	Jacobian of flux vector F or constant coefficient to Strouhal curve fit
B	Jacobian of flux vector G or inverse coefficient to Strouhal curve fit
b	Component of the control volume face velocity vector
c	Isentropic speed of sound
C	Linear coefficient to Strouhal curve fit
C_d	Coefficient of drag
C_l	Coefficient of lift
C_m	Coefficient of moment
D_x	Discrete operator approximating continuous derivative (x-direction)
D_y	Discrete operator approximating continuous derivative (y-direction)
E	Stagnation enthalpy
E_{N_l}	Error in lift coefficient magnitude
E_{N_m}	Error in moment coefficient magnitude
e	Residual averaging coefficient
F	Component of flux vector contained in spatial residual
f	Function that scales the speed of sound for Riemann invariants
G	Component of flux vector contained in spatial residual
G^m	Gain associated with step m of time advancement scheme
G_P	Gain associated with implicitly evaluated temporal derivative
G_Q	Gain associated with explicitly evaluated temporal derivative
G_R	Gain associated with spatial residual
F_c	Convective part of the flux vector

F_v	Viscous part of the flux vector
\hat{I}	Fourier coefficient of unsteady residual
J	Cell volume
k	Wavenumber
L_α	Lift transfer function
M_∞	Farfield Mach number
M_α	Moment transfer function
N	Surface normal
N_M	Cost associated with the number of multigrid cycles required for a converged solution
N_N	Cost associated with the number of solution instances required to resolve a solution oscillation
N_P	Cost associated with the number of periods required to reach a periodic steady state
n	Integer wavenumber
$\hat{P}_k^{\overline{m}}$	Temporal derivative used at step m of implicit time advancement
p	Pressure
Pr	Prandtl number
q	Heat flux vector
Q	Pressure switch for artificial dissipation scheme
$\hat{Q}_k^{\overline{m}}$	Temporal derivative used at step m of explicit time advancement
R	Spatial residual or normal to cell face orthogonal to boundary wall
R_c	Convective part spatial residual
R_v	Viscous part spatial residual
\hat{R}	Fourier coefficient of spatial residual
$\hat{R}_k^{\overline{m}}$	Spatial operator used at step m of time advancement scheme
Re	Reynolds number
S	Surface area or source term representing temporal derivative
S_t	Strouhal number
t	Time
Δt	Time step
T	Time period of the unsteady solution
T_s	Time required to time accurately drive a solution to its periodic steady state
U	Fluid velocity vector

u	Component of the fluid velocity vector
V	Volume
W	Solution
\bar{W}	Steady component of solution
\tilde{W}	Unsteady component of solution
\hat{W}	Fourier coefficient of solution
X	Position vector
x	Component of the position vector
y	Component of the position vector

Greek and Roman Symbols

α	Angle of attack or coefficient multiplying overall unsteady response
α_m	Mean angle of attack
α_0	Magnitude of dynamic angle of attack
δ_{ij}	Kronecker delta
ρ	Density
σ	Shear stress tensor
γ	Ratio of specific heats
Λ	Steady state CFL
λ	Eigenvalue of spatial discretization
λ_t	Eigenvalue of temporal discretization
λ_c	Convective component of eigenvalue
λ_v	Viscous component of eigenvalue
μ	Absolute viscosity
τ	Pseudo time
κ	Heat transfer coefficient
$\varepsilon^{(2)}$	Coefficient controlling first order dissipation
$\varepsilon^{(4)}$	Coefficient controlling third order dissipation
ε_N	Coefficient that scales course grid spectral viscosity
ξ	Coordinate direction on uniform grid orthogonal to η
η	Coordinate direction on uniform grid ξ

Subscripts and Superscripts

$()^*$	Dimensionless quantity
$()_\infty$	Characteristic quantity

Contents

Preface	iv
Acknowledgments	vi
Nomenclature	vii
1 Introduction	1
1.1 Motivation	1
1.2 Previous Work	5
1.2.1 Time-Accurate Methods	5
1.2.2 Periodic Methods	12
1.3 Approach	19
1.4 Overview	20
2 Methodology	22
2.1 Governing Equations	22
2.2 Numerical Method	24
2.2.1 Transforming the Equations into the Frequency Domain . . .	24
2.2.2 Pseudo-Time Derivative	27
2.2.3 Temporal Resolution	28
2.2.4 Variable Time Period	33
2.3 Convergence Acceleration of the Pseudo-Time Equation	36
2.3.1 Local Time Stepping	37
2.3.2 Residual Averaging	44

2.3.3	Multigrid	48
2.3.4	Coarse Grid Spectral Viscosity	51
2.3.5	Parallelization	54
3	Cylindrical Vortex Shedding	56
3.1	Test Case Description	56
3.2	Computational Grid	58
3.3	Cylinder Results	59
3.3.1	Variable-Time-Period Results	61
3.3.2	Fixed-Time-Period Results	65
3.4	Summary	66
4	Pitching Airfoil Experimental Validation	77
4.1	Introduction	77
4.2	Test Case Description	79
4.3	Computational Grid	81
4.4	Steady Inviscid Flow	87
4.5	Steady Viscous Flow	92
4.6	Unsteady Inviscid Flow	95
4.7	Unsteady Viscous Flow	108
4.8	Summary	111
5	Pitching Airfoil Numerical Validation	120
5.1	Choice of a Test Case	120
5.2	Computational Grid	121
5.3	Comparison With UFLO82 For Steady Flow	121
5.4	Comparison With UFLO82 For Unsteady Flow	124
5.5	Sensitivity of the NLFD Solver To Flow Conditions	140
5.5.1	Convergence Versus Reduced Frequency	140
5.5.2	Convergence Versus Dynamic Angle of Attack	141
5.6	Summary	141

6	Conclusions	145
6.1	Solver Implementation	145
6.2	Cylindrical Vortex Shedding	146
6.3	Pitching Airfoil	147
6.4	Computational Efficiency of the NLFD Method	148
6.5	Future Work	148
A	Discretization	150
A.1	Non-Dimensionalization	150
A.2	Spatial Operators	152
A.2.1	Convective Fluxes	152
A.2.2	Artificial-Dissipation Fluxes	152
A.2.3	Viscous Fluxes	154
A.2.4	Turbulence Model	157
A.3	Boundary Conditions	158
A.3.1	Wall	158
A.3.2	Farfield	162
B	Unsteady Channel Flow	166
B.1	Small Perturbation Theory	166
B.2	Channel Flow Results	169
B.3	Channel Flow Convergence	171
	Bibliography	174

List of Tables

1.1	Computational cost estimates for a multistage compressor and turbine calculation. Execution time estimates based on 750 processors running in parallel approximately 8 hours a day.	2
1.2	Results from a number of different runs produced by Mitchell’s direct numerical simulation of subsonic axisymmetric jets. The final column provides the ratio of time required to capture one oscillation for the lowest frequency in the solution (T) to the time required to reach a periodic steady state (T_s).	3
1.3	The error and stability characteristics of a number of explicit multistep methods. Note: $ip = \lambda \Delta t$. Source: Pulliam [43]	18
3.1	Pertinent features of the grids used in the cylinder simulations. The number of points in the boundary layer was calculated using the velocity profiles at the 12 o’clock position on the cylinder at a Reynolds number of 150.	59
3.2	Values of parameters used in the three-dimensional parametric survey, which combines both temporal and spatial resolution surveys over a range in Reynolds number.	61
3.3	Shift in Reynolds number between the NLFD calculations and the experimental dataset for all permutations of temporal and spatial resolution.	63
4.1	Hierarchy of governing equations used in the analysis of unsteady transonic flow over airfoils. Source: McCroskey [46].	78

4.2	Shock movement as a percentage of the airfoil chord for the different experimental test cases used in chapter 4.	79
4.3	Description of priority test cases used from AGARD Report 702. . . .	80
4.4	Description of test cases used within Green and Newman's steady experiment.	81
4.5	Description of meshes employed for the NLFD Euler calculations. . .	82
4.6	Description of meshes employed for the NLFD Navier-Stokes calculations.	82
4.7	Steady lift curve slope based on angles of attack less than 3° for all numerical and experimental results.	91
4.8	Steady lift curve slope based on angles of attack less than 3° for Navier-Stokes calculations (Reynolds number of 3.0 million) and experimental results.	93
4.9	Magnitude of zeroth harmonic in coefficient of lift divided by the magnitude of the zeroth harmonic in angle of attack $\left(\frac{\ \hat{C}_{l_0}\ }{\ \hat{\alpha}_0\ }\right)$ for various temporal resolutions employed by NLFD simulations on the 0012 airfoil.	97
4.10	Magnitude of fundamental harmonic in coefficient of lift divided by the magnitude of the fundamental harmonic in angle of attack $\left(\frac{\ \hat{C}_{l_1}\ }{\ \hat{\alpha}_1\ }\right)$ for various temporal resolutions employed by NLFD simulations.	97
4.11	Phase lag between the fundamental harmonic in lift coefficient and angle of attack $\left(\angle \hat{C}_{l_1} - \angle \hat{\alpha}_1\right)$ for various temporal resolutions employed by NLFD simulations on the 0012 airfoil.	98
4.12	Magnitude of zeroth harmonic in coefficient of lift divided by the magnitude of the zeroth harmonic in angle of attack for various temporal resolutions employed by the NLFD method $\left(\frac{\ \hat{C}_{l_0}\ }{\ \hat{\alpha}_0\ }\right)$	109
4.13	Magnitude of fundamental harmonic in coefficient of lift divided by the magnitude of the fundamental harmonic in angle of attack for various temporal resolutions employed by the NLFD method $\left(\frac{\ \hat{C}_{l_1}\ }{\ \hat{\alpha}_1\ }\right)$	109
4.14	Magnitude of phase shift between the fundamental harmonic in lift coefficient and angle of attack for various temporal resolutions employed by the NLFD method $\left(\angle \hat{C}_{l_1} - \angle \hat{\alpha}_1\right)$	110

5.1	Description of the test case used in the comparison of UFLO82 and the NLFD codes.	121
5.2	Cost estimation for dual time stepping codes like UFLO82.	125
5.3	Cost estimation for frequency domain codes like NLFD.	125
5.4	The number of solutions per wavelength required by UFLO82 to reach error levels equivalent to results produced by an NLFD code using one, two, and three time varying harmonics.	127
5.5	Variation in UFLO82 unsteady coefficient of lift as a function of temporal resolution.	128
5.6	The number of multigrid cycles required by the NLFD solver to obtain a solution at equivalent error levels.	129
5.7	NLFD execution time per multigrid cycle as a function of temporal resolution. Results are normalized by the execution time required by the steady state or zeroth mode case.	129
5.8	The number of time periods required by the UFLO82 solver to reach convergence. Separate data is provided for each temporal resolution listed as SPW or Solutions Per Wavelength.	130
5.9	Cost comparison between UFLO82 and NLFD codes using the error in coefficient of lift as the figure of merit.	131
5.10	Cost comparison between UFLO82 and NLFD codes using the error in coefficient of moment as the figure of merit.	131
B.1	Boundary conditions used for test case.	170
B.2	The number of multigrid cycles required to achieve machine zero convergence. Two separate cases are provided at different values of $\frac{T\bar{u}}{\Delta x}$. Within each case, the initial transients decay rate was varied by ten orders of magnitude.	173

List of Figures

1.1	A plot of mass flux as a function of time for the TFLO simulation of the Aachen experimental test rig.	4
1.2	Perturbed velocity at any point along the unsteady channel with a reduced frequency $\Omega = 50$	12
2.1	Simplified dataflow diagram of the time advancement scheme illustrating the pseudo spectral approach used in calculating the non-linear spatial operator R	26
2.2	Modified wavenumber analysis for the backward difference formula employed in UFLO82. By definition $N = \frac{2\pi}{k\Delta t}$. (a) Real component. (b) Imaginary component	30
2.3	(a) The minimum residual obtained in a fixed time period NLFD calculation as a function of the distance between Strouhal number specified <i>a priori</i> . The data is plotted using a logarithmic scaling for the vertical axis only. (b) The minimum residual obtained in a fixed time period NLFD calculation as a function of the distance between Strouhal number specified <i>a priori</i> and the Strouhal number produced from a GBVTP calculation. The data is plotted using a logarithmic scaling for both horizontal and vertical axes.	36
2.4	Complete dataflow diagram of the time-stepping method.	42
2.5	Stability diagrams of the magnitude of the overall time-stepping scheme in the complex $\lambda\Delta t$ plane. The coefficients for all above permutations: $\alpha = [\frac{1}{4}, \frac{1}{6}, \frac{3}{8}, \frac{1}{2}, 1]$ $\alpha_R = [1, 0, 0.56, 0.44]$	43

2.6	Residual averaging coefficient as a function of $k\Delta\tau_k$. These roots are calculated using a pseudo time step calculated based on the $CFL = 6$ and a rescaled eigenvalue $\lambda = 2$	47
2.7	System of grids used inside of a multigrid cycle. E, evaluate the change in the flow for one step; T, transfer the data without updating the solution. Source: Jameson [37].	49
2.8	Pictorial of multigrid restriction operator.	50
2.9	Pictorial of multigrid prolongation operator.	51
2.10	(a) Stability diagram of the explicit time-stepping scheme. The red and blue ellipses are the spectral footprint of the discretization with and without spectral viscosity respectively. (b) The gain of the time-stepping scheme along the paths of the spectral footprint.	53
2.11	Magnitude of all components of the unsteady density residual as a function of the multigrid cycle. (a) NLFD solver with coarse grid spectral viscosity (b) NLFD solver without coarse grid spectral viscosity.	54
2.12	Magnitude of the difference between solutions computed with and without spectral viscosity shown at time $t = \frac{6}{7}T$. (a) ρ (b) ρu (c) ρv (d) ρe	55
3.1	Smoke visualization of laminar vortex shedding in the wake of a cylinder at a Reynolds number of 140. Experiment conducted by Taneda [63] and reproduced by Van Dyke [17].	57
3.2	Top view of cylinder wake experiments using flow visualization techniques to identify parallel and oblique shedding modes. (a) Parallel mode. (b) Oblique modes. Source: Williamson [68].	58
3.3	Contours of entropy at a Reynolds number of 150 utilizing the 129x65 grid. (a) solution using 1 time varying harmonic. (b) solution using 7 time varying harmonics.	68
3.4	Contours of entropy at a Reynolds number of 150 utilizing the 385x161 grid. (a) solution using 1 time varying harmonic. (b) solution using 7 time varying harmonics.	69

3.5	Strouhal number as a function of Reynolds number plotted for different temporal resolutions. (a) 129x65 grid (b) 193x81 grid (c) 257x129 grid (d) 385x161 grid.	70
3.6	Strouhal number as a function of the spatial resolution plotted for different temporal resolutions. (a) $R_e = 60$ (b) $R_e = 90$ (c) $R_e = 120$ (d) $R_e = 150$	71
3.7	Strouhal number as a function of the temporal resolution plotted for different grid sizes. (a) $R_e = 60$ (b) $R_e = 90$ (c) $R_e = 120$ (d) $R_e = 150$	72
3.8	L^2 norm of the product of density and energy for each wavenumber as a function of Reynolds number. (a) 129x65 grid (b) 193x81 grid (c) 257x129 grid (d) 385x161 grid.	73
3.9	Base suction coefficient as a function of the Reynolds number plotted for different temporal resolutions. (a) 129x65 grid (b) 193x81 grid (c) 257x129 grid (d) 385x161 grid.	74
3.10	Base suction coefficient as a function of the spatial resolution plotted for different temporal resolutions. (a) $R_e = 60$ (b) $R_e = 90$ (c) $R_e = 120$ (d) $R_e = 150$	75
3.11	Base suction coefficient as a function of the temporal resolution plotted for different grid sizes. (a) $R_e = 60$ (b) $R_e = 90$ (c) $R_e = 120$ (d) $R_e = 150$	76
4.1	Nearfield resolution of 0012 grids used in the Euler calculations. (a) 161x33 points in O-mesh topology (b) 321x65 points in O-mesh topology (c) 257x33 points in C-mesh topology (d) 321x97 points in C-mesh topology.	83
4.2	Nearfield resolution of the 64A010 grids used in the unsteady Euler calculations. (a) 81x33 points in O-mesh topology (b) 161x33 points in O-mesh topology (c) 129x33 points in C-mesh topology (d) 193x49 points in C-mesh topology.	84
4.3	Nearfield resolution of the 0012 Navier-Stokes grids. (a) 257x49 points in C-mesh topology (b) 513x97 points in C-mesh topology.	85

4.4	Nearfield resolution of the 64A010 grids used in the unsteady Navier-Stokes calculations. (a) 129x33 points in O-mesh topology (b) 193x49 points in O-mesh topology (c) 257x65 points in C-mesh topology. . .	86
4.5	Lift coefficient as a function of angle of attack at a free-stream Mach number of 0.7 including both Euler NLFD and experimental data. (a) Experimental Reynolds number equal to 3.0 million. (b) Experimental Reynolds number equal to 9.0 million.	88
4.6	Drag coefficient as a function of angle of attack at a free-stream Mach number of 0.7 including both Euler NLFD and experimental data. (a) Experimental Reynolds number equal to 3.0 million. (b) Experimental Reynolds number equal to 9.0 million.	89
4.7	Moment coefficient as a function of angle of attack at a free-stream Mach number of 0.7 including both Euler NLFD and experimental data. (a) Experimental Reynolds number equal to 3.0 million. (b) Experimental Reynolds number equal to 9.0 million.	90
4.8	Force coefficient as a function of angle of attack at a free-stream Mach number of 0.7 and a Reynolds number of 3.0 million including both Navier-Stokes NLFD and experimental data. (a) Coefficient of lift. (b) Coefficient of drag. (c) Coefficient of moment.	94
4.9	Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Landon's 0012 experiment. (a) 161x33 O-mesh. (b) 321x65 O-mesh. (c) 257x33 C-mesh. (d) 321x97 C-mesh.	100
4.10	Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Davis's 64A010 experiment. (a) 81x33 O-mesh. (b) 161x33 O-mesh. (c) 129x33 C-mesh. (d) 193x49 C-mesh.	101
4.11	Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Landon's 0012 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics. . . .	102

4.12	Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Davis's 64A010 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics.	103
4.13	Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Landon's 0012 experiment. (a) 161x33 O-mesh. (b) 321x65 O-mesh. (c) 257x33 C-mesh. (d) 321x97 C-mesh.	104
4.14	Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Davis's 64A010 experiment. (a) 81x33 O-mesh. (b) 161x33 O-mesh. (c) 129x33 C-mesh. (d) 193x49 C-mesh.	105
4.15	Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Landon's 0012 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics. . . .	106
4.16	Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Davis's 64A010 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics. . . .	107
4.17	Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Landon's 0012 experiment. (a) 257x49 C-mesh. (b) 513x97 C-mesh.	111
4.18	Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Davis's 64A010 experiment. (a) 129x33 C-mesh. (b) 193x49 C-mesh. (c) 257x65 C-mesh.	113
4.19	Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Landon's 0012 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics. . . .	114
4.20	Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Davis's 64A010 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics.	115

4.21	Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Landon's 0012 experiment. (a) 257x49 C-mesh. (b) 513x97 C-mesh. . . .	116
4.22	Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Davis's 64A010 experiment. (a) 257x49 C-mesh. (b) 513x97 C-mesh. (c) 257x65 C-mesh.	117
4.23	Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Landon's 0012 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics.	118
4.24	Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Davis's 64A010 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics.	119
5.1	Nearfield resolution of 161x33 O-mesh used in the numerical comparison of UFLO82 and NLFD codes.	122
5.2	The absolute value of the difference between a UFLO82 and NLFD steady solution for Euler calculation on a 161x33 O-mesh. The different figures represent different components of the solution. (a) ρ (b) ρu (c) ρv (d) ρe	123
5.3	C_{l_1} error as defined by equation 5.2 between various temporal resolutions of the UFLO82 code and the control solution. The control solution is based on the coefficient of lift extracted from the final period of the most accurate solution (512 solutions per wavelength). (a) Entire time history of all data. (b) Magnified version of the same data focusing on less accurate solutions.	132

5.4	C_{m_1} error as defined by equation 5.2 between various temporal resolutions of the UFLO82 code and the control solution. The control solution is based on the coefficient of moment extracted from the final period of the most accurate solution (512 solutions per wavelength). (a) Entire time history of all data. (b) Magnified version of the same data focusing on less accurate solutions.	133
5.5	Asymptotic error in C_{l_1} computed at various temporal resolutions by the UFLO82 and NLFD codes. (a) C_{l_1} Error as a function of Δt on a fully logarithmic axis. (b) C_{l_1} Error as a function of the number of solutions used to resolve an oscillation in the solution using a logarithmic vertical scale.	134
5.6	Asymptotic error in C_{m_1} computed at various temporal resolutions by the UFLO82 and NLFD codes. (a) C_{m_1} Error as a function of Δt on a fully logarithmic axis. (b) C_{m_1} Error as a function of the number of solutions used to resolve an oscillation in the solution using a logarithmic vertical scale.	135
5.7	The number of multigrid cycles required by the UFLO82 code to reach equivalent error levels. (a) C_{l_1} . (b) C_{m_1}	136
5.8	The number of multigrid cycles required by the NLFD code to reach equivalent error levels. (a) C_{l_1} . (b) C_{m_1}	137
5.9	C_l error as a function of physical time in units of periods. (a) 45 Steps per wavelength (b) 125 Steps per wavelength (c) 244 Steps per wavelength.	138
5.10	C_m error as a function of physical time in units of periods (a) 18 Steps per wavelength (b) 45 Steps per wavelength (c) 123 Steps per wavelength.	139
5.11	The effects of reduced frequency on the convergence of the NLFD unsteady residual. (a) The maximum over all the wavenumbers of the absolute value of the residual as a function of the multigrid cycle. (b) The residual convergence rate per multigrid cycle for the first 140 cycles.	143

5.12	The effects of dynamic angle of attack on the convergence of the NLFD unsteady residual. (a) The maximum over all the wavenumbers of the absolute value of the residual as a function of the multigrid cycle. (b) The residual convergence rate per multigrid cycle for the first 140 cycles.	144
A.1	Convective Flux Control Volume	153
A.2	Viscous Flux Control Volume	155
A.3	Pictorial of cells at the wall for a curvilinear mesh.	159
A.4	Pictorial of cells in the farfield for a curvilinear mesh.	164
B.1	Pictorial of one-dimensional channel flow model problem.	167
B.2	Three-dimensional plot of channel velocity as a function of streamwise coordinate and time.	171
B.3	Three-dimensional plot of channel pressure as a function of streamwise coordinate and time.	172
B.4	The infinity norm of the difference between analytic and NLFD solutions as a function of exhaust pressure perturbation.	173
B.5	Absolute value of the residual as a function of the multigrid cycle for a number of exhaust pressure perturbations.	173

Chapter 1

Introduction

1.1 Motivation

The calculation of unsteady flows continues to present a severe challenge to Computational Fluid Dynamics (CFD). While preserving the accurate spatial discretizations associated with established steady-state solvers, unsteady codes also need to accurately resolve the time history of the solution. The challenge to CFD lies in this added dimension and its associated computational cost. The motivation behind this research has been to develop new numerical techniques that can mitigate this cost while maintaining a high quality representation of the solution.

Estimates from a turbomachinery calculation illustrate the extreme cost of computing a complex unsteady solution. An example is a proposed high fidelity simulation of both a complete turbine and compressor to be completed as part of Stanford University's contribution to the Accelerated Strategic Computing Initiative (ASCI). The unsteadiness in the solution stems from the relative motion of the rotating and fixed blades, and has a fundamental period which depends on the rate of rotation and the number of blade passages. Table 1.1 provides estimates for total processor utilization, and the overall time required to complete the job using one of the massively parallel computers owned by the Department of Energy [12]. It is obvious from the estimates that units of months and years are needed to quantify the amount of time required to complete these jobs. Given the magnitude of these times, unsteady CFD cannot

be appropriately used in an industrial setting as a design tool for applications of this type.

Component	Blade Rows	Grid Points (millions)	Percent Wheel	CPU Hours (millions)	Execution Duration (days)
Turbine	9	94	16	3.0	500
Compressor	23	NA	16	7.7	1300

Table 1.1: Computational cost estimates for a multistage compressor and turbine calculation. Execution time estimates based on 750 processors running in parallel approximately 8 hours a day.

The above example is representative of a variety of industrial applications for which the solution exhibits unsteady periodic behavior. The initial conditions determine the evolution of the solution over time. Typically, the initial transients decay until the solution reaches a periodic steady state where the solution repeats over a given time period. It is difficult to define precise mathematical criteria to determine whether a solution has reached a periodic state. In general, its statistics (mean and higher order moments integrated over time) will no longer be a function of time. The rate at which the calculation evolves from its initial solution to a periodic steady state is a function of the flow field physics, and for most complex flows is indeterminable *a priori*.

Many problems have multiple temporal scales associated with their physics. In the worst case the time scale of initial decay is much larger than the time scale of the periodic solution. In this situation, the decay of the initial transients represents the dominant computational cost in comparison with a single oscillation of the solution at its periodic steady state. For example, Mitchell [53] performed a direct numerical simulation of the acoustic field produced by an axisymmetric jet. Given an initial axial velocity profile, the shear layer of the jet was excited via small amplitude disturbances in the inflow boundary conditions. The shear layer rolled up into vortices which may or may not pair with one another as they propagate downstream. A fourth order Runge-Kutta scheme was used to calculate the evolution to a periodic steady state.

The ratio of the amount of time required to compute a single oscillation in the flow field versus the amount of time required to reach a periodic steady state is shown in table 1.2. The table presents the results of three different simulations performed at varying Mach numbers and shear layer widths. Given that the user only required a solution for a periodic steady state, it is evident that roughly 95 percent of the computational time was used to generate irrelevant results.

Description	Mach Number	$\frac{T}{T_s}$
Thick Shear Layer Single Pairing	0.4	5.8 %
Thick Shear Layer Single Pairing	0.8	3.6 %
Thin Shear Layer Dual Pairing	0.8	3.6 %

Table 1.2: Results from a number of different runs produced by Mitchell’s direct numerical simulation of subsonic axisymmetric jets. The final column provides the ratio of time required to capture one oscillation for the lowest frequency in the solution (T) to the time required to reach a periodic steady state (T_s).

Another example of a physical problem with multiple time scales is the numerical simulation of the flow through an axial flow turbine. Such a calculation has been performed by Yao *et al.* [77] as part of a verification process for the ASCI TFLO solver. The code performed an unsteady Reynolds Averaged Navier-Stokes (RANS) calculation on a $1\frac{1}{2}$ stage turbine modeled after an experimental test rig. Figure 1.1 shows the evolution of the mass flux into the machine during the process of the computation. A periodic steady state was reached after roughly 2500 time steps. The fundamental period of the oscillation of the solution at its converged state is resolved in about 80 time steps, a ratio is approximately 3.2 percent, which is similar to the results for the shear layer.

Unsteady flows can be divided into two main categories. The first category includes flows where the resolution of the initial transients is relevant, and accurate initial conditions are required. In the second category the user requires only the solution once it has reached a periodic steady state. As in the case of a steady flow




Figure 1.1: A plot of mass flux as a function of time for the TFLO simulation of the Aachen experimental test rig.

simulation, the final periodic state is independent of the initial conditions. Commonly the initial conditions are arbitrary guesses, although they may be derived by a lower order method. If a time-accurate solver is used in this case, the simulation of the evolution to a periodic steady state may be a function of non-physical initial conditions, and consequently deemed irrelevant.

The motivation of this research is to improve solver technologies for the second class of problems. There exist many physical problems within this class where boundary conditions force the unsteadiness at predetermined frequencies. Examples of this include the internal flows in turbomachinery, the external flow fields of helicopter blades or propellers, and certain aeroelastic computations. A second category exists where the unsteadiness is induced by instability waves within the flow field. The resulting periodic flow is the result of boundary conditions, but the unsteadiness is not forced at any predetermined frequency. Examples in this class include (but are obviously not limited to) vortex shedding behind a cylinder and other fluid dynamic cases involving separated flows and free shear layers. Estimates of the frequencies of the unsteadiness can be gained via experimental data or simplified analytic models.

However, the exact temporal frequencies for a given spatial discretization are usually impossible to determine *a priori*. The research presented in this thesis focuses on methods that are amenable to both situations.

1.2 Previous Work

Topics covered in the following sections include explicit, implicit, factored ADI, linearized frequency domain, and non-linear frequency domain schemes. While a variety of different governing equations may be used to model fluid mechanics (from linear potential methods to fully non-linear direct Navier-Stokes calculations), this work focuses only on solutions to the unsteady Euler and Navier-Stokes equations. Citations to investigators employing these schemes will originate from fields of study inside of unsteady computational fluid dynamics which include: aeroacoustics, aeroelasticity, turbomachinery and numerical investigations of turbulence.

1.2.1 Time-Accurate Methods

Most schemes, with the notable exception of the Lax-Wendroff scheme, can be constructed by discretizing in space while leaving the time term continuous to produce a system of coupled Ordinary Differential Equations (ODE) [22]. This is known as the semi-discrete form of the equations, and advancing this system forward in time is also known as the method of lines. Equation 1.1 shows the semi-discrete form of an arbitrary system of conservation laws in two spatial dimensions. The D_x and D_y operators are discrete approximations to the continuous derivatives operators $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$ respectively. The vectors F and G represent functions of the solution variable W constructed to conserve a property of the fluid (mass, momentum, energy etc)

$$\begin{aligned}\frac{dW}{dt} &= -R(W) \\ R(W) &= D_x F(W) + D_y G(W).\end{aligned}\tag{1.1}$$

The following sections focus on previous work where numerical schemes for this

approach have been applied to unsteady problems in fluid dynamics. Global time accuracy requires that approximation of the temporal derivative is at least second order accurate in the time step $\mathcal{O}(\Delta t^2)$.

Explicit Time-Stepping Schemes

The most straightforward approach to solving a system of ODEs via the method of lines is to employ an explicit time integration scheme, where the time derivative of the ODE is evaluated using known information. Explicit schemes can be classified into two basic types: multistep methods and predictor-corrector schemes. Two-step multistep methods can be generally described as:

$$(1 + \xi)W^{n+1} = [(1 + 2\xi)W^n - \xi W^{n-1}] + \Delta t[(1 + \varphi)\frac{dW^n}{dt} - \varphi\frac{dW^{n-1}}{dt}], \quad (1.2)$$

where the solution at the n^{th} time step is denoted as W^n . This equation can be extended to multistep methods using more than two steps. Similarly, predictor-corrector schemes (including Runge-Kutta) can be generally described by the following system of equations:

$$\begin{aligned} W^\beta &= W^n + \gamma t \alpha_1 \frac{dW^n}{dt} \\ W^\gamma &= W^n + \gamma t \left(\alpha_2 \frac{dW^n}{dt} + \beta_2 \frac{dW^\beta}{dt} \right) \\ W^\delta &= W^n + \gamma t \left(\alpha_3 \frac{dW^n}{dt} + \beta_3 \frac{dW^\beta}{dt} + \gamma_3 \frac{dW^\gamma}{dt} \right) \\ W^{n+1} &= W^n + \gamma t \left(\alpha_4 \frac{dW^n}{dt} + \beta_4 \frac{dW^\beta}{dt} + \gamma_4 \frac{dW^\gamma}{dt} + \delta_4 \frac{dW^\delta}{dt} \right). \end{aligned} \quad (1.3)$$

These schemes use the solution W^n at the n^{th} time step and non-physical solutions at fractional time steps denoted as W^β , W^γ and W^δ . Because of the similarity of the methods, this section will address issues associated with both of these groups.

The error and stability characteristics for a number of methods contained in both categories are provided in table 1.2.1. The second and third columns in this table

provide the order of the solution error in terms of the linearized form of the semi-discrete equation provided as:

$$\frac{dW}{dt} = \lambda W. \quad (1.4)$$

These columns define the relationship between the spectral radius, λ , and the solution error components in terms of their magnitude and phase contributions for the model convection problem

Using columns two and three, and following the arguments of Pulliam *et al.* [43], estimates for the efficiency of each scheme can be derived. Columns four and five provide the numerical error produced by a simulation which advances a linear system forward a fixed amount of time. The amount of work consumed by the different simulations is held constant. This is achieved by altering the time step Δt for each scheme such that the total number of residual evaluations remains fixed. The scheme with the least amount of error can be considered the most efficient; within this group the most efficient method is arguably the fourth order Runge-Kutta scheme. The final column provides a plot of the stability region of the scheme in the complex $\lambda \Delta t$ plane. The gray sections of these plots show the values of $\lambda \Delta t$ for which the numerical scheme will remain stable.

Given a rectangular cell (with width Δx and height Δy), an estimate of the spectral radius produced from a second order symmetric discretization of the Navier-Stokes equations is provided in equation 1.5. The variables u and v denote the components of the velocity vector in the x and y directions respectively, while μ is the absolute viscosity of the fluid

$$\lambda = \frac{|u| + c}{\Delta x} + \frac{|v| + c}{\Delta y} + \frac{4\mu}{\Delta x^2} + \frac{4\mu}{\Delta y^2} + \frac{2}{3} \frac{\mu}{\Delta x \Delta y}. \quad (1.5)$$

For time-accurate schemes the time step for all the cells in the mesh is the minimum of the time step calculated at each cell. For viscous calculations at large Reynolds numbers, the mesh needs to be significantly refined in the boundary/shear layers to resolve the gradients in the solution. The dominant contribution in the calculation of the spectral radius becomes the $\frac{1}{\Delta y^2}$ term which typically determines

the minimum time step. Given the stability regions of even the most efficient explicit schemes, a large number of time steps and work will be required to resolve a single oscillation in the solution. As the Reynolds number increases the grid refinement increases and the time step decreases. Eventually any further increases in the Reynolds number of the simulation will require more computational effort than what is available via modern computer hardware.

Successful calculations have been performed using explicit techniques for very low Reynolds number flows. For example, Kwak *et al.* [41] used a second order Adams-Bashforth method to perform unsteady computations of periodic turbulent flow matching experiments generating mesh turbulence at a Reynolds number of 34,000. More recently, Manning and Lele [44] performed unsteady calculations of a moving shock interacting with a shear layer. They used the third-order accurate RK scheme of Wray [76] to advance the viscous equations at a Reynolds number of 1000 based on the initial vorticity thickness. However, these calculations may represent the upper bound in the attainable Reynolds number for explicit schemes. This conclusion is reinforced by a study of flow separation behind a backward facing step performed at similar Reynolds numbers by Akselvoll and Moin [2]. Prior to the simulation, Akselvoll calculated the allowable time steps for a similar flow to demonstrate the dominance of the viscous terms on stability. Even at these low Reynolds numbers, Akselvoll adopted an implicit implementation of these terms to mitigate their restrictions on stability.

Implicit Time-Stepping Schemes

Implicit methods are more difficult to implement because the time derivatives of the discretization must be evaluated using future or unknown values of the solution. The benefit of these schemes is typically an expansion of the stability region, ultimately to the extent that they can be stable for arbitrarily large time steps whenever the true solution is non-increasing (A-stability). In contrast to the limited stability regions of the explicit schemes, A-stable methods are stable over the entire left half of the complex $\lambda\Delta t$ plane. The disadvantage of these schemes is that they require the solution of a system of (possibly non-linear) equations at each time step.

A popular implicit time-accurate scheme, the Crank-Nicholson method (also known as the "One-Step Trapezoidal" method), is defined as follows:

$$W^{n+1} = W^n + \frac{\Delta t}{2} \left[\frac{dW^{n+1}}{dt} + \frac{dW^n}{dt} \right]. \quad (1.6)$$

Following Briley and McDonald [10] or Beam and Warming [7] one can linearize the non-linear residual term using a Taylor series expansion as provided below:

$$R(W^{n+1}) = R(W^n) + \frac{\partial R(W^n)}{\partial W} (W^{n+1} - W^n) + \mathcal{O}(\Delta t^2). \quad (1.7)$$

Using equations 1.6, 1.7 and 1.1, one can write the Crank-Nicholson scheme as applied to a multidimensional system of equations in delta form:

$$\left[I + \frac{\Delta t}{2} \left(D_x \frac{\partial F}{\partial W} + D_y \frac{\partial G}{\partial W} \right) \right] \Delta W = -\Delta t R(W^n), \quad (1.8)$$

where

$$\Delta W = W^{n+1} - W^n. \quad (1.9)$$

If one were to directly solve this system of equations using LU factorization, the operation count would grow like $\mathcal{O}(N(b+1)^2)$ where b is the bandwidth of the left hand matrix and N is the number of equations [27]. In two dimensions the equations can be ordered so that the bandwidth is proportional to the product of the smallest dimension of the grid and the number of flow variables. In three spatial dimensions the bandwidth grows like the product of the two smallest dimensions and the number of flow variables. Consequently calculations on grids with even a moderate number of cells may become too expensive to compute. In this case the benefits afforded by the enlarged stability region are outweighed by the cost of directly inverting large matrices.

The first alternative to this approach is to use an iterative procedure and approximately solve the system of equations to a residual level that does not affect the overall accuracy of the solution. In this case, there is no need to linearize the temporal discretization scheme. A widely used approach is to add a pseudo time derivative

to the equations, and advance the solution in pseudo time until the residual of the discrete unsteady equation is negligible [34, 5, 13, 78]. This is commonly called the dual time-stepping method. It can be adapted to the Crank-Nicholson scheme shown in equation 1.6. However it is more commonly used with the backward difference formula:

$$\frac{dW^{n+1}}{d\tau} + \frac{3W^{n+1} - 4W^n + W^{n-1}}{2\Delta t} + R(W^{n+1}) = 0. \quad (1.10)$$

This has the advantage that it is both A-stable and damps the highest frequency modes. In the dual time-stepping method a nested system of loops is required to advance the solution. The inner loop is the set of iterations required by the solver to advance the solution forward in pseudo time to reduce the magnitude of the unsteady residual. The outer loop is the set of solutions at each physical instance in time. Extremely large physical time steps can be used due to the A-stable nature of the second order version of the time discretization.

Since the inner iterations are advanced in pseudo time, established iterative methods to accelerate the convergence (such as multigrid, residual averaging and local time stepping) can be applied without loss of temporal accuracy for the outer loop in physical time. These methods use explicit evaluations of the residual, avoiding the need to invert a linearized system of equations. The overall cost of the solution is proportional to the product of the work done per time step and the number of time steps required to drive the initial transients from the solution while accurately resolving it at its periodic steady state.

The second alternative is to factor the left hand side of equation 1.8, thus introducing an error term proportional to the product of Δt^2 and the change in the solution between time steps ΔW :

$$\begin{aligned} \left[I + \frac{\Delta t}{2} D_x \frac{\partial F}{\partial W} \right] \left[I + \frac{\Delta t}{2} D_y \frac{\partial G}{\partial W} \right] \Delta W = \\ -\Delta t R(W^n) - \frac{\Delta t^2}{4} D_x \frac{\partial F}{\partial W} D_y \frac{\partial G}{\partial W} \Delta W. \end{aligned} \quad (1.11)$$

Factoring the delta form of the equations has advantages in terms of convergence over factoring a non-delta form of the same equations [43]. Ignoring the factorization

error, one can solve the above equation in two steps, each requiring the solution of multiple block tridiagonal systems:

$$\begin{aligned} \left[I + \frac{\Delta t}{2} D_x \frac{\partial F}{\partial W} \right] \Delta \bar{W} &= -\Delta t R(W^n) \\ \left[I + \frac{\Delta t}{2} D_y \frac{\partial G}{\partial W} \right] \Delta W &= \Delta \bar{W}. \end{aligned} \quad (1.12)$$

The advantage of this approach is that the computational cost grows linearly with the number of grid points used in the calculation. The resulting scheme is nominally second order accurate with three sources of error: discretization error, linearization error, and factorization error. Unfortunately the factorization error can dominate at large CFL numbers, and the matrix inversion process is not easily amenable to parallel processing.

Commonly the time step that can be used with explicit algorithms is limited by the stability restrictions of the scheme. At first glance, A-stable schemes seem to be more efficient because they eliminate these restrictions. Unfortunately, there are complications when two separate time scales exist. Figure 1.2 shows the time history of the velocity at a point in the unsteady channel flow problem proposed in Appendix B. Note that the time scale associated with the decay of the initial transients is much larger than the time period of the oscillation of the solution in its periodic steady state. Although this is a model problem, similar phenomena are found in other applications such as aeroacoustics, aeroelasticity [15] and turbomachinery (see section 1.1). Regardless of the stability criteria, the time step is limited by the number of points required to resolve one time period (see modified wavenumber discussion in section 2.2.3). In this case, the time step is a fraction of the period of oscillation which in turn is a fraction of the time scale associated with the initial decay. If the time step is fixed over the duration of the calculation, the majority of work is spent resolving the solution during its initial decay.

In order to take full advantage of implicit schemes, one must adaptively change the time step during the evolution to a periodic steady state. The definition of robust criteria for setting the time step is difficult for non-linear problems in complex




Figure 1.2: Perturbed velocity at any point along the unsteady channel with a reduced frequency $\Omega = 50$.

geometries. Taking this approach to an extreme, one can solve for the time-averaged solution to initially seed the time-accurate calculation. This procedure is analogous to using an implicit time-accurate scheme with an infinite time step. Unfortunately, since the time-average is not a solution to the time-accurate equations, an initial transient will still be created, although it will be smaller than one created by a bad initial guess. In turbomachinery calculations this technique sometimes decreases the cost of the calculation by roughly half [12].

1.2.2 Periodic Methods

An alternative to the time-accurate methods discussed in previous sections is to assume that the solution is periodic over time. This is not a time-accurate approach in the sense that it does not admit transient decay as a component of the solution. Instead the method directly calculates the solution at its periodic steady state. Given the assumption of periodicity, trigonometric interpolants such as a Fourier series, can be used to accurately and efficiently represent the solution. The application of these

methods to a non-linear system of equations presents some difficulties. The following sections cover the different approaches in the application of these techniques culminating in the full Non-Linear Frequency Domain (NLFD) method, which is the subject of this research.

Linearized Frequency Domain Methods

A commonly-used approach is to linearize the governing equations by splitting the solution variable into steady \bar{W} and unsteady components \tilde{W} :

$$W = \bar{W} + \tilde{W}. \quad (1.13)$$

The magnitude of the unsteady (or perturbed) component is assumed to be much smaller than the magnitude of the steady component. Substitution of the split form of the solution into the governing equations leads to the following equation:

$$\begin{aligned} \frac{\partial \tilde{W}}{\partial t} + \frac{\partial F(\bar{W})}{\partial x} + \frac{\partial F(\bar{W}, \tilde{W})}{\partial x} + \frac{\partial F(\bar{W}, \tilde{W}^{2+})}{\partial x} \\ + \frac{\partial G(\bar{W})}{\partial y} + \frac{\partial G(\bar{W}, \tilde{W})}{\partial y} + \frac{\partial G(\bar{W}, \tilde{W}^{2+})}{\partial y} = 0. \end{aligned} \quad (1.14)$$

Note that the flux vectors F and G have been split into three components.

- $F(\bar{W})$ are the components of F depending only on the mean of the solution, \bar{W} .
- $F(\bar{W}, \tilde{W})$ are the components of F depending on the mean of the solution \bar{W} and terms which are first order in the unsteady variable \tilde{W} . This term and its analogue in the y-direction can be rewritten as the product of the Jacobian matrices, calculated using the mean flow values, and the perturbed variables as follows:

$$\begin{aligned} F(\bar{W}, \tilde{W}) &= \bar{A}\tilde{W} \\ G(\bar{W}, \tilde{W}) &= \bar{B}\tilde{W}. \end{aligned} \quad (1.15)$$

Where A and B are the Jacobian matrices

$$A = \frac{\partial F}{\partial W} \quad B = \frac{\partial G}{\partial W}. \quad (1.16)$$

- $F(\bar{W}, \tilde{W}^{2+})$ are the components of F depending on the mean of the solution \bar{W} and terms which are second order and greater in the perturbed variable \tilde{W} .

To simplify equation 1.14, assume that the mean flow is solved to a residual level that is negligible in comparison with the unsteady perturbations. Second, ignore all terms in the derivation that are second order in the perturbed variables ($F(\bar{W}, \tilde{W}^{2+})$ and $G(\bar{W}, \tilde{W}^{2+})$). This assumption is valid, given that the perturbed variables were assumed to be small in comparison to the mean flow, and hence first order terms will be dominant. The linearized form of the governing equations for both the time-averaged and unsteady modes can be stated as:

$$\begin{aligned} \frac{\partial F(\bar{W})}{\partial x} + \frac{\partial G(\bar{W})}{\partial y} &= 0 \\ \frac{\partial \tilde{W}}{\partial t} + \frac{\partial(\bar{A}\tilde{W})}{\partial x} + \frac{\partial(\bar{B}\tilde{W})}{\partial y} &= 0. \end{aligned} \quad (1.17)$$

The perturbed variable can be represented as a Fourier series

$$\tilde{W}(t) = \sum_{k=-\infty}^{\infty} \hat{W}_k e^{ikt}. \quad (1.18)$$

Taking advantage of orthogonality, substitution of the above series into the unsteady equation leads to the following form:

$$ik\hat{W}_k + \frac{\partial(\bar{A}\hat{W}_k)}{\partial x} + \frac{\partial(\bar{B}\hat{W}_k)}{\partial y} = 0. \quad (1.19)$$

The procedure is to first solve for the mean flow components of the solution. Then one can solve for any temporal frequency of the solution using the mean flow components to calculate the Jacobian matrices \bar{A} and \bar{B} . Due to the linearity of the governing equations, all of the temporal modes are uncoupled and can be calculated

independently and simultaneously. The cost of this solution procedure is proportional to the product of the number of temporal modes calculated and the cost of the steady solution. The scheme does have drawbacks. Applications in turbomachinery and aeroelasticity often require transonic solutions where the movement of the shock wave is non-negligible, making this approach inapplicable due to the assumptions of the linearization.

Deterministic-Stress Methods

Adamczyk [1] proposed several different linearizations and averaging operators of the velocity variable to form what he termed the *deterministic stress*. This term is similar to a Reynolds stress in that it attempts to quantify the effects of the unsteady field on the time averaged solution. Adamczyk proposed modeling these terms, but subsequent authors [25, 11] proposed calculating these terms with a modified version of a linearized frequency domain solver. The method begins with the split form of the governing equations presented in equation 1.14. This equation is then integrated over a period of time T equivalent to the time period of the fundamental harmonic. Terms which are periodic over this interval do not contribute to the integral and can be dropped, resulting in the following equation

$$\frac{F(\bar{W})}{\partial x} + \frac{G(\bar{W})}{\partial y} + \frac{1}{T} \int_0^T \frac{F(\bar{W}, \tilde{W}^{2+})}{\partial x} + \frac{G(\bar{W}, \tilde{W}^{2+})}{\partial x} dt. \quad (1.20)$$

The terms contained in the integral are labeled the deterministic stresses and represent the feedback that the unsteady modes have on the time-averaged solution. The unsteady modes are solved using the same linear equation presented above, except now the modes are coupled and require simultaneous solution of both the time-averaged and unsteady terms. Assuming the coupling process does not affect the convergence rate of the solver, the execution time required by the non-linear method should be equivalent to the execution duration of a linearized method. A slight penalty is incurred by the additional memory required to hold all the unsteady modes simultaneously.

Although some of the nonlinearities are addressed in the time-averaged solution,

the higher order terms are still neglected in the solution of the unsteady modes. The assumptions of the linearization still must apply for these modes, limiting the applicability of the method. In addition, the method couples only the time-averaged solution to the higher harmonics. No direct coupling between the harmonics is mathematically evident. This method will poorly represent physical problems where significant energy is transferred between the fundamental and higher harmonics. Despite these shortcomings, Adamczyk-type methods are widely used throughout the turbomachinery industry.

Non-Linear Frequency Domain Methods

None of the preceding methods can adequately account for the strong nonlinearities which are typically encountered in flows in turbomachinery and rotorcraft. This brings us to methods which represent the full non-linear equations in the frequency domain. These are the focus of the research reported in this thesis. The Harmonic Balance method proposed by Hall *et al.* [24] represents the first full non-linear method in the frequency domain.

Beginning with the governing equations presented in equation 1.1, both the residual and the solution can be represented by Fourier series in time (see equation 1.18). Taking advantage of the orthogonality of Fourier series, the resulting governing equation can be expressed in the frequency domain as:

$$ik\hat{W}_k + \hat{R}_k = 0. \quad (1.21)$$

Here, because $R(t)$ is a non-linear function of $W(t)$, \hat{R}_k can not be easily determined from \hat{W}_k . A more efficient approach is to obtain $W(t)$ as the transform of \hat{W}_k , use it to calculate $R(t)$, and obtain \hat{R}_k as the transform of $R(t)$.

Transforming equation 1.21 back into the temporal domain results in a coupled set of equations which represent the unsteady residual terms at each instance in the time history of the solution

$$S(t) + R(t) = 0, \quad (1.22)$$

where $S(t)$ is the spectral representation of $\frac{\partial W}{\partial t}$ under the assumption of periodicity

$$S(t) = \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} ik\hat{W}_k e^{ikt}. \quad (1.23)$$

With the addition of a pseudo-time term, established methods of convergence acceleration developed for steady flows can now be applied to this stationary system of equations. Hall adds pseudo-time derivatives to the coupled equations 1.22 to solve for the unsteady residuals in the time domain. In the present work pseudo-time derivatives are added directly to the equations 1.21 in the frequency domain. In contrast to conventional time-domain methods, the non-linear frequency domain approach in either of these variations constitutes an integrated space-time discretization in which the equations must simultaneously be solved for all the discrete samples in both space and time.

Assuming that all of the relevant modes in the solution and residual are represented, the method will account for all the nonlinearities in the governing equations. However, Hall *et al.* [23, 24] and McMullen *et al.* [48, 49] show that a limited number of modes are often sufficient to capture the dominant physics in the solution. Limiting the number of modes has the obvious advantage of decreasing the number of instances of the unsteady residual that need to be calculated. If all the modes of the solution converge as quickly as a similar steady state calculation, then the cost of the calculation is the product of the cost of a steady solution and the number of instances used in the time series of the unsteady residual.

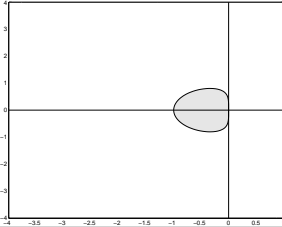
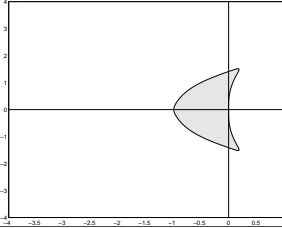
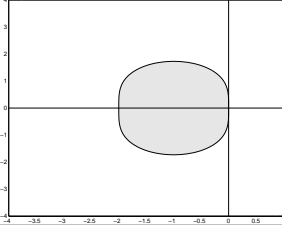
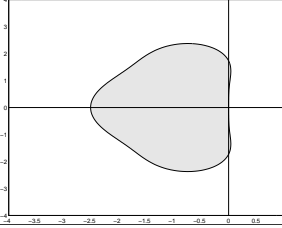
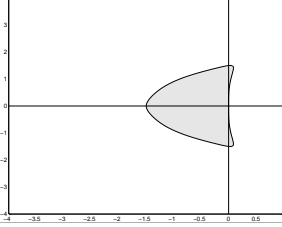
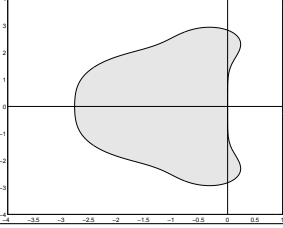
Scheme	Local Mag	Error Phase	Global Error Mag	Error Phase	Stability Diagram
Adams Bashforth	$\frac{p^4}{4}$	$\frac{5p^3}{12}$	2.18e-02	9.79	
Leapfrog	$\frac{p^4}{4}$	$\frac{p^3}{12}$	9.17e-02	4.45	
RK-2	$\frac{p^4}{8}$	$\frac{p^3}{6}$	8.31e-02	14.5	
RK-3	$\frac{p^4}{24}$	$\frac{p^5}{30}$	7.63e-02	2.48	
Hyman	$\frac{p^4}{12}$	$\frac{4p^5}{45}$	4.43e-02	1.30	
RK-4	$\frac{p^6}{144}$	$\frac{p^5}{120}$	2.25e-02	1.96	

Table 1.3: The error and stability characteristics of a number of explicit multistep methods. Note: $ip = \lambda\Delta t$. Source: Pulliam [43]

1.3 Approach

The main objective of this work is to improve the efficiency of periodic steady state flow solvers. While the Non-Linear Frequency Domain method leads to an integrated space-time discretization, this research is focused primarily on the representation of the temporal derivatives. The approach which has been adapted is to modify a steady state code with its existing spatial operators by adding an implementation of the Non-Linear Frequency Domain method. The efficiency and accuracy of the new method have been verified for three different model problems. Each model problem demonstrates a different aspect of the functionality of the new solver.

- Linearized unsteady channel flow was chosen as the initial model problem. An analytic solution derived in the context of linearized perturbation theory is available for comparison. The decay rates of the initial transients can be expressed as a simple exponential function and the convergence rates of the NLFD solver can be shown to be independent of this parameter.
- The second model problem was the laminar vortex shedding behind a circular cylinder. Primarily the case demonstrates the efficiency of the solver in representing complex flow field physics with a limited number of temporal modes. Since the shedding frequency is unknown *a priori*, it also serves to test the application of gradient-based-variable-time-period methods. These methods were developed by this research to iteratively find the exact time period of the discrete equations given an initial guess in its vicinity.
- The third model problem was transonic flow about a pitching airfoil. Due to the popularity of this case in aeroelastic research, this research has the luxury of comparing NLFD results to multiple experiments and time-accurate codes. The case demonstrates the efficiency and accuracy of the solver in capturing the fully non-linear physics of moving shock waves.

The emphasis of this research is on algorithm development. It is not the main intent of this work to provide additional understanding of physical phenomena. As such, the model problems chosen in this work are relatively simple. The intent was to

reduce the physical complexity of the model problems to assist in the isolation of the algorithm strengths and weaknesses.

The other criterion used in the selection of the model problems was the existence of analytic, computational or experimental data previously obtained in independent investigations. A time-accurate code had already been developed for the case of the pitching airfoil, providing the opportunity to compare both the accuracy and efficiency of the different techniques.

1.4 Overview

The research has produced several noteworthy results.

- Harmonic balance methods, as proposed by K. Hall [23, 24] for the Euler equations, have been recast in the frequency domain and extended to the laminar and turbulent Navier-Stokes equations. Modifications to convergence acceleration methods (multigrid, residual averaging, and pseudo-local time stepping) have been successfully applied to guarantee convergence rates equivalent to state-of-the-art steady flow solvers.
- A gradient-based-variable-time-period method has been developed to allow frequency domain methods to be applied to classes of problems where the dominant frequency is unknown *a priori*. The method is shown to improve the quality of the solution over fixed-time-period computations, which do not completely minimize the magnitude of the unsteady residual.
- Comparisons with experimental datasets and other numerically generated results have been made. The experimental comparisons support the conclusions that the NLFD methods can accurately describe fully non-linear problems using a limited number of modes. Comparison to state of the art time-accurate codes show that NLFD methods offer significant savings in computational cost.
- The convergence rate of the Non-Linear Frequency Domain method is shown

to be independent of the decay rates of the initial transients. Because time-accurate solvers are accurately capturing the decay of the initial transients, the computational cost of these solvers are direct functions of this decay rate. Alternatively, NLFD methods solve only for solutions that exhibit periodic behavior.

Chapter 2 documents the numerical methods employed in this research. In addition to describing the frequency domain method, this chapter outlines aspects of the spatial discretization, boundary conditions, and convergence acceleration techniques used in the solver. Chapter 3 presents results from laminar vortex shedding behind a cylinder. This chapter compares NLFD solutions (with and without variable time period methods) to experimental data sets. The goal is to verify the accuracy of the NLFD and variable-time-period methods. Chapter 4 compares NLFD solutions to experimental data from a transonic pitching airfoil at a fixed frequency. This chapter begins by comparing steady NLFD results with experimental data in an attempt to verify the spatial accuracy of the grids employed. These same grids are the basis for unsteady NLFD computations which are compared to different sets of unsteady experimental data. The goal of this chapter is to verify the accuracy of the NLFD method in a non-linear environment containing significant shock movements. Chapter 5 compares the efficiency of a NLFD solver to time-accurate codes. To eliminate the impact of spatial error on the comparison, the chapter begins by demonstrating the equivalence of the spatial operators contained in the two codes. An efficiency comparison is performed based on the amount of work needed to reach equivalent temporal error levels.

Chapter 2

Methodology

2.1 Governing Equations

For an arbitrary volume of fluid Ω , conservation of mass, momentum and energy can be expressed in the following integral form:

$$\frac{d}{dt} \int_{\Omega} W dV + \oint_{\partial\Omega} \vec{F} \cdot \vec{N} ds = 0. \quad (2.1)$$

This integral form contains volumetric integrals denoted as $\int_{\Omega} dV$ and surface integrals denoted as $\oint_{\partial\Omega} ds$. The outward pointing unit normal to the surface is defined as \vec{N} . For a two-dimensional flow, the physical properties of the fluid (density ρ , Cartesian velocity components u_i , and stagnation energy E) are collected into the state vector

$$W = \begin{bmatrix} \rho \\ \rho u_1 \\ \rho u_2 \\ \rho E \end{bmatrix}. \quad (2.2)$$

The transport and/or production of these properties is accounted for in the flux vector \vec{F} . For convenience this overall flux is split into convective, \vec{F}_c , and viscous, \vec{F}_v , components

$$\vec{F} = \vec{F}_c - \vec{F}_v. \quad (2.3)$$

The convective fluxes will include those terms normally associated with the Euler equations and the fluxes associated with a moving control volume. The velocity of the surface of the control volume is denoted as \vec{b} . Using indicial notation these terms can be defined for each coordinate direction, i , in two-dimensional space as:

$$F_{c_i} = \begin{bmatrix} \rho & (u_i - b_i) \\ \rho u_1 & (u_i - b_i) + \delta_{1i} p \\ \rho u_2 & (u_i - b_i) + \delta_{2i} p \\ \rho E & (u_i - b_i) + u_i p \end{bmatrix}. \quad (2.4)$$

The viscous fluxes will include the stress tensor terms associated with viscous dissipation

$$F_{v_i} = \begin{bmatrix} 0 \\ \sigma_{i1} \\ \sigma_{i2} \\ u_j \sigma_{ij} + q_i \end{bmatrix}. \quad (2.5)$$

Closure for the energy and viscous stress tensor terms is provided by the following relationships:

$$\begin{aligned} p &= (\gamma - 1)\rho \left[E - \frac{1}{2}(u_j u_j) \right] \\ \sigma_{ij} &= \mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) + \lambda \delta_{ij} \frac{\partial u_k}{\partial x_k} \\ q_i &= \kappa \frac{\partial T}{\partial x_i} = \frac{\gamma}{\gamma - 1} \frac{\mu}{P_r} \frac{\partial}{\partial x_i} \left(\frac{p}{\rho} \right). \end{aligned} \quad (2.6)$$

For all the analyses presented in this thesis the ratio of specific heats, γ , and the Prandtl number, P_r , are held constant at 1.4 and 0.72 respectively. The relationship between absolute viscosity μ and temperature T is determined by Sutherland's law.

Following the analysis presented in appendix A.1 the conservation laws can be stated in dimensionless form as:

$$\frac{d}{dt^*} \int_{\Omega} W^* dV^* + \oint_{\partial\Omega} \vec{F}_c^* \cdot \vec{N} ds^* = \sqrt{\gamma} \frac{M_{\infty}}{Re_{\infty}} \oint_{\partial\Omega} \vec{F}_v^* \cdot \vec{N} ds^*. \quad (2.7)$$

2.2 Numerical Method

2.2.1 Transforming the Equations into the Frequency Domain

This section presents the complete discretization used to approximate the conservation laws presented in equation 2.7. These equations are approximated using the finite volume approach where the continuous surface integrals are represented by a discrete summation of fluxes across a finite number of faces on the control volume:

$$\oint_{\partial\Omega} \vec{F} \cdot \vec{N} ds = \sum_{cv} \vec{F} \cdot \vec{S}. \quad (2.8)$$

The approximation of the flux vector, \vec{F} , that ensures numerical stability is the subject of shock capturing theory, which has been developed over the last two decades. This theory is well established and is not the focus of this research. Consequently, the implementation of these schemes and all other spatial operators used in this research are described in the appendixes.

Having approximated all the spatial operators, the remaining term in the conservation laws is the temporal derivative of the volumetric integral of the solution. This is approximated as the product of the cell volume with the temporal derivative of the average of the solution over the cell. Adding the spatial and temporal operators together results in the complete approximation to the governing equations:

$$V \frac{\partial W}{\partial t} + \sum_{cv} \vec{F}_c \cdot \vec{S} - \sum_{cv} F_d - \frac{\sqrt{\gamma} M_\infty}{Re_\infty} \sum_{cv} \vec{F}_v \cdot \vec{S} = 0. \quad (2.9)$$

The spatial operator R is introduced as a function of space and time including both the convective and dissipative fluxes.

$$R = \sum_{cv} \vec{F}_c \cdot \vec{S} - \sum_{cv} F_d - \frac{\sqrt{\gamma} M_\infty}{Re_\infty} \sum_{cv} \vec{F}_v \cdot \vec{S}. \quad (2.10)$$

The term includes not only the natural viscous dissipation, \vec{F}_v , but also any artificial dissipation, F_d , provided by shock capturing schemes and turbulence modeling. Taking advantage of this simplified notation, a semi-discrete form of the governing equations can be written as:

$$V \frac{\partial W}{\partial t} + R = 0. \quad (2.11)$$

Assuming that the solution W and spatial operator R are periodic in time both can be represented by separate Fourier series:

$$\begin{aligned} W &= \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} \hat{W}_k e^{ikt} \\ R &= \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} \hat{R}_k e^{ikt} \end{aligned} \quad (2.12)$$

where,

$$i = \sqrt{-1}. \quad (2.13)$$

These discrete Fourier transforms can be substituted into the semi-discrete form of the governing equations provided by equation 2.11, and the time derivative of the state variable can be moved inside the series summation. Taking advantage of the orthogonality of the Fourier terms results in a separate equation for each wavenumber, k , in the solution

$$ikV\hat{W}_k + \hat{R}_k = 0. \quad (2.14)$$

Here, however each coefficient \hat{R}_k of the transform of the residual depends on all the coefficients \hat{W}_k , because $R(W(t))$ is a non-linear function of $W(t)$. Thus (2.14) represents a non-linear set of equations which must be iteratively solved. The solver attempts to find a solution, W , that drives this system of equations to zero for all wavenumbers, but at any iteration in the solution process the unsteady residual, \hat{I}_k , will be finite:

$$\hat{I}_k = ikV\hat{W}_k + \hat{R}_k. \quad (2.15)$$

The nonlinearity of the unsteady residual stems from the spatial operator. There

are two approaches to calculating the spatial operator expressed in the frequency domain. The first uses a complex series of convolution sums to calculate \hat{R}_k directly from \hat{W}_k . Such an approach was discussed in Hall's introductory paper on Harmonic Balance techniques [24]. Hall justly discarded the approach due its massive complexity (considering artificial dissipation schemes and turbulence modeling) and cost that scales quadratically with the number of modes N .

The alternative proposed by Hall and implemented by this research is to use a pseudo-spectral approach that relies on the computational efficiency of the Fast Fourier Transform (FFT). A diagram detailing the transformations used by the pseudo spectral approach is provided in figure 2.1.

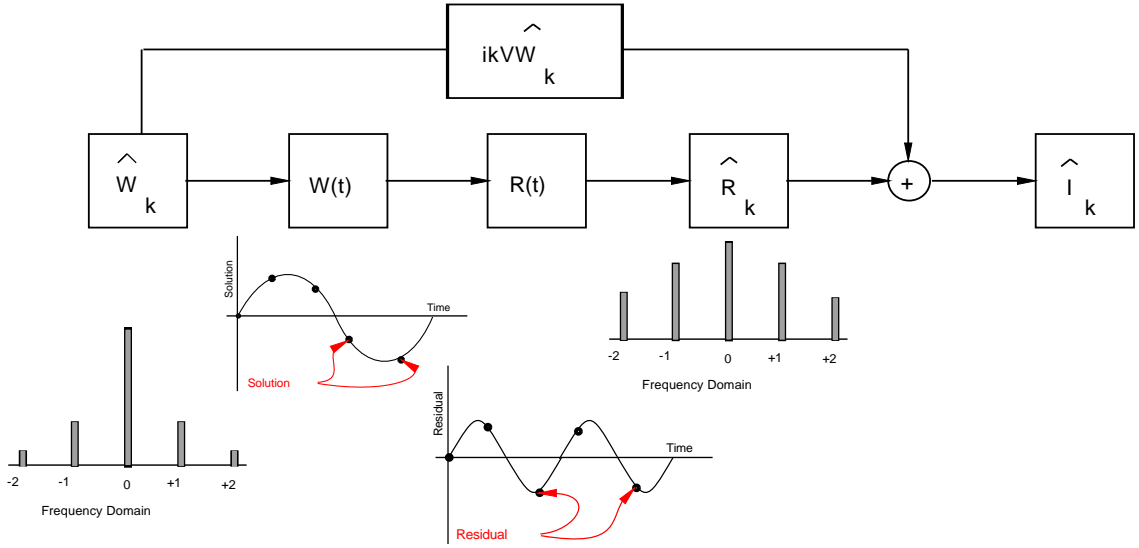


Figure 2.1: Simplified dataflow diagram of the time advancement scheme illustrating the pseudo spectral approach used in calculating the non-linear spatial operator R .

The pseudo-spectral approach begins by assuming that \hat{W}_k is known for all wavenumbers. Using an inverse FFT, \hat{W}_k can be transformed back to the physical space resulting in a state vector $W(t)$ sampled at evenly distributed intervals over the time period. At each of these time instances the steady-state operator $R(W(t))$ can be computed. A FFT is then used to transform the spatial operator to the frequency domain where \hat{R}_k is known for all wavenumbers. The unsteady residual \hat{I}_k can then

be calculated by adding \hat{R}_k to the spectral representation of the temporal derivative $ikV\hat{W}_k$.

The cost of evaluating the spatial operator is the product of the cost of evaluating a steady-state spatial operator and the number of time instances used to represent the solution N . The cost of the FFT is proportional to $N \ln(N)$. For most realistic values of N ($N = 1 \rightarrow 10$) the cost of the pseudo-spectral approach is dominated by the cost associated with calculating the spatial operator. Although this estimate is not exact, timing estimates provided in section 5.4 substantiate these statements.

One of the advantages of the pseudo-spectral approach is its flexibility in admitting different forms of the non-linear operators. This research uses finite volume formulations, but application of the pseudo-spectral approach is equally well suited for finite difference or other types of spatial operators. In addition, turbulence models which would be difficult to explicitly express in the frequency domain are easily admitted in the pseudo-spectral approach.

2.2.2 Pseudo-Time Derivative

This section discusses the basic strategy used by the solver to find a solution that minimizes the magnitude of the unsteady residual. A simple iterative approach for solving this equation at all wavenumbers except for the mean value ($k = 0$) would be to solve for the temporal derivative as

$$\hat{W}_k = \frac{-1}{ikV} \hat{R}_k. \quad (2.16)$$

Given a linearization of the spatial operator, $\hat{R}_k^n = \lambda \hat{W}_k^n$, an equation defining the growth factor of the solution is provided as:

$$\frac{\hat{W}_k^{n+1}}{\hat{W}_k^n} = \frac{-1}{ikV} \lambda. \quad (2.17)$$

The magnitude of this growth factor tends to infinity as kV tends to zero. Consequently the stability characteristics of this approach are poor for the low frequency cases ($k \ll 1$).

Instead of directly solving equation 2.14 a pseudo-time derivative can be added, and a time-stepping scheme can be employed to numerically integrate the resulting equations

$$V \frac{\partial \hat{W}_k}{\partial \tau} + \hat{I}_k = 0. \quad (2.18)$$

The application of the pseudo-time derivative is consistent with established convergence acceleration techniques used to solve steady-state problems. In the NLFD case, an unsteady residual exists for each wavenumber used in the solution and the pseudo-time derivative acts as a gradient to drive the absolute value of all of these components to zero simultaneously.

2.2.3 Temporal Resolution

This section investigates the accuracy of the temporal discretization of the NLFD method relative to another discretization used in a common time-accurate method. To facilitate this discussion we will base the following analysis on the Harmonic Balance form of the equations, which yield solutions equivalent to the NFD approach. As stated in section 1.2, the Harmonic Balance form is

$$S(t) + R(t) = 0, \quad (2.19)$$

where $S(t)$ is the spectral representation of $\frac{\partial W}{\partial t}$ under the assumption of periodicity

$$S(t) = \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} ik \hat{W}_k e^{ikt}. \quad (2.20)$$

The time-accurate code proposed as an example for this section approximates the temporal derivative using a second-order implicit discretization of the following form:

$$\frac{3W^{n+1} - 4W^n + W^{n-1}}{2\Delta t} + R(W^{n+1}) = 0. \quad (2.21)$$

At issue is the accuracy of $S(t)$ in the NLFD method relative to the 3-4-1 difference operator in the time-accurate approach. The comparison begins with a modified-wavenumber analysis where a periodic basis function e^{ikt} containing only one mode is substituted into the difference operator resulting in an expression for the modified wavenumber k' :

$$\begin{aligned} D_t \hat{W}_k^n &= \frac{[3 - 4\cos(k\Delta t) + \cos(2k\Delta t)] + [4\sin(k\Delta t) - \sin(2k\Delta t)]i}{2\Delta t} \hat{W}_k^n \\ &= k' \hat{W}_k^n. \end{aligned} \quad (2.22)$$

According to the Sampling theorem [54] the discrete Fourier transform exactly represents any signal if the frequency of this mode lies within the sampling frequency, k_c :

$$\begin{aligned} D_t \hat{W}_k^n &= ik \hat{W}_k^n \\ k' &= ik \quad \text{where } k < k_c. \end{aligned} \quad (2.23)$$

The sampling frequency is defined by the Nyquist limit of two points per wavelength. Since the exact derivative is purely imaginary, any real components of the modified wavenumber are error and represent the dissipative effects of the scheme. The dispersive error is associated with any imaginary component that deviates from the linear relationship specified in equation 2.23.

The results of the modified-wavenumber analysis for the two temporal operators are plotted in figure 2.2. The horizontal axis is quantified in terms of the product of the wavenumber and time step ($k\Delta t$), which is inversely related to the points per wavelength

$$N = \frac{T}{\Delta t} = \frac{2\pi}{k\Delta t}. \quad (2.24)$$

Vertical lines have been drawn to show values of points per wavelength. The horizontal axis is terminated at the Nyquist sampling limit of two points per wavelength. The figure shows that significant variations between the NLFD and 3-4-1 operators begin at approximately eight points per wavelength.

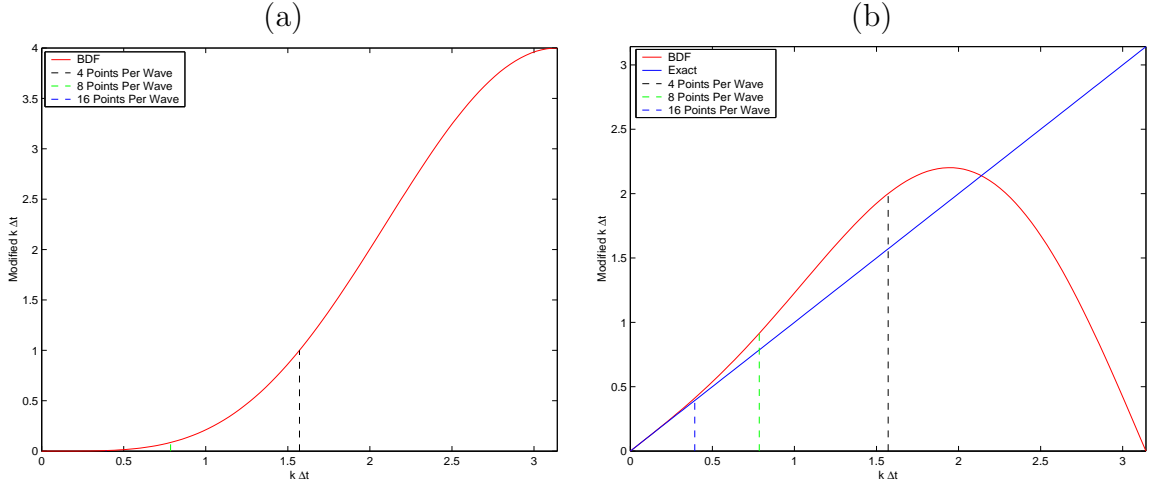


Figure 2.2: Modified wavenumber analysis for the backward difference formula employed in UFLO82. By definition $N = \frac{2\pi}{k\Delta t}$. (a) Real component. (b) Imaginary component

The modified wavenumber analysis gives no insight into the accuracy of the spectral representation for solutions containing frequencies outside of the sampling frequency of the series. The energy from these unresolved modes is transferred back into the solution in a corrupting process known as aliasing. In this case, the overall accuracy is a function of the convergence rate of Fourier series, which has been estimated by Snider [60]. He proceeds in two stages. First he considers the rate of convergence of a true Fourier series,

$$S_n(t) = \frac{A_0}{2} + \sum_{r=1}^n (a_r \cos(rt) + b_r \sin(rt)). \quad (2.25)$$

to a smooth periodic function, $g(t)$, as the number of terms is increased. Here the true Fourier coefficients are

$$\begin{aligned} a_r &= \frac{1}{\pi} \int_0^{2\pi} g(t) \cos(rt) dt \\ b_r &= \frac{1}{\pi} \int_0^{2\pi} g(t) \sin(rt) dt. \end{aligned} \quad (2.26)$$

It follows from repeated integration by parts that

$$\begin{aligned} |a_r| &\leq \frac{2M}{r^K} \\ |b_r| &\leq \frac{2M}{r^K} \end{aligned} \quad (2.27)$$

where M is a bound on the K^{th} derivative of $g(t)$

$$|g^K(t)| \leq M. \quad (2.28)$$

Thus the rate of decay of the coefficients is very fast if K is large. The error between $S_n(t)$ and $g(t)$ may now be estimated as the remainder of the series, consisting of the terms from $n+1$ to ∞ . This yields the result

$$|S_n(t) - g(t)| \leq \frac{2M}{K-1} \frac{1}{n^{K-1}} \quad (2.29)$$

Snider then considers the additional error when the true Fourier coefficients are replaced by the interpolated coefficients to give

$$C_n(t) = \frac{A_0}{2} + \sum_{r=1}^{n-1} (A_r \cos(rt) + B_r \sin(rt)) + \frac{A_n}{2} \cos(nt) \quad (2.30)$$

where,

$$\begin{aligned} A_r &= \frac{1}{n} \sum_{j=0}^{2n-1} g(t_j) \cos(rt_j) \\ B_r &= \frac{1}{n} \sum_{j=0}^{2n-1} g(t_j) \sin(rt_j) \end{aligned} \quad (2.31)$$

which are provided by the discrete transform. Here the result of the aliasing is manifested in the formulas

$$A_r = a_r + \sum_{l=1}^{\infty} (a_{2ln-r} + a_{2ln+r})$$

$$B_r = b_r + \sum_{l=1}^{\infty} (b_{2ln-r} + b_{2ln+r}) \quad (2.32)$$

in which the higher harmonics are absorbed into the interpolation coefficients, A_r and B_r . The rapid decay of the higher order coefficients of signal $g(t)$ possessing K continuous derivatives finally results in a bound for the error between the exact and the discrete representations:

$$\begin{aligned} |g(t) - C_n(t)| &< \frac{13 \max g^{(K)}(t)}{n^{K-1}} \\ &= \mathcal{O}\left(\frac{1}{n^{K-1}}\right). \end{aligned} \quad (2.33)$$

This estimate, which includes the effects of aliasing, asserts that a discrete Fourier series will converge to a smooth function ($k \rightarrow \infty$) at a rate faster than any power of n . This exponential convergence rate is faster than convergence rates which could be obtained with any finite difference operator in the temporal domain. In this sense the spectral representation provides the optimal time accuracy.

The NLFD method requires the user to specify a finite number of temporal modes to be used in the representation of the solution. In an effort to reduce the computational cost, users typically specify a small number of modes, deliberately choosing to ignore the effects of harmonics outside those used by the Fourier series. In this case, the NLFD method falls under a general class of schemes known as reduced order methods [8].

In non-linear environments energy from one temporal mode is transferred to other harmonics. The exact solution contains energy distributed through a large number of modes, many of which go unresolved in the numerical simulation. The non-linear processes which transfer energy between modes vary with the problem. Since the error is a function of this transfer mechanism, the application of the NLFD scheme as a reduced order method must be verified on a case by case basis. This verification process should categorize the type of results accurately predicted for each different flow field. Capturing only the dominant modes in the solution field can provide accurate estimates for global properties such as C_l and C_d , but may produce inferior

results for the local unsteady pressure coefficients [48].

2.2.4 Variable Time Period

Unsteady flows can be subdivided into two classes. The first class includes flows where the boundary conditions force the unsteadiness at predetermined frequencies. Examples of this include the internal flows of turbomachinery, the external flow fields of helicopter blades or propellers, and certain aeroelastic-elastic computations. Problems within the forced classification are the easiest to solve, since the time period is known at the beginning of the NLFD simulation.

The second general class of unsteady flows are those where the instabilities in the equations of fluid mechanics allow unsteadiness to evolve into the flow field. Examples in this class include (but are obviously not limited to) vortex shedding behind a cylinder and other fluid dynamic cases involving separated flows and free shear layers. Using experimental or simplified analytic models, the temporal frequencies can be estimated *a priori* for this second class. However, any approximation of this type will not be exactly equivalent to the natural frequency of the discretized equations.

A major outcome of this research has been the development of a gradient approach for this second class of problems that iteratively determines the time period of the fundamental harmonic. In practice, the user provides an initial guess in the vicinity of the final answer and the gradient information is used to change the estimate to the time period after each iteration in the solution process. Convergence is achieved when changes in the fundamental time period and all of the residual levels in the unsteady equations are negligible.

A derivation of the Gradient-Based-Variable-Time-Period (GBVTP) method begins by noting that the wavenumber k is calculated by normalizing the sinusoidal period of oscillation 2π with the time period of interest T

$$k = \frac{2\pi n}{T}. \quad (2.34)$$

The unsteady residual in equation 2.18 can then be written as a function of the time

period T

$$\hat{I}_n = \frac{i2\pi nV}{T} \hat{W}_n + \hat{R}_n. \quad (2.35)$$

The process of finding a solution to the unsteady flow equations is regarded as an optimization problem where the magnitude of the unsteady residual is minimized. A gradient of this residual with respect to the time period can be calculated and used to iteratively modify this time period until the magnitude of all components of the unsteady residual are negligible.

Because \hat{I}_n is a complex quantity, the figure of merit will be defined as the square of the magnitude of this quantity. Using the chain rule, a gradient of this cost function with respect to the time period can be written as

$$\frac{1}{2} \frac{\partial |\hat{I}_n|^2}{\partial T} = \hat{I}_{nr} \frac{\partial \hat{I}_{nr}}{\partial T} + \hat{I}_{ni} \frac{\partial \hat{I}_{ni}}{\partial T}. \quad (2.36)$$

The quantity \hat{I}_n is already calculated while monitoring the convergence of the solution (note that the real and imaginary parts of \hat{I}_n are \hat{I}_{nr} and \hat{I}_{ni} respectively). The partial derivative terms can be expanded as:

$$\begin{aligned} \frac{\partial \hat{I}_{nr}}{\partial T} &= \frac{2\pi nV \hat{W}_{ni}}{T^2} \\ \frac{\partial \hat{I}_{ni}}{\partial T} &= -\frac{2\pi nV \hat{W}_{nr}}{T^2}. \end{aligned} \quad (2.37)$$

The formulas can be further simplified by introducing cross product notation with the Fourier coefficients of the solution and residual written as vectors

$$\begin{aligned} \vec{W}_n &= \hat{W}_{nr} \mathbf{e}_1 + \hat{W}_{ni} \mathbf{e}_2 \\ \vec{I}_n &= \hat{I}_{nr} \mathbf{e}_1 + \hat{I}_{ni} \mathbf{e}_2. \end{aligned} \quad (2.38)$$

Using this notation the gradient can be expressed as the magnitude of the cross

product of the above vectors

$$\frac{1}{2} \frac{\partial |\hat{I}_n|^2}{\partial T} = \frac{2\pi n V}{T^2} |\vec{I}_n \times \vec{W}_n|. \quad (2.39)$$

The time period can be updated using the gradient information by selecting a stable step ΔT

$$T^{n+1} = T^n - \Delta T \frac{\partial |\hat{I}_n|^2}{\partial T}. \quad (2.40)$$

Typically one can start with an initial guess in the vicinity of the final answer for the time period. An unsteady flow solution can be obtained by solving the unsteady equations to some residual level. The above gradient can then be used to adjust the time step at each iteration in the solution. The solution and gradients are hence simultaneously updated, and a final solution can be calculated to any arbitrary residual level.

A numerical experiment was performed to validate the GBVTP method. The NLFD code was used to calculate the flow over a cylinder (see chapter 3) at a Reynolds number of 105 in both fixed- and variable-time-period modes. While driving the residual to *machine zero*, the GBVTP algorithm determined the exact Strouhal number for the discretized equations. The fixed-time-period calculations were seeded with Strouhal numbers over a range of values in the vicinity of the Strouhal number determined by the GBVTP method. The minimum residual obtained during each fixed time period calculation is provided in figure 2.3. The data within this figure is plotted in two different formats. The first format shows the minimum residual of the fixed-time-period calculation as a function of the Strouhal number. A dramatic change in the minimum residual levels occurs at the Strouhal number near the value calculated by the GBVTP method. To determine the exact relationship, the same data was replotted as a function of the distance between the Strouhal number of the fixed-time-period method and the value from the GBVTP method. The plots show that to obtain machine zero level residuals using fixed-time-period calculations the Strouhal number specified by the user *a priori* must approximate the GBVTP value to machine




Figure 2.3: (a) The minimum residual obtained in a fixed time period NLFD calculation as a function of the distance between Strouhal number specified *a priori*. The data is plotted using a logarithmic scaling for the vertical axis only. (b) The minimum residual obtained in a fixed time period NLFD calculation as a function of the distance between Strouhal number specified *a priori* and the Strouhal number produced from a GBVTP calculation. The data is plotted using a logarithmic scaling for both horizontal and vertical axes.

accuracy. The results from this numerical experiment demonstrate the capability of the GBVTP algorithm to find this value. In addition, since there is only one minimum in the residual values, the data supports the notion that there is a unique Strouhal number associated with the exact solution of the discretized equations.

2.3 Convergence Acceleration of the Pseudo-Time Equation

With the NLFD method, solving the unsteady problem is similar to the process of solving a steady-state system since the governing equations and solution expressed in the frequency domain are time invariant. Accelerating the convergence of this stationary system of equations is achieved through a suite of methods originally developed for steady flow solvers including local pseudo-time stepping, residual averaging and multigrid. These techniques need to be modified for the NLFD method to ensure

that the convergence rates of the NLFD solver are equivalent to steady-state codes. In addition, the effects of unsteady terms in the governing equations can have a negative impact on the convergence rates of the NLFD solver. To mitigate these effects, new methods like coarse grid spectral viscosity had to be developed in the course of this research to regain the original steady convergence rates. The following sections describe the implementation of these techniques in the NLFD solver.

2.3.1 Local Time Stepping

A modified Runge-Kutta (RK) time-stepping scheme is used to advance the solution forward in pseudo time [45]. Separate coefficients are used for the dissipative and convective components of the spatial operator at each stage of the RK scheme. These are based on a set of coefficients optimized to accelerate convergence for steady-state flows. Since the pseudo-time equation is being used to solve equations 2.15 in the frequency domain, the transformations between the frequency and the time domain illustrated in figure 2.1 are applied at every stage of the RK scheme.

A varying local time step dictated by numerical stability is used in each cell. Unlike an explicit time-accurate scheme where the global time step is determined by the minimum of the time steps for all the cells, the local time stepping scheme maximizes the correction each cell can make per time step. The pseudo-time evolution does not affect the temporal accuracy of the NLFD formulation, which is only a function of the number of modes used to represent the solution and residual.

Previous methods to solve unsteady flow problems have employed both explicit [34] and implicit [50] approaches in the treatment of the diagonal terms associated with the temporal derivatives. The explicit approach bases the calculation of the temporal derivative on known values of the solution, while the implicit approach uses a future value of the solution in this calculation. Following the work of Melson *et al.* [50] this section presents a stability analysis that handles both implicit and explicit time-stepping schemes to compare their relative merits for the solution of the pseudo-time equation 2.18. Either permutation of the time-stepping scheme can be expressed by

the following set of equations:

$$\begin{aligned}
\hat{W}_k^0 &= \hat{W}_k^n \\
(1 + \gamma \alpha^m \Delta \tau_k \lambda_t) \hat{W}_k^m &= \hat{W}_k^0 - \alpha^m \left[-\gamma \hat{P}_k^{\overline{m-1}} + \hat{Q}_k^{\overline{m-1}} + \hat{R}_k^{\overline{m-1}} \right] \\
m &= 1, 2, \dots, M \\
\lambda_t &= ikV \\
\hat{W}_k^{n+1} &= \hat{W}_k^M.
\end{aligned} \tag{2.41}$$

The following list describes the different variables in this scheme.

- γ is the variable that switches the nature of the method from implicit to explicit: for explicit schemes $\gamma = 0$, for implicit schemes $\gamma = 1$.
- $\Delta \tau_k$ is the pseudo time step calculated as a function of the stability region of the scheme.
- α^m is the set of coefficients multiplying the overall unsteady residual. α_P^m , α_Q^m and α_R^m are coefficients used in the calculation of the \hat{P}_k , \hat{Q}_k and R_v terms respectively.
- \hat{W}_k^n and \hat{W}_k^{n+1} are the Fourier coefficients of the solution at pseudo-time steps n and $n + 1$ respectively.
- \hat{W}_k^m is the Fourier coefficient of the solution at stage m of the Runge-Kutta scheme with a total of M stages.
- $\hat{R}_k^{\overline{m-1}}$ is the Fourier coefficient of the transformed spatial operator. Since the spatial operator contains non-linear terms a pseudo-spectral approach has been adopted. The convective and dissipative terms are calculated at each instance of the solution in its time history. The convective fluxes are calculated using only the solution at the current stage of Runge-Kutta scheme:

$$R_c^m = \sum_{cv} \vec{F}_c(W^m) \cdot \vec{S}. \tag{2.42}$$

Depending on the α_R^m coefficients, the overall viscous flux is formed as a linear combination of the viscous flux calculated using the current stage of the solution and the overall viscous flux calculated at a previous stage of the Runge-Kutta scheme following Martinelli and Jameson [45]:

$$\begin{aligned} R_v^{\overline{m-1}} &= \alpha_R^m R_v^{m-1} + (1 - \alpha_R^m) R_v^{\overline{m-2}} \\ R_v^m &= \sum_{cv} \vec{F}_v(W^m) \cdot \vec{S}. \end{aligned} \quad (2.43)$$

The different time instances of the viscous and convective fluxes are added together to form the complete time history of the spatial operator. This term is then transformed into the frequency domain via a FFT and directly substituted into the time-stepping scheme:

$$\begin{aligned} R^{\overline{m-1}} &= R_c^{m-1} + R_v^{\overline{m-1}} \\ \hat{R}_k^{\overline{m-1}} &= \frac{1}{N} \sum_{j=0}^{N-1} R_j^{\overline{m-1}} e^{ikt_j}. \end{aligned} \quad (2.44)$$

- $\hat{P}_k^{\overline{m-1}}$ is the implicit temporal derivative of the solution expressed for each wavenumber in the frequency domain. This term is calculated in a manner similar to the evaluation of the overall viscous flux discussed in the previous subsection. Depending on the α_P^m coefficient, the overall temporal derivative is a linear combination of the temporal derivative evaluated using the solution at the current stage and the overall temporal derivative calculated at a previous stage in the Runge-Kutta scheme:

$$\hat{P}_k^{\overline{m-1}} = \alpha_P^m \lambda_t \Delta \tau_k \hat{W}_k^{m-1} + (1 - \alpha_P^m) P^{\overline{m-2}}. \quad (2.45)$$

- $\hat{Q}_k^{\overline{m-1}}$ is the explicit temporal derivative of the solution calculated for each wavenumber in the frequency domain. This term is similar to $\hat{P}_k^{\overline{m-1}}$ except that

it is based on a different set of coefficients α_Q^m :

$$\hat{Q}_k^{\overline{m-1}} = \alpha_Q^m \lambda_t \Delta \tau_k \hat{W}_k^{m-1} + (1 - \alpha_Q^m) \hat{Q}_k^{\overline{m-2}}. \quad (2.46)$$

The stability of the scheme can be addressed using a linearized analysis. Assume that the iterates of the convective and viscous residual can be represented as products of a complex eigenvalue and the solution at the same stage:

$$\begin{aligned} \hat{R}_{c_k}^m &= \lambda_c \hat{W}_k^m \\ \hat{R}_{v_k}^m &= \lambda_v \hat{W}_k^m. \end{aligned} \quad (2.47)$$

In a similar manner, gain functions for stage m of the Runge-Kutta scheme can be defined for the solution G^m , the two temporal derivatives G_P^m , G_Q^m , and the viscous residual G_R^m as:

$$\begin{aligned} \hat{W}_k^m &= G^m \hat{W}_k^0 \\ \hat{P}_k^{\overline{m}} &= G_P^m \hat{W}_k^0 \\ \hat{Q}_k^{\overline{m}} &= G_Q^m \hat{W}_k^0 \\ \hat{R}_{v_k}^{\overline{m}} &= G_R^m \hat{W}_k^0. \end{aligned} \quad (2.48)$$

Substituting the definitions for the linearized residual and gain functions into the equation for the time-stepping scheme results in a recursive equation for the different gain functions:

$$\begin{aligned} G_P^{m-1} &= \alpha_P^m \Delta \tau_k \lambda_t G^{m-1} + (1 - \alpha_P^m) G_P^{m-2} \\ G_Q^{m-1} &= \alpha_Q^m \Delta \tau_k \lambda_t G^{m-1} + (1 - \alpha_Q^m) G_Q^{m-2} \\ G_R^{m-1} &= \alpha_R^m \Delta \tau_k \lambda_v G^{m-1} + (1 - \alpha_R^m) G_R^{m-2} \\ G^m &= \frac{1 - \alpha^m (-\gamma G_P^{m-1} + G_Q^{m-1} + \Delta \tau_k \lambda_c G^{m-1} - G_R^{m-1})}{1 + \alpha^m \Delta \tau_k \gamma \lambda_t}. \end{aligned} \quad (2.49)$$

Due to the recursive nature of the equations listed above, some initial conditions are

also defined as:

$$\begin{aligned}
 G^0 &= 1 \\
 G_P^{-1} &= 0 \\
 G_Q^{-1} &= 0 \\
 G_R^{-1} &= 0.
 \end{aligned} \tag{2.50}$$

The goal is to quantify the magnitude of the overall gain $\|G^M\|$ for an M stage Runge-Kutta scheme defined using a given set of coefficients α^m , α_P^m , α_Q^m and α_R^m . This overall gain can be evaluated for each permutation of the eigenvalue trio, $\lambda_t\Delta\tau$, $\lambda_c\Delta\tau$ and $\lambda_v\Delta\tau$:

$$\lambda\Delta\tau = \lambda_t\Delta\tau + \lambda_c\Delta\tau + \lambda_v\Delta\tau. \tag{2.51}$$

It is assumed that the convective eigenvalue is purely imaginary, and the viscous eigenvalue is purely real. The eigenvalue associated with the temporal derivative has been previously defined as purely imaginary.

Figure 2.5 provides a contour plot of the gain magnitude, when it is less than one, for four different schemes. These plots are generated in the complex $\lambda\Delta\tau$ plane, by fixing $\lambda_t\Delta\tau$ and varying $\lambda_c\Delta\tau$ and $\lambda_v\Delta\tau$ over a range of values.

The upper left corner of this figure shows the stability region for an explicit scheme that does not lag the calculation of the temporal derivative at each stage. The plot is consistent with stability diagrams presented by Martinelli [45].

The upper right corner shows a stability diagram for the same coefficients as the previous case except that the temporal derivative is evaluated implicitly ($\gamma = 1$). Melson *et al.* [50] used a similar implicit technique to integrate a temporal eigenvalue with a real component. They showed that the implicit technique expands the stability region of the original scheme. Unfortunately, in the case of an imaginary component, the complex denominator acts to rotate the stability region about the origin. Consequently this technique is not suitable for the NLFD formulation because many eigenvalues associated with the spatial discretization lie close to the imaginary axis and hence outside the stability region of the implicit method.

The lower left corner presents the stability region based on a explicit scheme with a *lagged* approach to the temporal derivative. The effect is to rotate the stability contours in the opposite direction as the implicit technique effectively suffering the same deficiencies.

The lower right corner presents the stability region based on an implicit scheme with a similar *lagged* approach to the temporal derivative. At first glance, the combination of lagging and implicit handling would seem to offset each other, providing a stability region similar to the fully explicit scheme in the upper left corner. On closer inspection, the stability contours near the origin move to the left of the imaginary axis, making the scheme useless for spatial discretizations with higher order artificial dissipation.

Given the deficiencies of the implicit and lagged approaches, the NLFD code has been implemented with a explicit scheme where the temporal derivative is completely reevaluated at each step in the scheme ($\alpha_Q = [1, 1, 1, 1, 1]$). However, the set of coefficients associated with the viscous terms lag the evaluation of the spatial operator. This complicates the dataflow diagram for the pseudo-spectral approach shown in figure 2.1. A modified dataflow diagram including these lagged operations is included in figure 2.4. This diagram shows that the spatial operator, R , calculated in physical space is a function of the spatial operator at the previous stage and the spatial operator calculated using the current stage of the solution. The current iterate of the spatial operator is transformed back into the frequency domain using a FFT where it is added to frequency domain representations of the temporal derivative and the initial solution to produce a new stage of the solution.

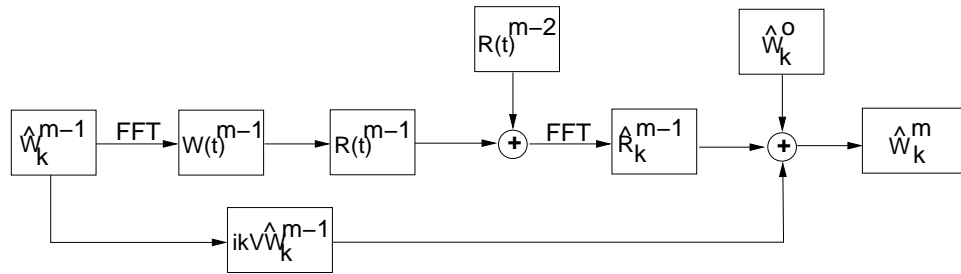


Figure 2.4: Complete dataflow diagram of the time-stepping method.

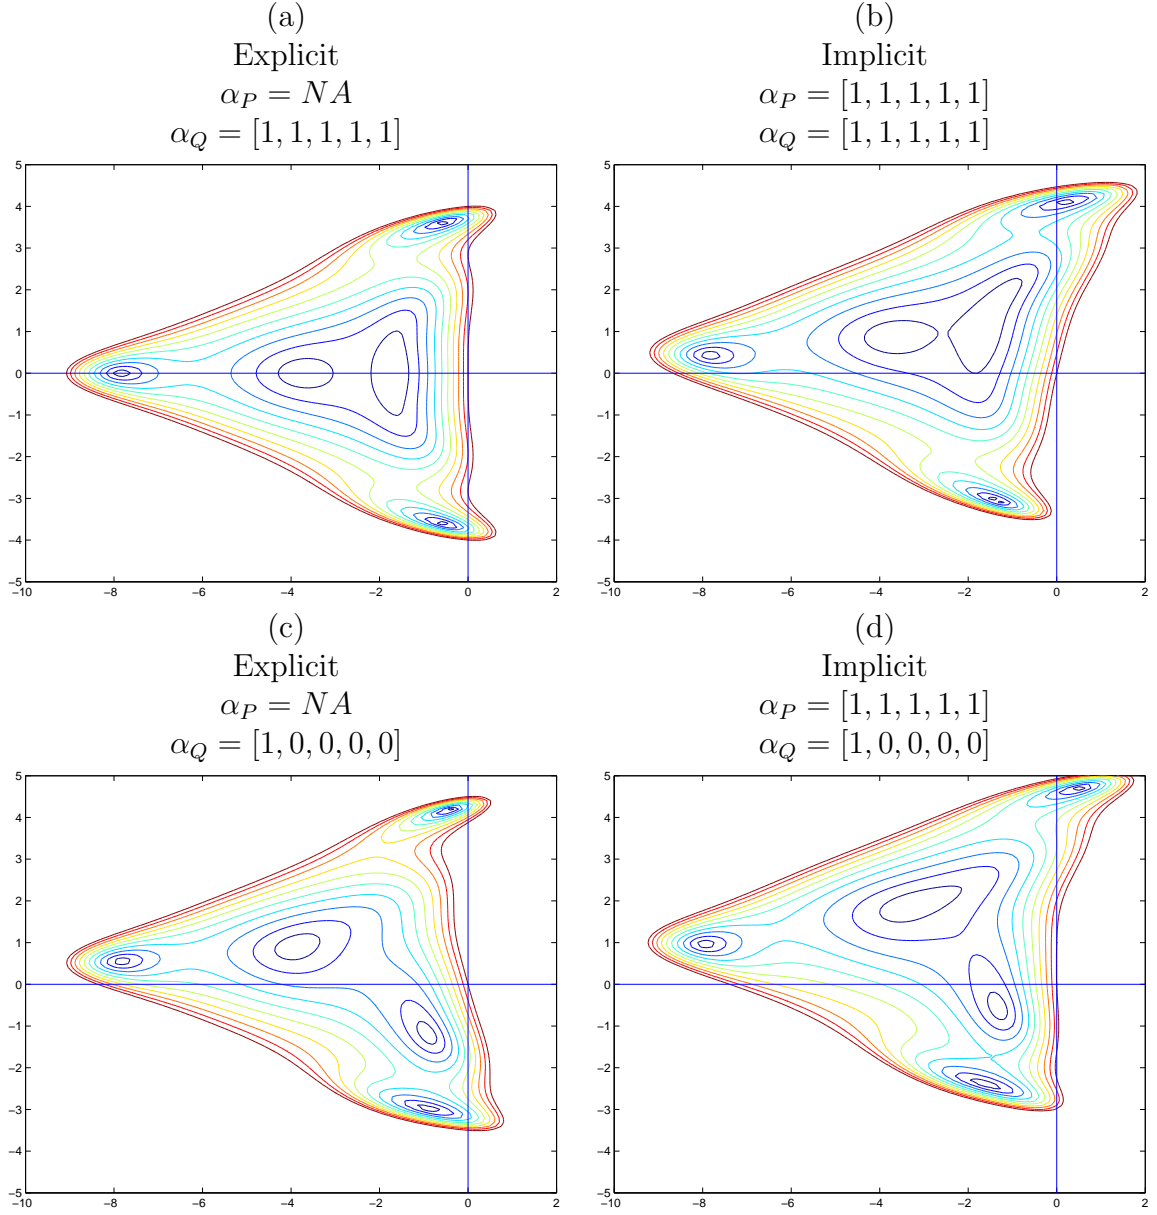


Figure 2.5: Stability diagrams of the magnitude of the overall time-stepping scheme in the complex $\lambda\Delta t$ plane. The coefficients for all above permutations: $\alpha = [\frac{1}{4}, \frac{1}{6}, \frac{3}{8}, \frac{1}{2}, 1]$ $\alpha_R = [1, 0, 0.56, 0.44]$.

2.3.2 Residual Averaging

An implicit smoothing operator is applied to the unsteady residual in a process known as residual averaging. The coefficients of this implicit smoothing operator are chosen such that larger time steps can be used while maintaining the overall stability of the numerical scheme. The additional cost of applying the operator is offset by the improved convergence rates per multigrid cycle.

One of the goals of this research is ensure that the cost of a NLFD simulation is the product of the cost of an equivalent steady calculation and the number of instances used to represent the solution. With respect to a steady computation, the NLFD code has some additional overhead due to the costs of forming the $ikV\hat{W}$ term and performing the Fast Fourier transforms on the solution and residual. If these costs are small in comparison with the cost of evaluating the spatial residual terms then the aforementioned goal can be met. To ensure this, one should avoid any numerical operations done in the unsteady code that are not done in the steady solver. With this in mind, this section proposes a residual averaging method which averages only in the spatial directions.

Inspection of the unsteady residual presented in equation 2.18 shows that the temporal derivative term is proportional to the product of the temporal frequency and cell volume. For low frequency cases on a well refined grid these terms are second order and their impact on the stability of the numerical method can be neglected. However, the multigrid process transfers the solution to a very coarse mesh where these terms may dominate. If one does not adjust the time step accordingly, the instabilities created on the coarse mesh will cause instabilities on the fine mesh ultimately corrupting the solution.

The remainder of this section will use the advection equation as a model to derive stability estimates for the residual averaging operators. In space and time a continuous form of implicit residual averaging can be written as follows:

$$\left(1 - e^{\frac{\partial^2}{\partial x^2}}\right) \Delta\tau \bar{I} = \Delta\tau \left(\frac{\partial u}{\partial t} + c \frac{\partial u}{\partial x}\right). \quad (2.52)$$

For consistency with the NLFD method, assume that both the solution and residual

can be represented by a Fourier series in time. Applying central difference operators to the spatial derivatives results in a discrete approximation to the continuous equation

$$-e\Delta\tau_k\hat{\hat{I}}_{k_{i+1}} + (1+2e)\Delta\tau_k\hat{\hat{I}}_k - e\Delta\tau_k\hat{\hat{I}}_{k_{i-1}} = -\Delta\tau_k\left(ik\hat{u}_k + c\frac{\hat{u}_{k_{i+1}} - \hat{u}_{k_{i-1}}}{2\Delta x}\right). \quad (2.53)$$

A von Nuemann analysis of the resulting equation provides the spectral footprint of the time advancement scheme. In semi-discrete form the averaged residual can be written as the product of an eigenvalue λ and the solution

$$\begin{aligned} \Delta\tau_k\hat{\hat{I}}_k &= \lambda i\hat{u}_k \\ |\lambda| &= \frac{k\Delta\tau_k + \Lambda\sin(\rho)}{1+2e(1-\cos(\rho))} \\ \rho &= \omega\Delta x \end{aligned} \quad (2.54)$$

The effects of dissipation are ignored by this analysis by only quantifying the length of this eigenvalue along the imaginary axis. After some algebra, the above equation can be rewritten in quadratic form for $\sin(\rho)$

$$(\pi^2 + 4e^2)\sin^2(\rho) + 2a\pi\sin(\rho) + (a^2 - 4e^2) = 0 \quad (2.55)$$

where,

$$\begin{aligned} \Lambda &= \frac{c\Delta\tau_k}{\Delta x} \\ \pi &= \frac{\Lambda}{|\lambda|} \\ \gamma &= \frac{k\Delta x}{c} \\ a &= \pi\gamma - 2e - 1. \end{aligned} \quad (2.56)$$

The quadratic equation yields two separate roots. The negative root is chosen for this application because as the unsteady term γ goes to zero the steady-state residual

averaging equations are regained

$$\sin(\rho) = \frac{-a\pi - 2e\sqrt{\pi^2 - a^2 + 4e^2}}{\pi^2 + 4e^2}. \quad (2.57)$$

The frequency of the discrete solution ρ that maximizes the length of the eigenvalue $|\lambda|$ is now determined by forming the gradient of the eigenvalue with respect to the frequency and setting it equal to zero

$$\frac{\partial |\lambda|}{\partial \rho} = 0 = \frac{(k\Delta\tau_k + \Lambda\sin(\rho))2e\sin(\rho) + (1 + 2e(1 - \cos(\rho)))\Lambda\cos(\rho)}{(1 + 2e(1 - \cos(\rho)))^2}. \quad (2.58)$$

After some algebra, a similar polynomial in $\sin(\rho)$ can be derived as:

$$((\gamma 2e)^2 + (1 + 2e)^2)\sin^2(\rho) + 8\gamma e^2\sin(\rho) - (1 + 4e) = 0. \quad (2.59)$$

In this case, the positive root is selected to regain the steady-state equations as the unsteady terms go to zero

$$\sin(\rho) = \frac{-4\gamma e^2 + (1 + 2e)\sqrt{(\gamma 2e)^2 + (1 + 4e)}}{(\gamma 2e)^2 + (1 + 2e)^2}. \quad (2.60)$$

Subtracting equation 2.57 from equation 2.60 produces an expression for the residual averaging coefficient that rescales the length of the eigenvalue to a user specified value that ensures numerical stability. Unfortunately the resulting polynomial is large, and a symbolic math program was needed to simplify the expression. There is more than one root for the residual averaging coefficient e ; the largest root is expressed as:

$$e = \frac{1}{4} \left(\frac{\pi^2 - (\pi\gamma)^2 + 2\pi\gamma - 1}{1 - \pi\gamma} \right). \quad (2.61)$$

Figure 2.6 plots all the roots of the residual averaging coefficient as a function of the unsteadiness of the problem quantified using the product $k\Delta\tau_k$. The largest root identified by equation 2.61 provides the most physical damping, and should be the one selected to ensure stability.

Note that the residual coefficient defined in equation 2.61 diverges to infinity as

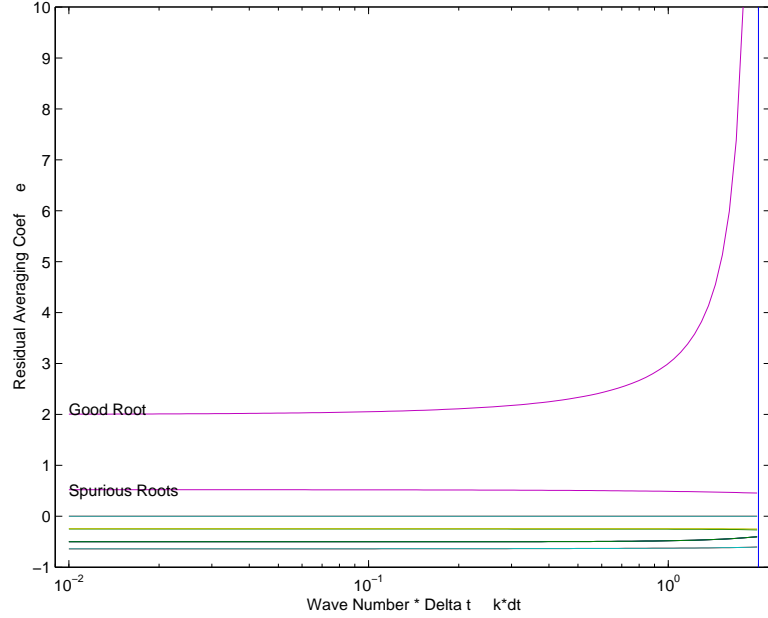


Figure 2.6: Residual averaging coefficient as a function of $k\Delta\tau_k$. These roots are calculated using a pseudo time step calculated based on the $CFL = 6$ and a rescaled eigenvalue $\lambda = 2$.

the product $\pi\gamma$ goes to one. This non-dimensionalized parameter can be simplified by replacing π and γ with the fundamental variables used in their definitions (see equation 2.56):

$$\pi\gamma = \frac{c\Delta\tau_k}{\Delta x |\lambda|} \frac{k\Delta x}{c} = \frac{k\Delta\tau_k}{|\lambda|}. \quad (2.62)$$

Setting the left hand side of the above expression to one and solving will define a maximum permissible time step, which in essence is a CFL limiter

$$\Delta\tau_k < \frac{|\lambda|}{k} \quad for \quad k > 0. \quad (2.63)$$

Without the extensive algebra, a similar result can be obtained using asymptotic methods. If the temporal derivative is large in comparison to the spatial residual then equation 2.54 is simplified to

$$|\lambda| = \frac{k\Delta\tau_k}{1 + 2e(1 - \cos(\rho))}. \quad (2.64)$$

By inspection, the frequency ρ that maximizes the length of the residual is zero. Substituting this result back into the original equation results in an expression identical to equation 2.63.

In practice the residual averaging coefficient defined in equation 2.61 is not used in the spatial operators. Instead, the steady-state values for these coefficients are used in conjunction with the time step or CFL limiter. The unsteady contribution is not smoothed or resized by the spatial operators, and by itself can exceed the stability region of the time advancement scheme. In this case, the CFL limiter rescales the length of the eigenvalue to ensure stability.

2.3.3 Multigrid

The multigrid scheme accelerates convergence to a steady state in pseudo-time by moving the solution and residual through a sequence of successively coarser grids, following an idea proposed by Fedorenko [18] and improved by Brandt [9]. Typically one starts with a solution defined on a fine mesh and projects it onto a coarser mesh with the goal of gaining a low-cost correction to the fine grid solution that eliminates the low frequency error component. Many combinations of grids have been proposed; some of the V and W cycles used in this research are depicted in figure 2.7.

The pseudo-time derivative employed by the time-stepping scheme is modified to ensure that the change in the coarse grid solution provides an effective correction to the fine grid solution

$$-\frac{\partial \hat{W}_k^m}{\partial \tau} = A_{2h}^h \hat{I}_k(W_h) - \hat{I}_k(A_{2h}^h W_h^M) + \hat{I}_k(W_{2h}^m). \quad (2.65)$$

In the above equation, A_{2h}^h is the transfer operator which moves a solution or residual from a fine mesh to a coarser one. This transfer operator is typically referred to as a restriction or aggregation [38] operator. Figure 2.8 shows the zone of influence for this operator and introduces the subscript notation used in its defining equations. The restriction operator for a steady-state code based on a volumetric average [45]




Figure 2.7: System of grids used inside of a multigrid cycle. E, evaluate the change in the flow for one step; T, transfer the data without updating the solution. Source: Jameson [37].

can be stated as:

$$W_{I,J} = \frac{V_{i,j}W_{i,j} + V_{i,j+1}W_{i,j+1} + V_{i+1,j}W_{i+1,j} + V_{i+1,j+1}W_{i+1,j+1}}{V_{i,j} + V_{i,j+1} + V_{i+1,j} + V_{i+1,j+1}}. \quad (2.66)$$

In unsteady cases where the mesh is rigidly rotating, only the solution terms are functions of time. Multiplying the above equation by e^{-ikt} and summing over all the instances in the time history of the solution will result in a restriction operator for each wavenumber used in the solution. The resulting equation mimics the original since all the operators are linear

$$\hat{W}_{kI,J} = \frac{V_{i,j}\hat{W}_{k i,j} + V_{i,j+1}\hat{W}_{k i,j+1} + V_{i+1,j}\hat{W}_{k i+1,j} + V_{i+1,j+1}\hat{W}_{k i+1,j+1}}{V_{i,j} + V_{i,j+1} + V_{i+1,j} + V_{i+1,j+1}}. \quad (2.67)$$

The restriction operators for the residual are a subset of the ones used for the solution.




Figure 2.8: Pictorial of multigrid restriction operator.

In this case, all the volumetric terms in the numerator are equal to one and the sum of the volumes in the denominator is also equal to one.

The other operator in the multigrid scheme is the prolongation operator, which transfers a change in the solution from a coarse grid to a finer one. The correction to the fine mesh is calculated using bi-linear interpolation based on the change in the solution at the closest cells within the coarser mesh. Figure 2.9 shows the zone of influence for this operator and introduces the subscript notation used in its mathematical description. In the steady case, the coefficients for the prolongation operator can be derived using a two-dimensional Taylor series expansion of the coarse grid solution

$$\Delta W_{i,j} = \frac{9}{16}\Delta W_{I,J} + \frac{3}{16}\Delta W_{I,J+1} + \frac{3}{16}\Delta W_{I+1,J} + \frac{1}{16}\Delta W_{I+1,J+1}. \quad (2.68)$$

As in the case of restriction, the linear prolongation operators can easily be expressed in the frequency domain, resulting in a similar equation for each wavenumber.




Figure 2.9: Pictorial of multigrid prolongation operator.

2.3.4 Coarse Grid Spectral Viscosity

During tests of the multigrid solver it was noticed that the convergence rate of the solver exhibits a weak dependence on the frequency of the unsteady solution: as this frequency increases the convergence rate decreases. Noting equation 2.69, the effect of frequency on the temporal derivative is obvious. It can be argued that if the convergence rate of the steady problem ($k = 0$) is considered optimal, then any increase in frequency moves the set of equations away from the steady-state solution and negatively impacts convergence. In order to mitigate these effects, additional terms have been added to the coarse grid residual:

$$V \frac{d\hat{W}_k}{d\tau} + ikV\hat{W}_k + \hat{R}_k = \varepsilon_N V (ik)^{2s} \hat{W}_k. \quad (2.69)$$

Terms of this type were originally introduced in spectral viscosity methods proposed by Gelb and Tadmor [19, 62]. The intent of these schemes was to add even powered derivatives of the solution to damp higher frequencies while maintaining spectral accuracy. However, in our approach these additional terms only affect the residual on

coarser meshes in the multigrid process, and do not affect the converged answer on the fine mesh. By damping high frequency errors they improve the projected coarse mesh correction to the finest solution.

An analysis similar to that proposed in section 2.3.1 was performed to determine the gain of the time-stepping scheme as a function of the spatial frequencies in the discretized solution. For the purposes of this analysis the spectral footprint of the overall discrete operator is assumed to be

$$\lambda\Delta\tau = \lambda_t\Delta\tau - \frac{3}{8}(1 - \cos(\rho))^2 + 2.5\sin(\rho)i. \quad (2.70)$$

This spectral footprint is drawn on the stability diagrams for an explicit scheme in the left-hand plot of figure 2.10. The blue line is calculated using a $\lambda_t\Delta\tau = \frac{i}{2}$ which is representative of an NLFD scheme without spectral viscosity. The red line is calculated with $\lambda_t\Delta\tau = -\frac{1}{2} + \frac{i}{2}$ corresponding to an NLFD with spectral viscosity. Obviously, the addition of spectral viscosity shifts the ellipses farther to the left in the complex $\lambda\Delta\tau$ plane. The gain along the path of the ellipse is provided in the right-hand plots within figure 2.10. The data shows that the shift in the spectral footprint induced by the spectral viscosity lowers the gain of the time-stepping scheme for the majority of frequencies in the solution. Lower gains increase the damping and improve the efficiency of the multigrid solver. However, the stability region of the time-advancement scheme limits the magnitude of this shift, which is proportional to the magnitude of the spectral viscosity term.

The above results are for a single value of spectral viscosity relative to the other components of the spectral footprint. Given that this approach affects only the convergence and not the final answer, a number of schemes of different order, with different coefficients and cutoff frequencies were tested. The cutoff frequency represents the lowest frequency to which the spectral viscosity term was applied; typically the fundamental harmonic. Within limits dictated by stability, the lower order dissipation schemes, including the zeroth order ($s = 0$), outperformed the higher order approaches. However, a uniform approach of equal efficiency for all the different test cases was not found, and some level of tuning was required to maintain optimal

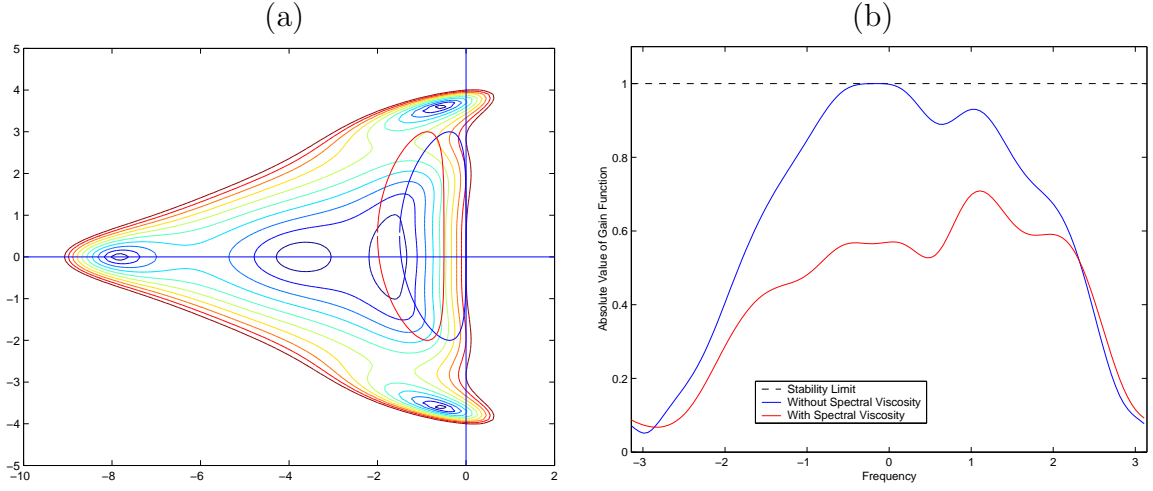


Figure 2.10: (a) Stability diagram of the explicit time-stepping scheme. The red and blue ellipses are the spectral footprint of the discretization with and without spectral viscosity respectively. (b) The gain of the time-stepping scheme along the paths of the spectral footprint.

convergence rates.

To illustrate the effects of this method on a NLFD calculation of the unsteady Euler equations, a convergence study was performed using a test case which highlights the dependence between convergence and frequency. The case uses a 193x49 C-mesh around a 64A010 pitching airfoil with boundary conditions approximating Davis's experiment described by table 4.3. The only deviation from the parameters defining this case are an increase in the dynamic angle of attack to $\pm 3.0^\circ$. This was chosen to increase the nonlinearities in the flow field ultimately providing a more challenging test case.

Figure 2.11 shows the magnitude of all components of the unsteady residual as a function of the multigrid cycle. The plots in this figure show the residual convergence for solvers with and without coarse grid spectral viscosity. Overplotted on these graphs is the convergence rate of the NLFD solver for the analogous steady problem lacking the pitching motion of the airfoil. Without spectral viscosity, the convergence rate of the unsteady NLFD solver is not equivalent to the steady-state case. However, the left-hand figure shows that the optimal convergence rate of the steady solver can

be reclaimed for unsteady calculations through the use of spectral viscosity.

In the current formulation the coarse grid spectral viscosity does not affect the fine grid solution. This is confirmed by figure 2.12, which shows the difference between the solutions with and without spectral viscosity. The subplots are generated for each solution variable at a consistent point in time in the rotation of the airfoil. The magnitude of the differences are consistent with the accuracy limitations of the floating point math used in each simulation; effectively indicating that the difference between the solutions is *machine zero*.




Figure 2.11: Magnitude of all components of the unsteady density residual as a function of the multigrid cycle. (a) NLFD solver with coarse grid spectral viscosity (b) NLFD solver without coarse grid spectral viscosity.

2.3.5 Parallelization

Except for the residual averaging operations, the explicit nature of the time advancement scheme has facilitated the parallelization of this code. Following the work of Alonso *et al.* [4], the current solver decomposes the grid into sub-blocks which are distributed on different processors. The problem decomposition is spatially based with all the instances of the solution at a given point held on the same processor. A dual

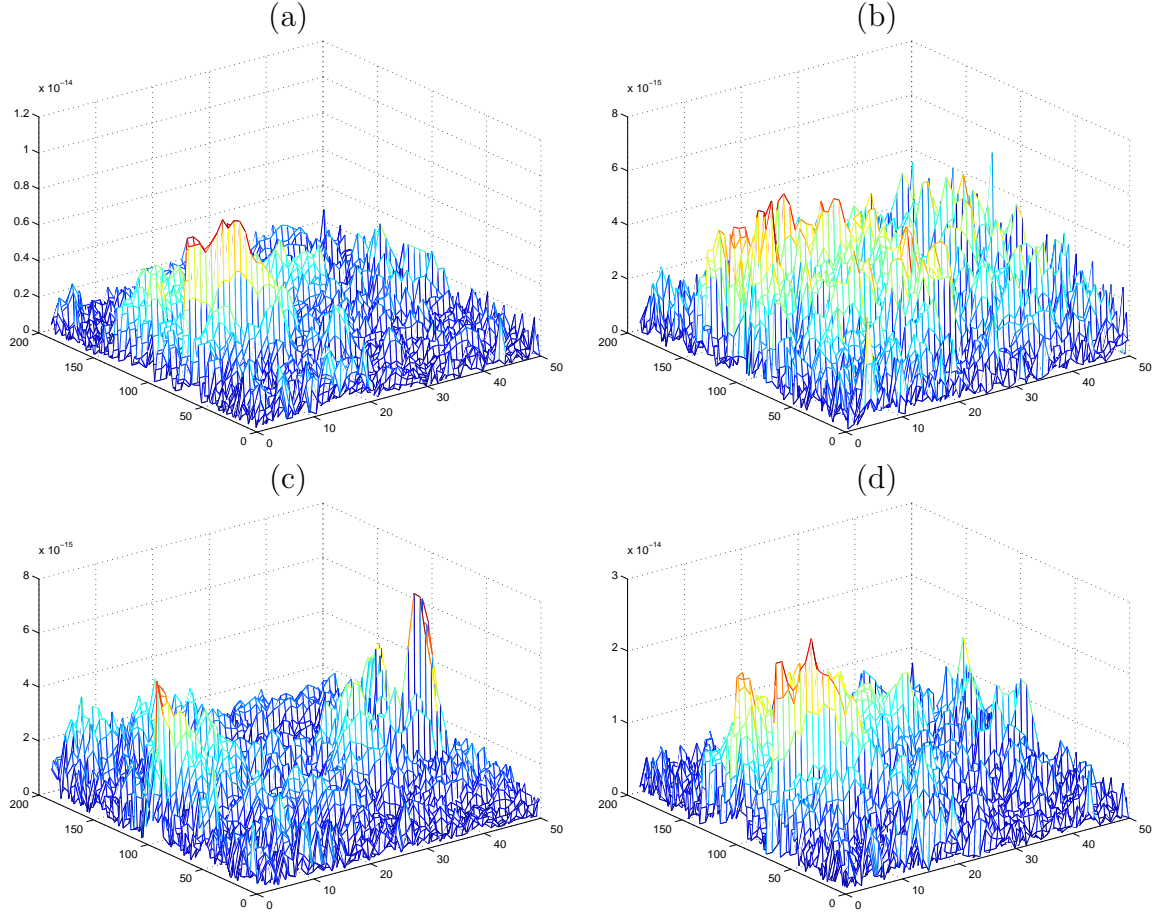


Figure 2.12: Magnitude of the difference between solutions computed with and without spectral viscosity shown at time $t = \frac{6}{7}T$. (a) ρ (b) ρu (c) ρv (d) ρe .

halo scheme serves to locally cache the state vector from neighboring grids. MPI libraries implement the actual communication between processors. Both the amount of work and inter-block communication scale linearly with the number of wavenumbers the user specifies *a priori*. This yields ratios of computation to communication that are consistent with steady-state codes using similar spatial decompositions. Consequently, the parallel efficiencies achieved in steady-state solvers can be obtained for NLFD methods.

Chapter 3

Cylindrical Vortex Shedding

This chapter addresses the first of three test cases which have been used to evaluate and also validate the NLFD method. Simulations of vortex shedding in low Reynolds number flows past cylinders demonstrate the efficiency of the NLFD method in representing complex flow field physics with a limited number of temporal modes. Since the shedding frequency is unknown *a priori*, vortex shedding is also a good test case for the application of Gradient-Based-Variable-Time-Period methods.

3.1 Test Case Description

Laminar vortex shedding occurs behind a circular cylinder over a range of Reynolds numbers between 40 and 194 [73]. Unsteady vortices alternately shed from the leeward surface of the cylinder forming an ordered wake known as the von Kármán vortex street. Figure 3.1 illustrates this phenomena with a picture of an experiment using smoke to visualize the vortex structure.

Laminar vortex shedding has been investigated in numerous modern experimental studies. Despite the accuracy of modern experimental equipment, the value of the universal Strouhal-Reynolds number relationship varied by approximately 20 percent between experiments conducted at different laboratories. This variation has been attributed to the existence of oblique shedding modes that introduce three-dimensional effects in a flow field assumed to be invariant along the length of the cylinder [69].



Figure 3.1: Smoke visualization of laminar vortex shedding in the wake of a cylinder at a Reynolds number of 140. Experiment conducted by Taneda [63] and reproduced by Van Dyke [17].

Figure 3.2 captures a *top down* view from two experiments using flow visualization techniques to observe the difference between the shedding modes. Using the experimentally observed oblique angle, Williamson proposed an equation to correct oblique data to yield parallel results. More importantly he demonstrated an experimental apparatus to correct the end boundary conditions in order to force parallel shedding modes. Experiments after 1988 that purposefully induce parallel shedding, reduce the original scatter in the data to 1 percent [68]. Of note is an experiment by Norberg [57] which used different end boundary conditions to ensure parallel shedding, yet yielded remarkable agreement with Williamson’s data.

This discussion is relevant because independent experimental data is compared to the two-dimensional numerical results provided by this research. To make a consistent comparison the physical wake found in the experiment should be approximately two-dimensional. Given the work of Williamson in this field, his experimental data will be used exclusively for the comparisons presented later in this section.

In addition to comparisons with experimental results comparisons have also been made with the numerical results of Henderson [26]. His code uses spectral elements with an eighth order basis. To verify the accuracy of his results, Henderson performed




Figure 3.2: Top view of cylinder wake experiments using flow visualization techniques to identify parallel and oblique shedding modes. (a) Parallel mode. (b) Oblique modes. Source: Williamson [68].

a spatial and temporal resolution study for flows at higher Reynolds numbers that included additional scales of motion. In comparison, the NLFD code uses a lower order spatial discretization with a limited number of temporal modes that purposely under-resolve the solution. A comparison of results between these two codes confirm the accuracy of the NLFD method and highlight the benefits to be gained from higher order methods and their associated cost.

3.2 Computational Grid

A numerical resolution survey was conducted using several different grids. Table 3.1 provides a list of the grids and their pertinent features. All the grids have a circular topology with adjacent radial grid lines subtending arcs of equal magnitude. An exponential function stretches the grid in the radial direction. This stretching function and the quadratic relationship between the cell area and radius, combine to create a mesh with poor spatial resolution in the mid and farfield. This coarse spatial resolution in combination with the artificial dissipation schemes employed by this research act to quickly diffuse the structure of the vortices as they travel downstream of the cylinder. Within six chords of the cylinder the vortices contained in the most

Circumferential Points	Radial Points	Boundary Layer Points	Boundary Distance (Chords)
129	65	12	200
193	81	15	200
257	129	19	200
385	161	23	200

Table 3.1: Pertinent features of the grids used in the cylinder simulations. The number of points in the boundary layer was calculated using the velocity profiles at the 12 o'clock position on the cylinder at a Reynolds number of 150.

accurate numerical solutions showed little definition, while figures 3.1 and 3.2 picture experimental conditions where well-formed vortices exist much farther downstream.

A poorly resolved farfield will have some effect on the prediction of statistics like Strouhal number and base suction coefficient. However, the results will show that these parameters can be accurately predicted using the grids presented in this section. In addition, these simulations take advantage of the diffusive character of the farfield of the grid and implement relatively simple boundary conditions. The degree to which these conditions reflect energy is no longer a critical parameter given the *de facto* sponge created by diffusion and the extended distance of the boundary.

3.3 Cylinder Results

This section presents results from NLFD calculations of the cylinder test case and compares them with similar data from experiments and other numerical computations. Temporal and spatial resolution surveys are included to show the convergence these results exhibit as a function of diminishing numerical error. Residual convergence data related to solver performance will not be presented for this case. A perfectly symmetric spatial discretization and initial conditions produce a symmetric wake after driving the residual close to machine zero. Depending on the grid and Reynolds number, the random effects of machine error may be enough to stimulate the instability in the shear layer generating a von Kármán vortex street. It was more

efficient to stop the solver, manually perturb the solution, and restart the solution process than to wait for instabilities to arise from a machine-accurate symmetric solution. Because of this manual intervention, the residual histories are functions of the relative time and magnitude of the perturbation process. This made comparison to other codes difficult.

The temporal resolution survey covers four solutions using one, three, five and seven time varying harmonics. The flow field, as it exists in nature, obviously contains a larger number of modes. During the computation, these unresolved modes transfer energy back into lower frequencies in a corrupting process known as aliasing. This effect is a function of the magnitude of the unresolved modes, which are unknown in the absence of an exact solution. This resolution survey looks at solution changes with varying resolution in order to identify an answer that is independent of aliasing affects.

The spatial resolution survey employs the four grids presented in table 3.1. The physics of unsteady flows dictate that the spatial and temporal modes are coupled. Increasing the spatial resolution resolves additional spatial modes corresponding to an increased number of temporal modes that affects the aliasing issues addressed earlier. To ensure a systematic search of the solution space, the survey computes every permutation of the four temporal and four spatial discretizations.

The resolution survey is performed at 10 equally spaced Reynolds numbers within the range of Reynolds numbers admitting laminar vortex shedding. The results show that the relative magnitude of the higher harmonics grows with increasing Reynolds number. The changing distribution of energy between modes affects aliasing and further tests the capabilities of the NLFD method to accurately resolve the solution.

To summarize the above discussion, table 3.2 identifies all values of the three different parameters used in these surveys. Including all the permutations, 160 different solutions were calculated. A byproduct of this three-dimensional parameter survey is that a large number of solutions will statistically give a better estimate of accuracy; eliminating any random effects that can occur with a single point comparison.

In order to evaluate the effect of the GBVTP algorithm two separate sets of results are provided. The first set uses a fixed-time-period algorithm where the period of

Variable	Values
Temporal Resolution	1,3,5,7
Grid Resolution	129x65,193x81,257x129,385x161
Reynolds Number	60,70,80,90,100,110,120,130,140,150

Table 3.2: Values of parameters used in the three-dimensional parametric survey, which combines both temporal and spatial resolution surveys over a range in Reynolds number.

the fundamental harmonic is specified *a priori* based on the experimental data of Williamson *et al.* [69, 75, 74, 70]. The second set of data uses the same experimental data as the initial guess to the shedding frequency, and employs the GBVTP method to iteratively determine the exact fundamental frequency for the discretized equations. A comparison of the base suction coefficient extracted from these two data sets will quantify the relative merits of the GBVTP method.

3.3.1 Variable-Time-Period Results

This section presents results from NLFD calculations using the GBVTP method. A comparison between variable-time-period and fixed-time-period methods is presented in a subsequent section, so only contour and Strouhal number data will be provided here. Given the large number of solutions contained in this survey, this section does not provide contour plots for every computation. Instead, figures 3.3 and 3.4 provide entropy contours at the extremes in temporal and spatial resolution showing the qualitative difference in the solution with changing numerical accuracy. Because the energy contained in the higher harmonics increases with increasing Reynolds number, these contours have been provided for the most challenging case; a Reynolds number of 150.

The diffusiveness of the O-mesh topology with increasing radius is evident from the plots. This is epitomized by the coarse mesh, where significant diffusion of the vortex structure occurs after only four chords. These vortices show no well defined small scale features like a braid region, and large scale features like the overall shape of the vortex are poorly defined.

Of interest is the change in the single harmonic solution between the coarse and fine grids. One would expect to see an improvement in the qualitative features as the grids are refined, but in this case the opposite is true. Although the coarse grid results show the effects of dissipation, the vortex structure is largely intact. The fine grid case exhibits spurious features including a split of the vortex during the *roll up* process. These finer spatial resolutions support more spatial modes, which correspond to an increased number of temporal modes; heightening the effects of aliasing. Instead of improving the answer, the additional spatial resolution exacerbates the corruptive process.

This section also provides a more quantitative comparison of the numerical and experimental results based on Strouhal number. Given the three dimensionality of this parametric study, these statistics are provided in a number of different formats. Figure 3.5 provides the Strouhal frequency statistics as a function of Reynolds number. Each subplot presents results from a different grid while each line represents the locus of solutions computed at a single temporal resolution over a range in Reynolds numbers. For comparative purposes, Williamson's [69, 72, 70] experimental data is overplotted on the graphs with circles. At the low end of the Reynolds number range, several of the coarse grid calculations did not support an unsteady solution. These coarser grids have substantial amounts of artificial dissipation which effectively lowers the Reynolds number of the computation. This shift can move a computation from a Reynolds number supporting unsteady flow into a regime where only steady symmetric separation is admitted.

For consistency with William's experimental report [71] the following function is used to define the Strouhal/Reynolds number relationship for the numerical data:

$$S_t = A + \frac{B}{Re} + CRe. \quad (3.1)$$

For every spatial and temporal resolution, ten solutions were computed over a range of Reynolds numbers. For each numerical solution, the Strouhal and Reynolds numbers form a dependent pair that can be used to solve for the coefficients A,B,C in equation 3.1. The coefficients are determined by solving a system of equations (shown by

	129x65 Grid	193x81 Grid	257x129 Grid	385x161 Grid
1 Harmonic	50.3	44.7	31.1	19.3
3 Harmonic	68.5	33.5	8.7	3.3
5 Harmonic	68.3	33.6	8.5	2.9
7 Harmonic	68.3	33.6	8.6	2.9

Table 3.3: Shift in Reynolds number between the NLFD calculations and the experimental dataset for all permutations of temporal and spatial resolution.

equation 3.2) to minimize the sum of the square of the error between the Strouhal data and equation 3.1 evaluated at each of the Reynolds numbers:

$$\begin{bmatrix} \sum_{i=1}^N \frac{1}{R_{e_i}^2} & \sum_{i=1}^N \frac{1}{R_{e_i}} & N \\ \sum_{i=1}^N \frac{1}{R_{e_i}} & N & \sum_{i=1}^N R_{e_i} \\ N & \sum_{i=1}^N R_{e_i} & \sum_{i=1}^N R_{e_i}^2 \end{bmatrix} \begin{bmatrix} B \\ A \\ C \end{bmatrix} = \begin{bmatrix} \sum_{i=1}^N \frac{S_{t_i}}{R_{e_i}} \\ \sum_{i=1}^N S_{t_i} \\ \sum_{i=1}^N S_{t_i} R_{e_i} \end{bmatrix}. \quad (3.2)$$

Using the coefficients from the curve fit process of the numerical data, a second equation is proposed to investigate the Reynolds number shift between the experimental and numerical results

$$S_t = A + \frac{B}{R_e + D} + C(R_e + D). \quad (3.3)$$

The shift operator D , can be chosen such that the error between the curve fit equation (using the coefficients A, B, C derived from the numerical data) and the experimental results is minimized. This parameter represents a statistically averaged Reynolds number shift between the experimental and numerical datasets, and is provided in table 3.3.

Because of the limited number of converged solutions, the data for the coarsest mesh is suspect. In general, the Reynolds number shift decreases by a factor of four for every doubling in the grid size. This shift can be attributed to the truncation error of the baseline discretization, which is second order accurate. This error component has a dissipative influence lowering the effective Reynolds number of the numerical solution.

Spatial Convergence

The same data provided in figure 3.5 can be replotted as a function of spatial resolution. Figure 3.6 shows Strouhal data versus the number of points in a grid at four different Reynolds numbers (60,90,120,150). Most importantly, the numerical data shows convergence to the experiment as the grid spacing is decreased. However, the Strouhal number results continue to change even on the finest grids reinforcing the results presented in table 3.3. Unfortunately, the spatial resolution survey has not converged to an answer independent of grid size.

Temporal Convergence

This subsection replots the same data from previous sections except that here it is represented as a function of temporal resolution. Figure 3.7 shows the Strouhal number as a function of the temporal resolution at four different Reynolds numbers (60,90,120,150). Again, the numerical data shows convergence to the experiment as the temporal resolution is increased. In fact, most of the data shows little variation after using just three time-varying modes. In a small number of cases, the data computed at the highest Reynolds number and spatial resolution shows a small improvement when using more than three modes.

This variation can be explained by the data presented in figure 3.8. The L^2 norm of the energy within the solution for each harmonic defined is quantified on the vertical axis of this figure by the following equation:

$$\|E_k\| = \left[\int_{\Omega} |\hat{\rho} e_k|^2 dV \right]^{\frac{1}{2}}. \quad (3.4)$$

The aforementioned figure plots this norm as a function of the Reynolds number for the four different grids used in this survey. These plots are based on data compiled from the seven mode calculations which were the most physically representative solutions generated. They show that the relative energy in the higher harmonics increases as the Reynolds number increases. At the upper end of the Reynolds number range, the unresolved modes alias more energy back into the lower modes leading to a variation in the predicted Strouhal number between the three and five mode solutions.

At the lower end of the Reynolds number range, the higher modes carry substantially less energy and fewer modes are required to accurately predict Strouhal number and other statistics.

3.3.2 Fixed-Time-Period Results

This section compares the results of calculations using variable- and fixed-time-period methods to show their relative merits. Two sets of data have been generated. The first set consists of all the spatial, temporal and Reynolds number permutations using the GBVTP method outlined in section 2.2.4. The second set contains the exact same permutations as in the first, except that the time period is held fixed throughout the calculation. The Strouhal value for these calculations is derived from a curve fit (the expression is provided in equation 3.1) based on Williamson’s experimental data.

Because the fixed-time-period calculations use the experimental Strouhal data, a comparison between these two data sets based on Strouhal number would simply replicate the results presented in section 3.3.1. Instead the comparison will be based on the value of the base suction pressure coefficient, which is the negative mean pressure coefficient at the far leeward edge of the cylinder.

Figure 3.9 shows the base suction coefficient as a function of Reynolds number. Each subplot presents results from a different grid, while each line connects all the data points calculated using the same temporal resolution over a range in Reynolds numbers. In general, the difference between the NLFD and experimental datasets decreases as the Reynolds number decreases. Surprisingly, this difference on the finest mesh is smaller than the difference between the Henderson results and the experimental data over the entire range in Reynolds numbers.

In general, the fixed-time-period results at a given spatial and temporal resolution are farther from the experimental results than similarly resolved variable-time-period results. The GBVTP algorithm effectively moves the numerical results closer to the experimental data representing a positive influence on the accuracy of the NLFD method.

Spatial Convergence

The same data provided in figure 3.9 can be replotted as a function of spatial resolution. Figure 3.10 plots the base suction coefficient versus the number of points in the grid at four different Reynolds numbers (60,90,120,150). The spatial resolution study shows that unlike the Strouhal results, additional increases in grid size will have a negligible effect on the suction coefficient data.

The plots also shows that the positive effect of the GBVTP algorithm is most noticeable on the coarser meshes. For example, the coarsest grid results at a Reynolds numbers of 120 and 150, shows that a single mode GBVTP calculation outperforms any temporal resolution calculated using a fixed-time-period method. The difference between fixed- and gradient-time-period results diminishes as the spatial resolution is improved and can be considered negligible on the finest grid case.

Temporal Convergence

This subsection replots the same data as in the previous sections, except here it is presented as a function of temporal resolution. Figure 3.11 plots the base suction coefficient as a function of the temporal resolution at four different Reynolds numbers (60,90,120,150). Consistent with the Strouhal results, there is a slight variation between the three and five mode results on the finest mesh at the highest Reynolds number. The cause of this variation has already been discussed in section 3.3.1. In general, the base suction coefficient has converged to engineering accuracy using just three time varying harmonics in the representation of the solution.

3.4 Summary

The NLFD method has been successfully applied to the unsteady laminar Navier-Stokes equations for low Reynolds number flows behind cylinders. Grid resolution surveys show that spatial convergence was achieved for the base suction coefficient results but not for the Strouhal number results. The temporal resolution surveys show that predictions to engineering accuracy can be achieved using between one

and three time varying modes. Surprisingly, the NLFD results are closer to the experimental data than the Henderson results calculated using a very high order method. The Gradient-Based-Variable-Time-Period method (GBVTP) predicted the Strouhal number of the discrete equations, and in doing so improved the quality of the solution over fixed-time-period solvers.

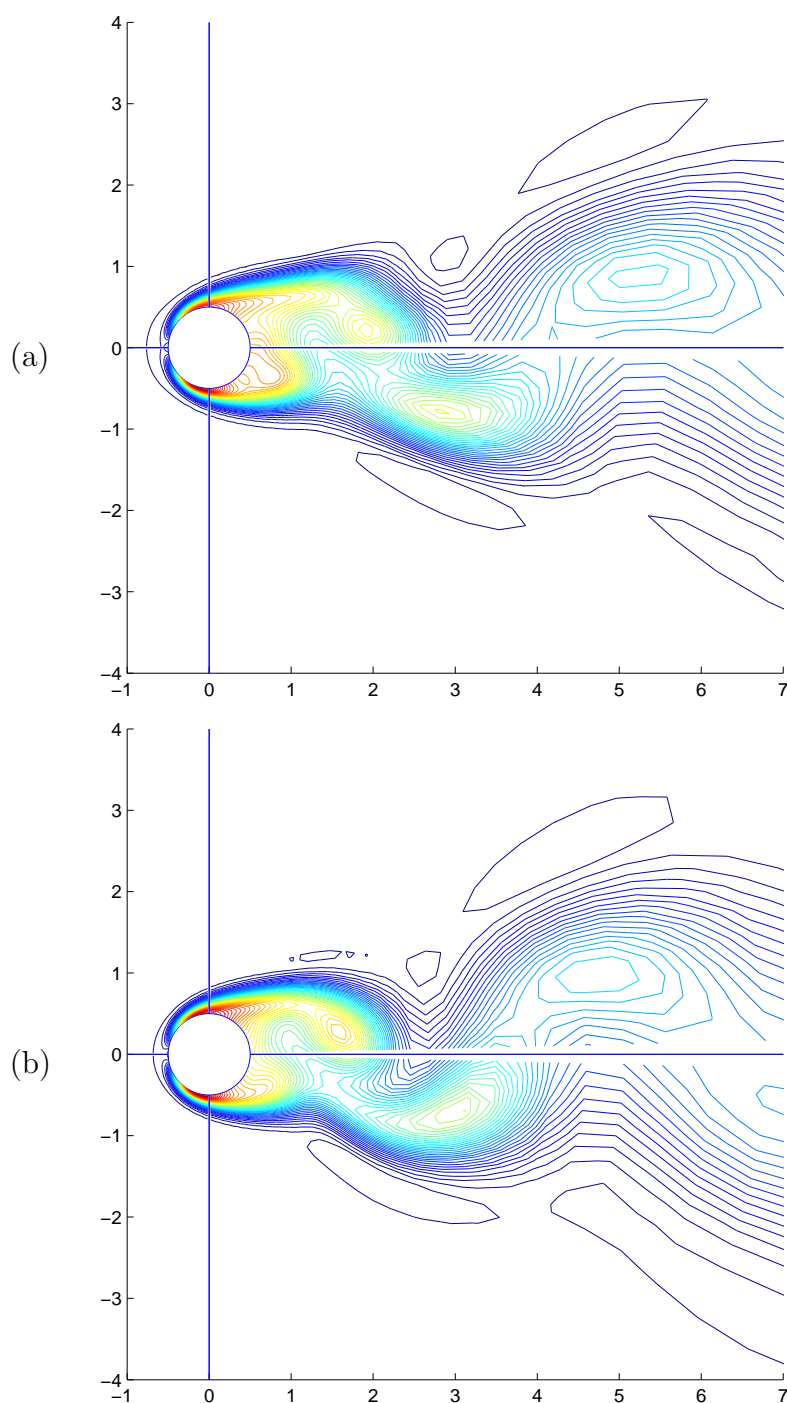


Figure 3.3: Contours of entropy at a Reynolds number of 150 utilizing the 129x65 grid. (a) solution using 1 time varying harmonic. (b) solution using 7 time varying harmonics.

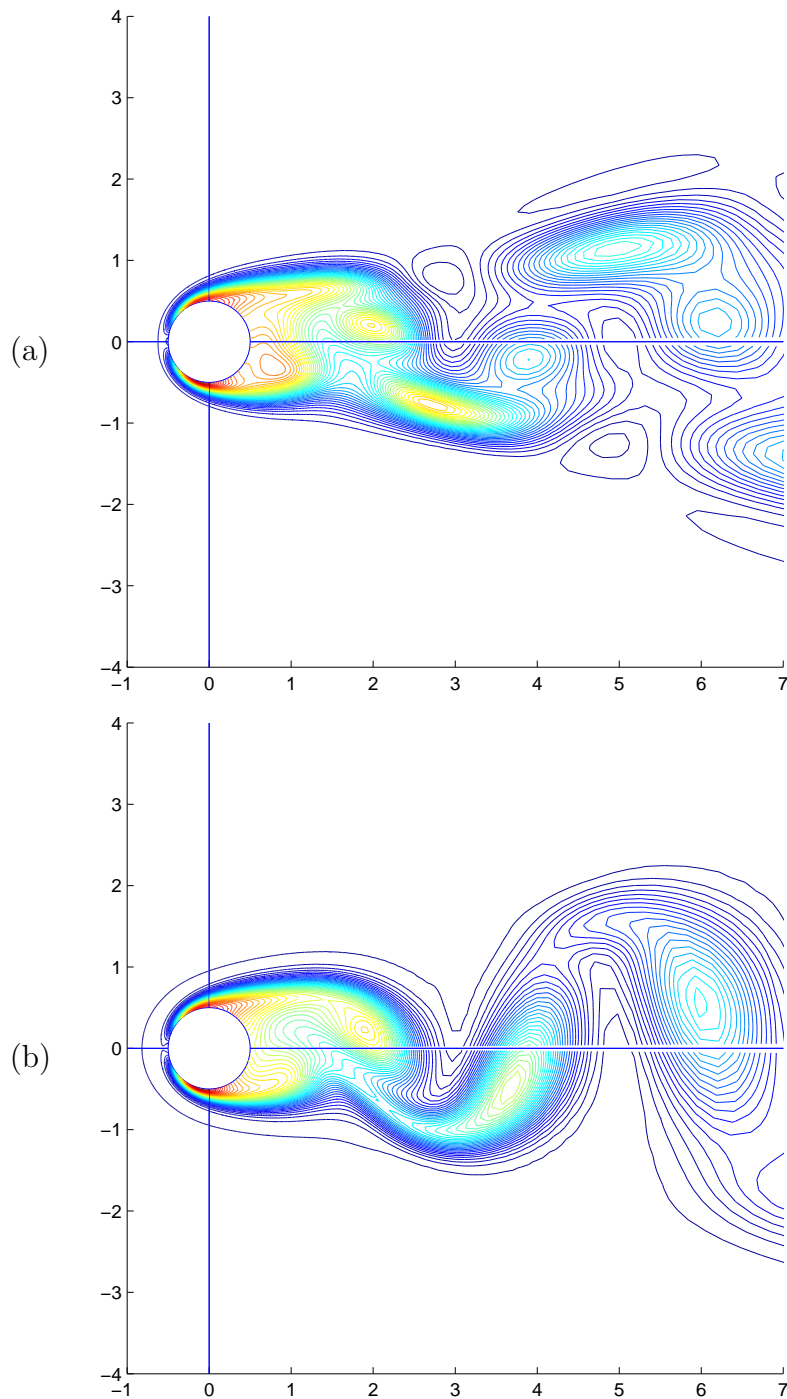


Figure 3.4: Contours of entropy at a Reynolds number of 150 utilizing the 385x161 grid. (a) solution using 1 time varying harmonic. (b) solution using 7 time varying harmonics.

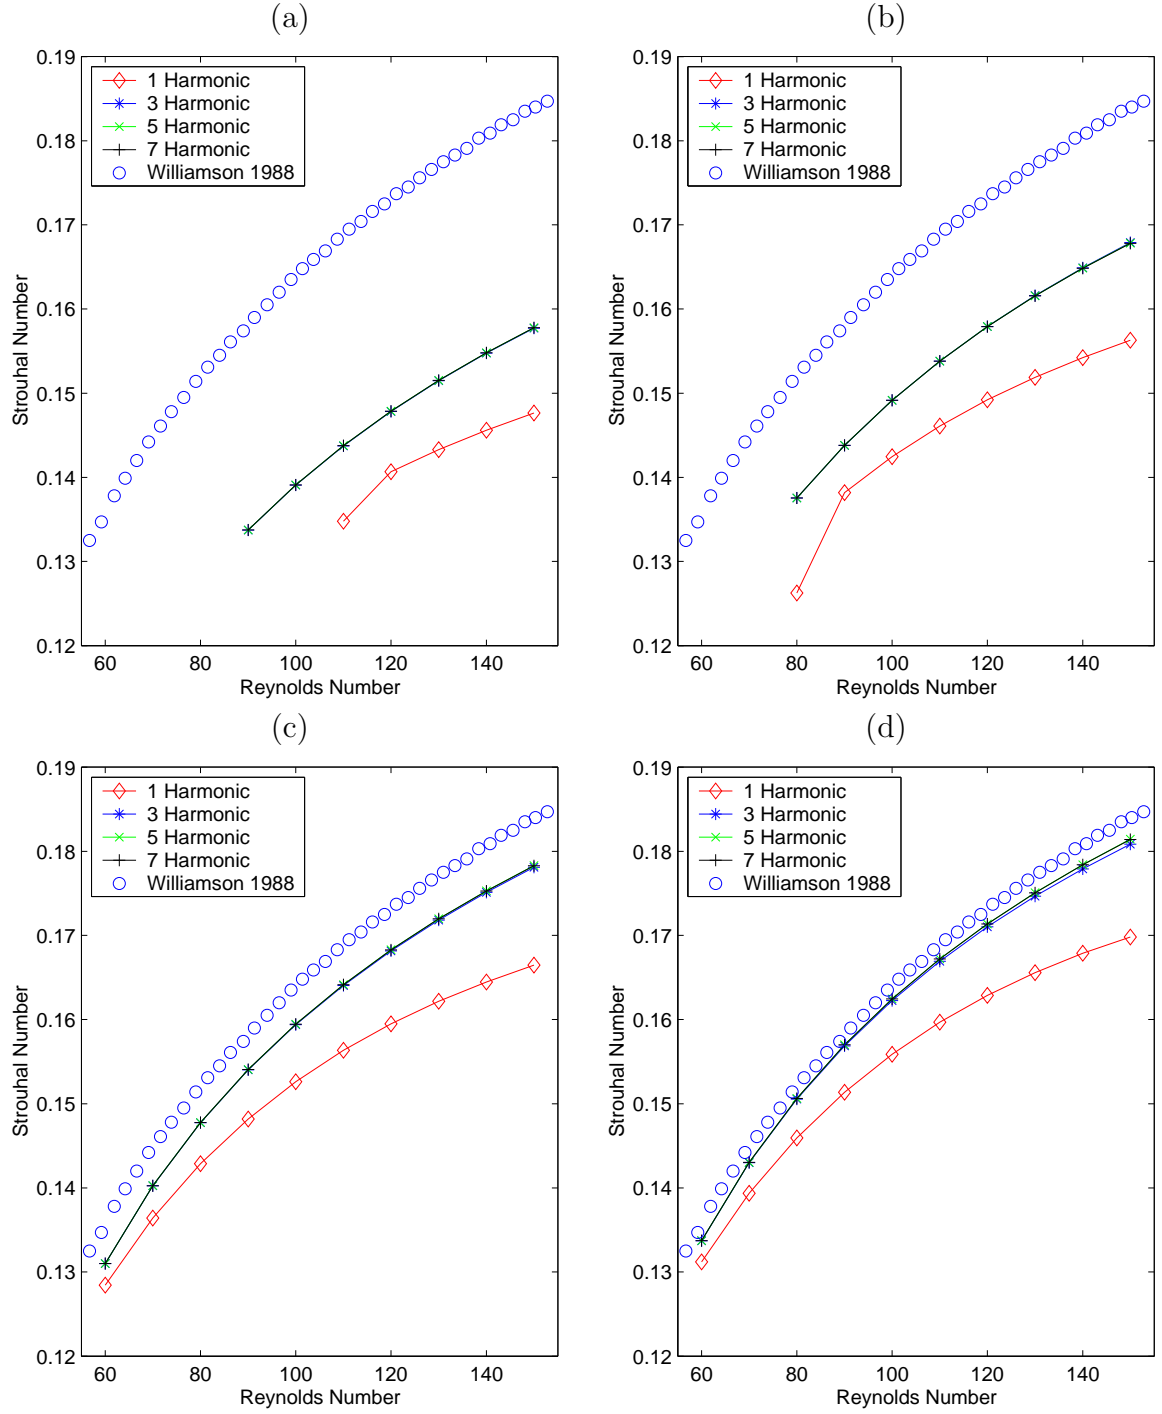


Figure 3.5: Strouhal number as a function of Reynolds number plotted for different temporal resolutions. (a) 129x65 grid (b) 193x81 grid (c) 257x129 grid (d) 385x161 grid.

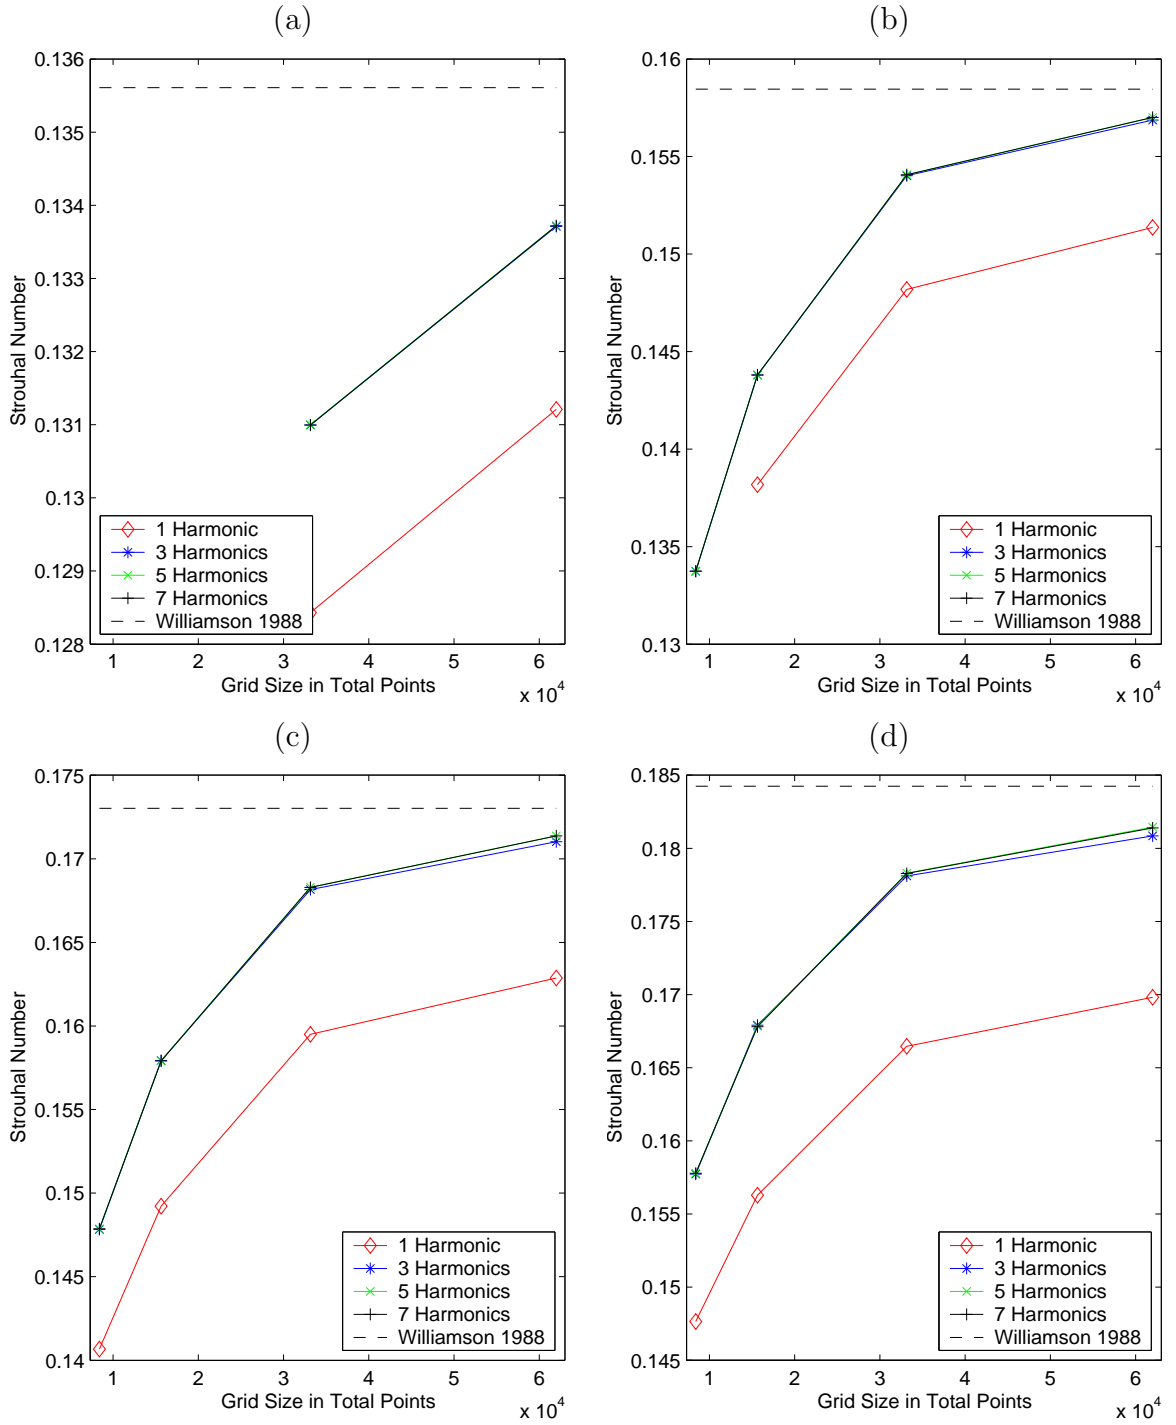


Figure 3.6: Strouhal number as a function of the spatial resolution plotted for different temporal resolutions. (a) $R_e = 60$ (b) $R_e = 90$ (c) $R_e = 120$ (d) $R_e = 150$.

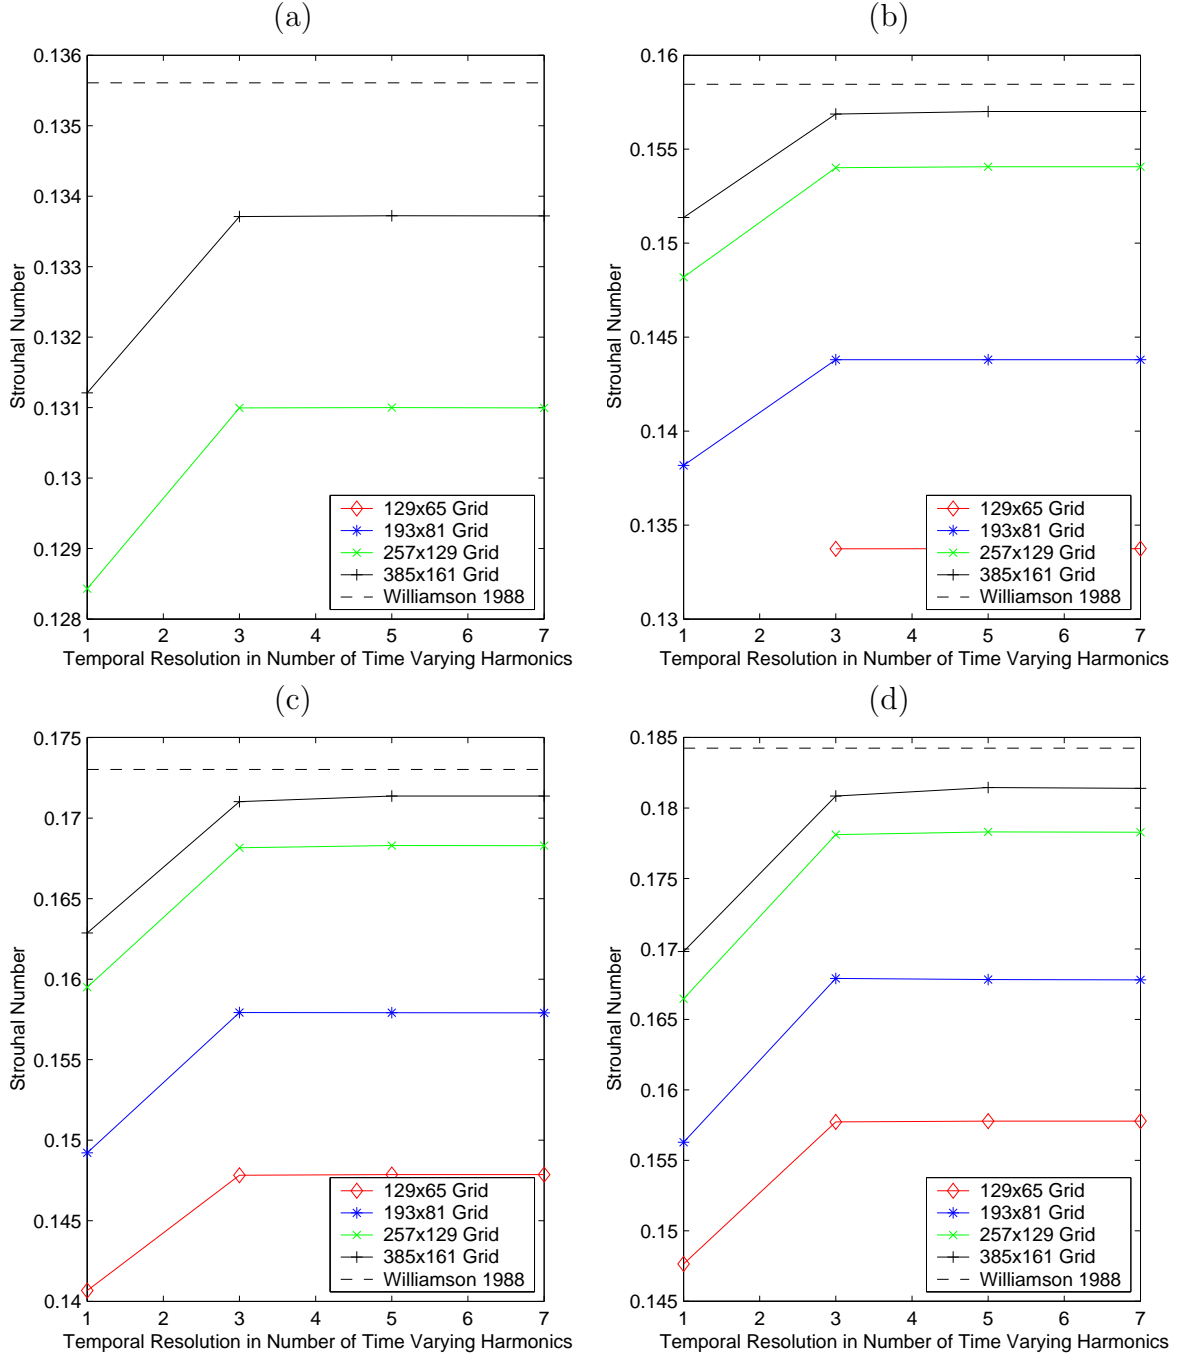


Figure 3.7: Strouhal number as a function of the temporal resolution plotted for different grid sizes. (a) $R_e = 60$ (b) $R_e = 90$ (c) $R_e = 120$ (d) $R_e = 150$.




Figure 3.8: L^2 norm of the product of density and energy for each wavenumber as a function of Reynolds number. (a) 129x65 grid (b) 193x81 grid (c) 257x129 grid (d) 385x161 grid.

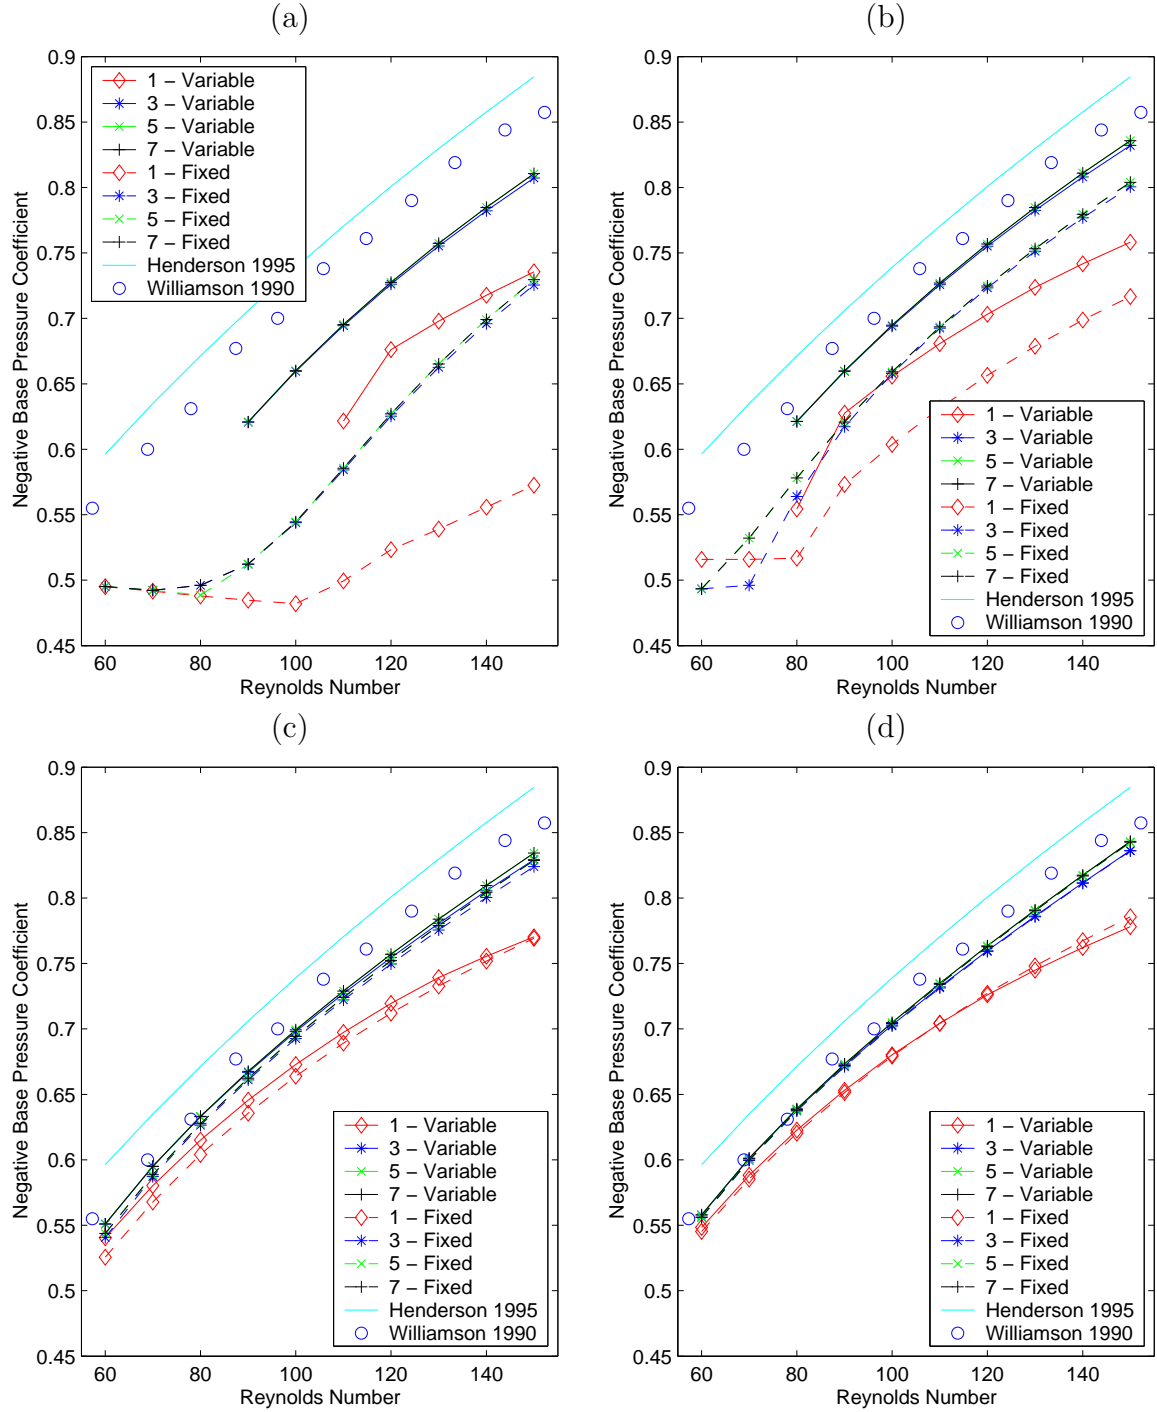


Figure 3.9: Base suction coefficient as a function of the Reynolds number plotted for different temporal resolutions. (a) 129x65 grid (b) 193x81 grid (c) 257x129 grid (d) 385x161 grid.

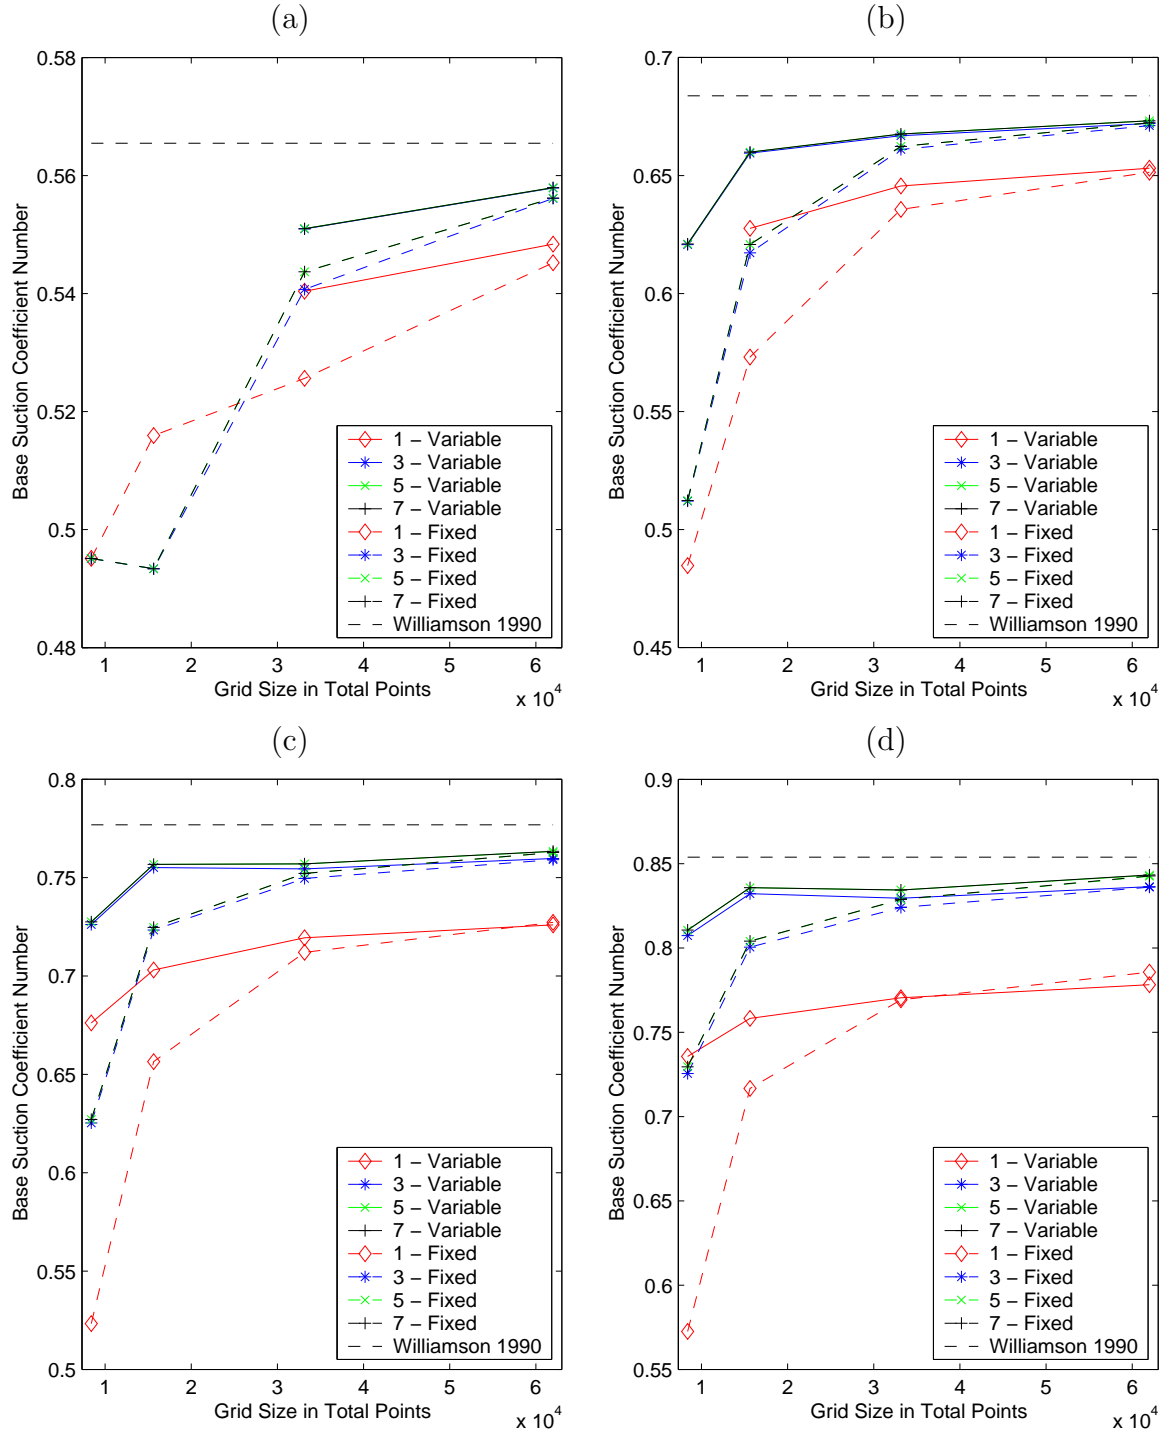


Figure 3.10: Base suction coefficient as a function of the spatial resolution plotted for different temporal resolutions. (a) $Re = 60$ (b) $Re = 90$ (c) $Re = 120$ (d) $Re = 150$.

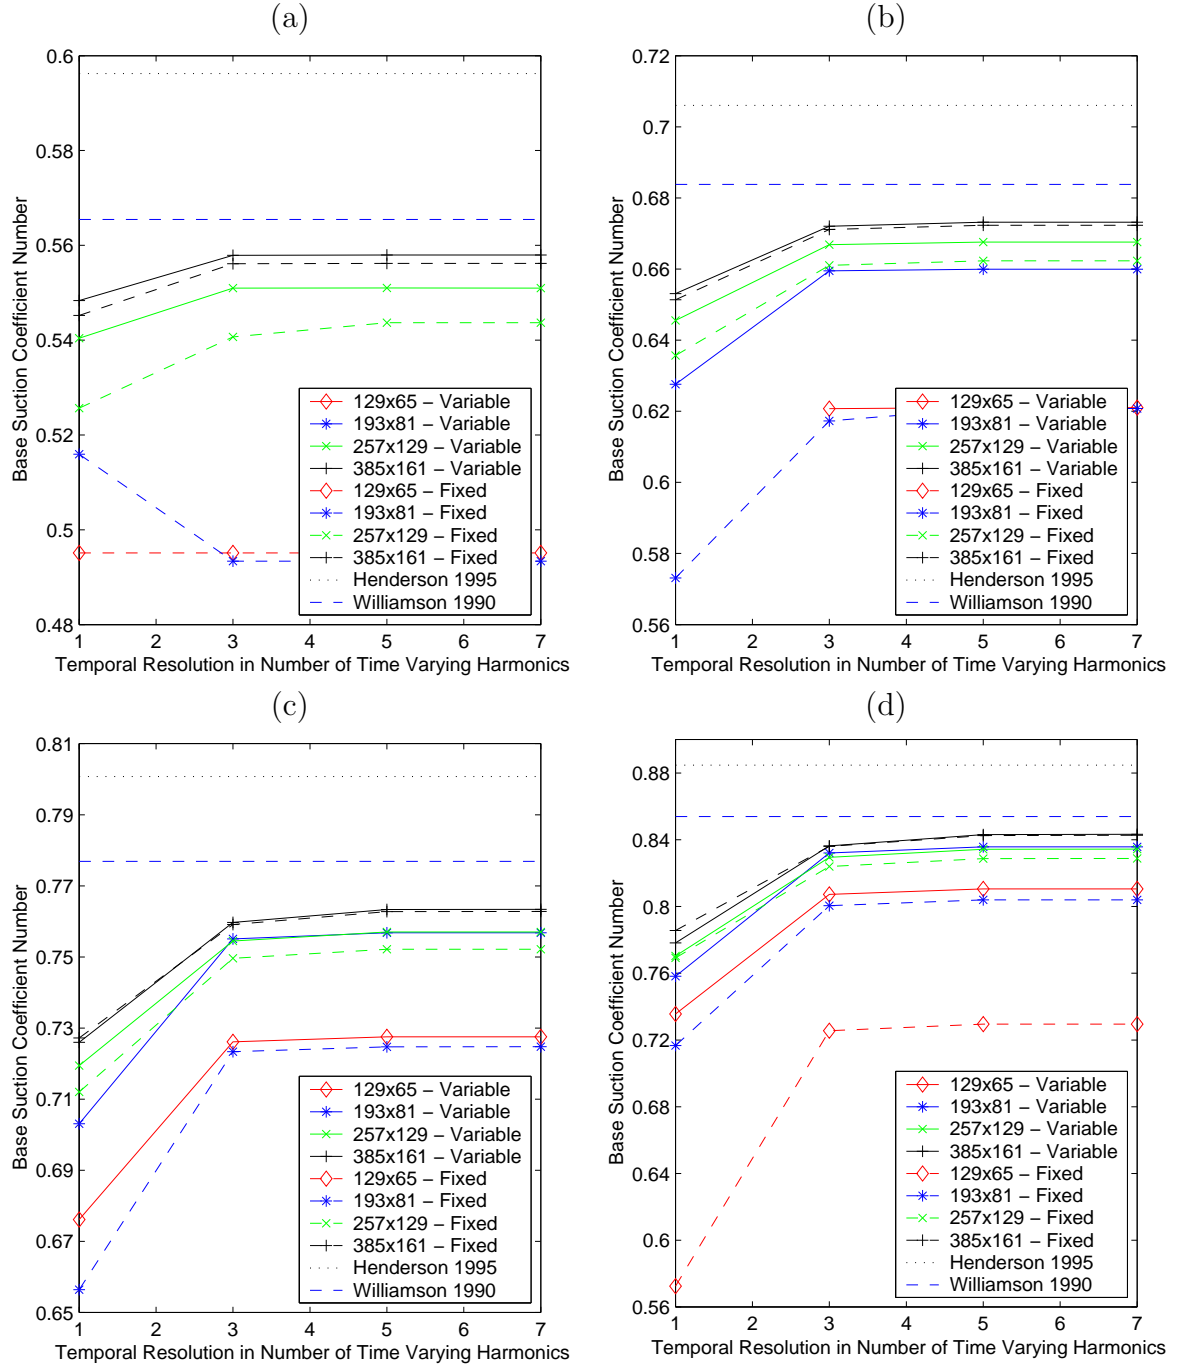


Figure 3.11: Base suction coefficient as a function of the temporal resolution plotted for different grid sizes. (a) $R_e = 60$ (b) $R_e = 90$ (c) $R_e = 120$ (d) $R_e = 150$.

Chapter 4

Pitching Airfoil Experimental Validation

This chapter addresses the second of three test cases which have been used to evaluate and validate the NLFD method. The numerical results produced by the NLFD code are compared with results obtained from independently conducted experiments of a pitching airfoil rotating periodically about its quarter chord at a given frequency. Farfield quantities like angle of attack, velocity and thermodynamic properties of the fluid are held constant in time. Due to the transonic environment, the solutions will contain shock waves that move periodically with some phase lag to the angle of attack changes [6]. The convergence of these results is assessed by surveys of spatial and temporal resolution based on global force coefficients.

4.1 Introduction

Table 4.1 is taken from McCroskey's review of airfoils in unsteady flows [46]. It lists the different governing equations used in these simulations and the physical conditions that limit their range of applicability. The approaches with the highest fidelity and computational cost are at the top of the list. This section presents the results of simulations using both the Euler and the Reynolds averaged Navier-Stokes equations using a Baldwin-Lomax turbulence model. Given the alternative choices, the test

cases should be representative of the capabilities of the governing equations. At issue is the level of nonlinearity in the problem which simplified models do not account for.

Equations	Remarks
Navier-Stokes with turbulence modeling	Strong shock waves and separation
Thin-layer Navier-Stokes with turbulence modeling	Strong shock waves, moderate separation
Euler	Inviscid, rotational, strong shock waves
Full potential	Weak shock waves, exact airfoil boundary conditions
Small disturbance Non-linear Time linearized	Simplified grids and boundary conditions Moderate amplitude Small amplitude shock motion, harmonic or indicial
Local linearizations	Small amplitudes, harmonic or indicial
Transonic linear theory	High-frequency, stationary shock waves

Table 4.1: Hierarchy of governing equations used in the analysis of unsteady transonic flow over airfoils. Source: McCroskey [46].

Citing work by Tijdemann and Seabass [64], Ashley [6] and Nixon [56], McCroskey [46] asserts “In the absence of boundary-layer separation, the motion of the shock wave is the essential feature that makes the unsteady transonic problem non-linear”. Dowell *et al.* [15, 16] suggest that solutions transition from the linear to the non-linear regime as the shock motion increases above 5 percent of the chord. The shock motion is a function of the magnitude of the dynamic angle of attack and the reduced frequency. In general, this motion and its resulting nonlinearities grow with increasing rotation magnitudes and decreasing reduced frequencies. Approximately linear solutions may be regained when the rotation angles become very small or the reduced frequencies become large.

Using the criteria presented above, the nonlinearity of the test cases can be established by analyzing the numerical simulations presented in subsection 4.7. Table 4.2 shows the magnitude of the shock movement based on inviscid NLFD solutions calculated with the highest temporal and spatial resolutions. The data supports the assertion that the cases used in this section are adequate tests for the non-linear

capabilities of the NLFD solver.

Investigator	Case Identifier	Shock Movement (percentage of chord)
Davis	CT Case 6	7.4%
Landon	CT Case 1	7.9%

Table 4.2: Shock movement as a percentage of the airfoil chord for the different experimental test cases used in chapter 4.

The effects of viscosity on airfoils in unsteady flows is primarily a function of the Reynolds number. This dimensionless quantity affects the shape of the boundary layer, the location of its transition/separation points, and the interaction between shock waves and the boundary layer [64]. Strong viscous effects are typified by flows containing shock-induced boundary layer separation and trailing-edge separation [46]. Numerical simulations of these flows are extremely sensitive to the turbulence modeling approach [61]. This is not a focus of this research. To isolate these effects, one of the test cases presented later in this section is a fully attached flow, while a second case exhibits mild separation affects.

4.2 Test Case Description

The experimental data was published as part of AGARD report 702. Two different airfoils, the NACA 64A010 and the NACA 0012, were tested by Davis [14] and Landon [42] respectively. Each study provides data for a variety of different test cases. This research uses only the priority cases containing the highest fidelity data published by each experimentalist. The important parameters used in the description of these priority cases are summarized in table 4.3.

The data for the 64A010 airfoil is for a transonic symmetric airfoil oscillating at a low reduced frequency over a limited range in angle of attack. Because of this fairly small variation in angle of attack, the numerical results are considered to be less sensitive to the choice of turbulence model than the data for the 0012 airfoil. Another

Description	Variable	Davis Experiment	Landon Experiment
AGARD Case Number		Data Set 2 CT Case 6 Dynamic Index 55	Data Set 3 CT Case 1
Airfoil		NACA 64A010	NACA 0012
Mean Angle of Attack	α_m	0.0°	2.89°
Angle of Attack Variation	α_0	$\pm 1.01^\circ$	$\pm 2.41^\circ$
Reynolds Number	Re_∞	12.56×10^6	4.8×10^6
Mach Number	M_∞	0.796	0.6
Reduced Frequency	k_c	0.202	0.0808

Table 4.3: Description of priority test cases used from AGARD Report 702.

factor contributing to its selection was its popularity in the numerical-analysis community. Numerous existing numerical results are available for comparison with those produced by the NLFD method.

In comparison to the data for the 64A010, the data for the 0012 was obtained at a lower Reynolds number but with an angle of attack variation approaching the stall boundary. The increased sensitivity of these results to the choice of turbulence model is a drawback. In the following sections, we first verify the accuracy of the spatial discretization of the NLFD scheme by running it as a steady-state solver and comparing its results to experimental and numerical datasets. These comparisons confirm that the grids used for the unsteady cases provide sufficient resolution. In addition, the comparisons exhibit the sensitivity of the results to the choice of the mathematical model.

Experimental results exist for the 64A010 airfoil in steady flow [40]. However, one of the advantages of using the 0012 airfoil is the large number of experimental studies that have been performed. McCroskey's [47] review identifies the major sources of error that exist in 50 independent 0012 wind-tunnel experiments and categorizes those experiments by their ability to mitigate the effect of these error sources. The dataset of Green and Newman [21] was selected for use in this paper, because it was the only transonic source in *group 1* of McCroskey's review. This *group 1* classification was

described as containing only “...the experiments which clearly stand out as having been conducted with the utmost care and/or as most nearly eliminating the important sources of wind-tunnel errors”. The test cases within this experiment, which are used as a basis for comparison with the numerical results are described in table 4.4.

Description	Variable	Green-Newman Experiment
Airfoil		NACA 0012
Mean Angle of Attack	α_m	$-2^\circ \leftrightarrow 10^\circ$
Reynolds Number	Re_∞	3.0×10^6 9.0×10^6
Mach Number	M_∞	0.7
Slotted-Wall Tunnel Runs		119 134
Adaptive-Wall Tunnel Runs		201

Table 4.4: Description of test cases used within Green and Newman’s steady experiment.

4.3 Computational Grid

Two separate grid configurations have been used in the simulations solving the Euler equations. The first grid configuration is an O-mesh generated by conformal mapping procedures. The second topology is a C-mesh which was generated using a hyperbolic grid generation tool. Table 4.5 provides a list of the different grids and their defining characteristics used in the Euler studies. The steady 0012 results presented in later sections show little variation in force coefficients over the range of spatial resolutions. Given this observation, the calculations for the 64A010 have been performed with grids that are generally coarser with larger cell heights at the airfoil wall. Figures 4.1 and 4.2 provide a nearfield picture of the grids for all the combinations of the different airfoils and grid sizes.

A C-mesh grid topology has been used exclusively for the Navier-Stokes calculations. The grids were generated with the same hyperbolic mesh generation tool used

Airfoil	Topology	Dimensions	Mean Boundary Distance (Chords)	Mean Grid Spacing at Wall (Chords)
0012	O-mesh	161x33	26	0.0090
0012	O-mesh	321x65	26	0.0045
0012	C-mesh	257x33	21	0.0035
0012	C-mesh	321x97	22	0.0017
64A010	O-mesh	81x33	26	0.0090
64A010	O-mesh	161x33	26	0.0090
64A010	C-mesh	129x33	21	0.0087
64A010	C-mesh	193x49	22	0.0087

Table 4.5: Description of meshes employed for the NLFD Euler calculations.

Airfoil	Topology	Dimensions	Boundary Distance (Chords)	Distance from First Cell to Wall (y^+)
0012	C-mesh	257x49	16	1.12
0012	C-mesh	513x97	18	1.10
64A010	C-mesh	129x33	15	11.6
64A010	C-mesh	193x49	12	6.9
64A010	C-mesh	257x65	12	3.8

Table 4.6: Description of meshes employed for the NLFD Navier-Stokes calculations.

for the Euler C-mesh grids. In contrast to the Euler cases, a significant grid refinement normal to the wall was employed to capture the viscous boundary layer. The average height of the first cell adjacent to the wall is quantified in table 4.6 in terms of y^+ units based on a Reynolds number of 3.0 million for 0012 airfoils and 12.5 million for the 64A010 airfoils. Other descriptive parameters including grid dimensions and farfield boundary distances are also provided in table 4.6. Figures 4.3 and 4.4 provide a nearfield picture of the grids for all permutations of spatial resolutions and airfoils.




Figure 4.1: Nearfield resolution of 0012 grids used in the Euler calculations. (a) 161x33 points in O-mesh topology (b) 321x65 points in O-mesh topology (c) 257x33 points in C-mesh topology (d) 321x97 points in C-mesh topology.




Figure 4.2: Nearfield resolution of the 64A010 grids used in the unsteady Euler calculations. (a) 81x33 points in O-mesh topology (b) 161x33 points in O-mesh topology (c) 129x33 points in C-mesh topology (d) 193x49 points in C-mesh topology.




Figure 4.3: Nearfield resolution of the 0012 Navier-Stokes grids. (a) 257x49 points in C-mesh topology (b) 513x97 points in C-mesh topology.

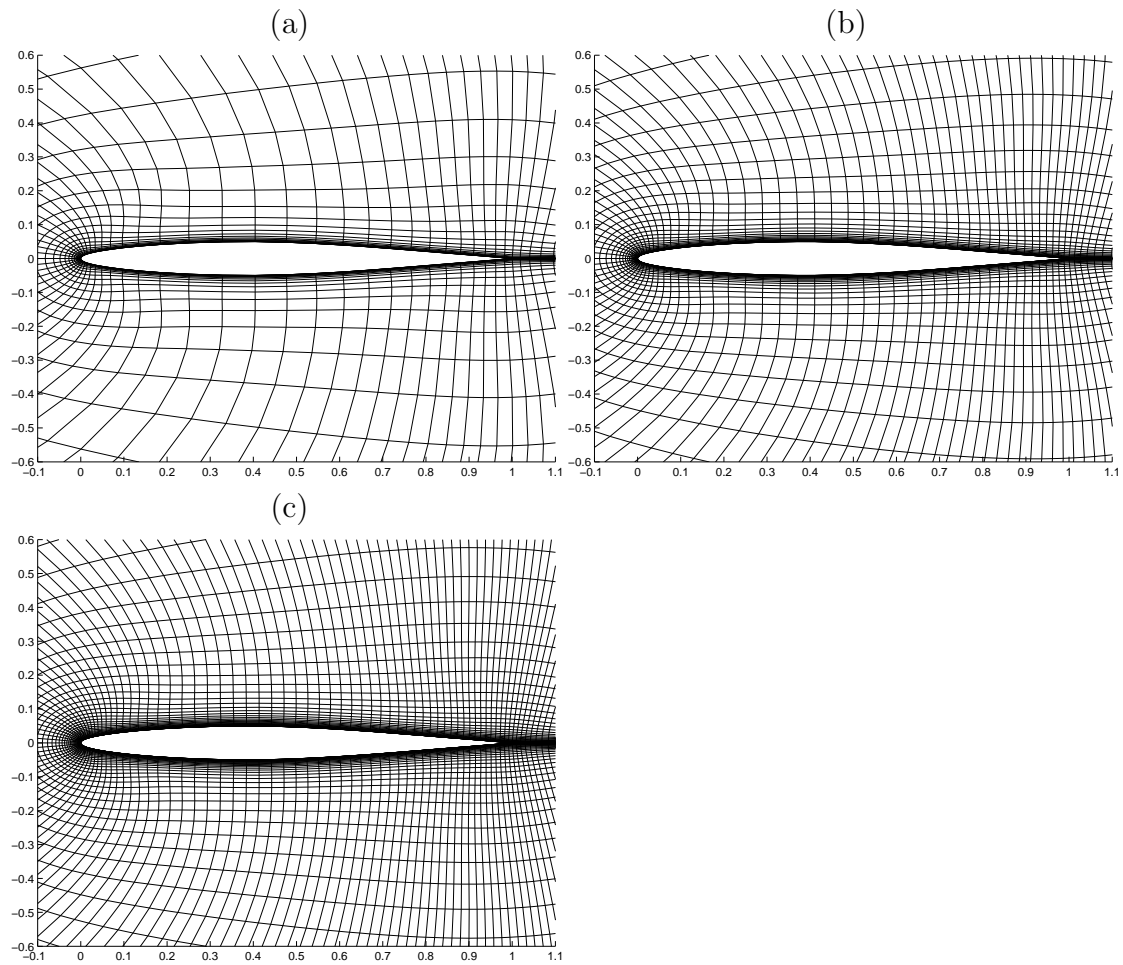


Figure 4.4: Nearfield resolution of the 64A010 grids used in the unsteady Navier-Stokes calculations. (a) 129x33 points in O-mesh topology (b) 193x49 points in O-mesh topology (c) 257x65 points in C-mesh topology.

4.4 Steady Inviscid Flow

This section compares the experimental results of Green and Newman [21] to steady Euler solutions produced by the NLFD code on the four 0012 grids introduced in the previous section. The following subsections provide numerical results for this airfoil, while excluding any steady results for the 64A010 airfoil. Matching the experimental test conditions, the farfield Mach number was set to 0.7 for the numerical simulations. Due to the range in angle of attack used by the unsteady experiments, the angle of attack for the steady numerical calculations will be limited to between -2° and 6° (calculated in 1° increments). Figures 4.5 through 4.7 compare numerical results from the NLFD code to the two different sets of experimental data collected at Reynolds numbers of 3.0 and 9.0 million. The steady results at 3.0 million are included due to their proximity to Landon's unsteady results at 4.8 million which are the focus of this research. The additional dataset at 9.0 million provides trends in the force coefficients for increasing Reynolds number. This trend would be expected to approach the solutions of the Euler equations that represent the limit as the Reynolds number increases to infinity.

There are six different sets of experimental data points in these figures. They represent all the permutations of three different wind-tunnel tests using two iterative passes in the correction process of the force coefficients. The experimental results were collected in the NASA Langley 0.3 meter Transonic Cryogenic Tunnel (TCT) using both slotted- and adaptive-wall technologies to minimize interference effects. Results numbered 200 and greater are those collected using adaptive walls while those results numbered less than 200 are the slotted-wall experiments. The raw force coefficients are corrected using a nonlinear post-test wall-interference assessment/correction (WIAC) code. The code solves the two-dimensional transonic small disturbance equations using three different sets of boundary conditions. The three different solutions are repeated until the camber line of an equivalent inviscid body matches the actual test geometry [21, 52, 20]. Different passes from this scheme are provided in figures 4.5 through 4.7 to show changes between iterations in the correction process.

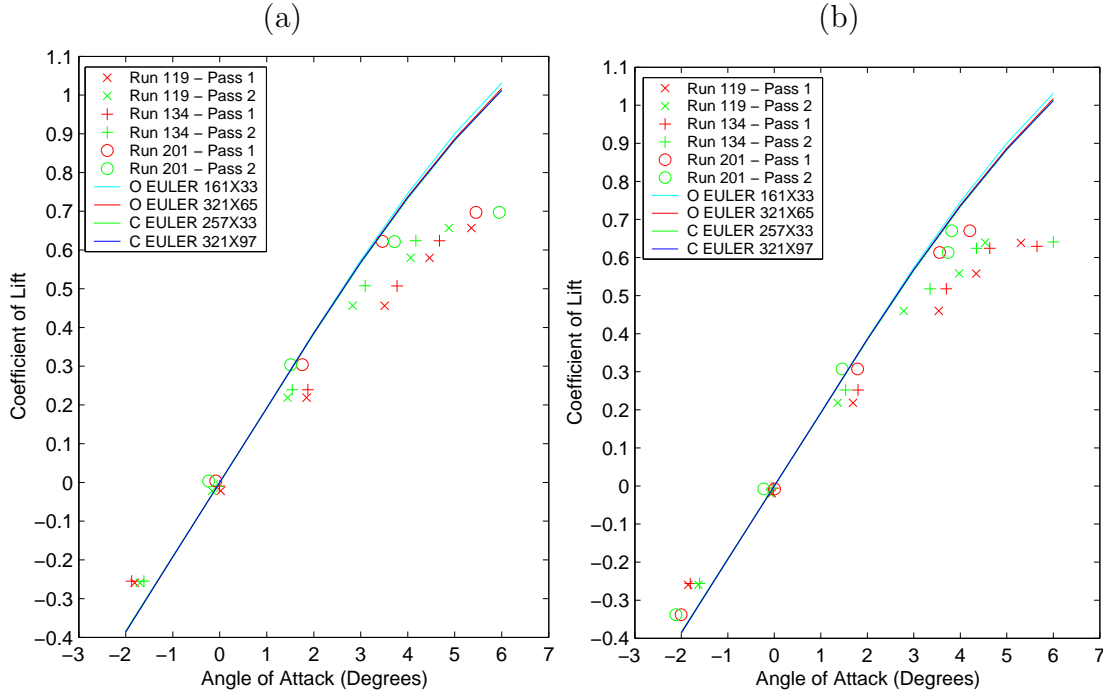


Figure 4.5: Lift coefficient as a function of angle of attack at a free-stream Mach number of 0.7 including both Euler NLFD and experimental data. (a) Experimental Reynolds number equal to 3.0 million. (b) Experimental Reynolds number equal to 9.0 million.

Coefficient of Lift Results

For the entire range in angles of attack, the numerical C_l results show little variation over the range of different spatial resolutions and grid topologies. Below 3° angle of attack, the Euler results consistently over-predict the magnitude of the experimental lift curve slope. Table 4.7 provides an estimate of this slope using a least-squares fit based on *pass 2* data points with an angle of attack less than 3° .

Above 3° angle of attack, the lift curve slope of the experimental data begins to decay. This is consistent with the hypothesis that the flow contains a shock induced separation. This observation is supported by subsequent Navier-Stokes calculations where at 3° angle of attack, a small region of separation immediately following the shock is predicted. As expected, the numerical Euler results do not predict the shock

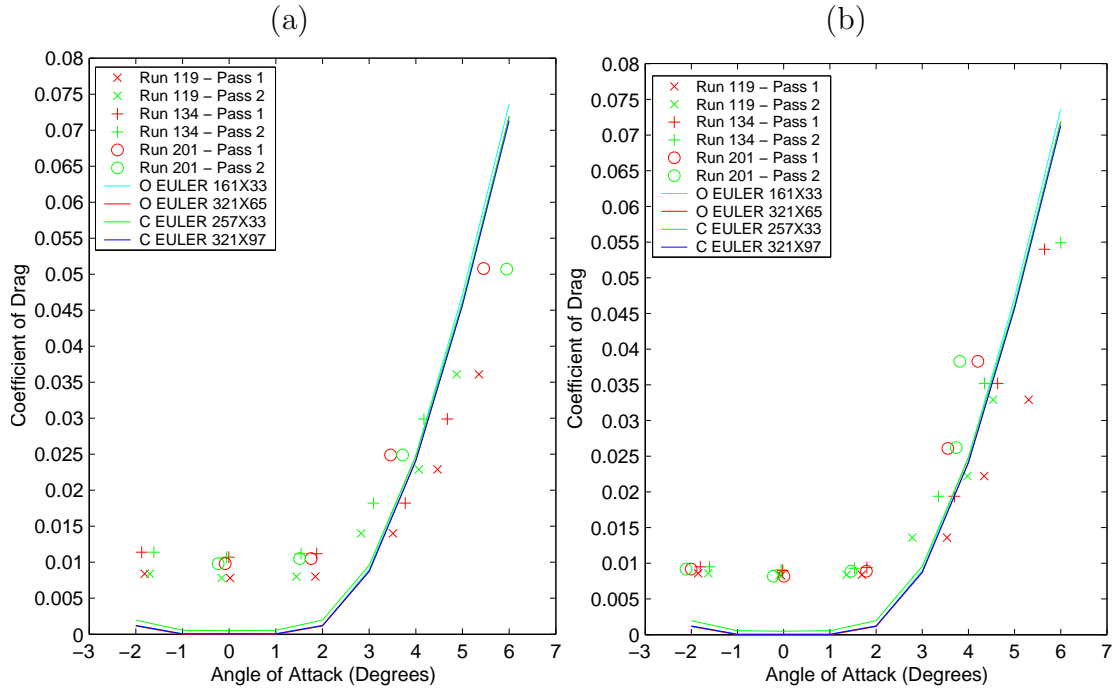


Figure 4.6: Drag coefficient as a function of angle of attack at a free-stream Mach number of 0.7 including both Euler NLFD and experimental data. (a) Experimental Reynolds number equal to 3.0 million. (b) Experimental Reynolds number equal to 9.0 million.

induced flow separation, consequently providing poor results for angles of attack above 3° .

Two trends can be noted based on inspection of the data contained in table 4.7 and figure 4.5. As expected, increasing the Reynolds number shifts the experimental datasets closer to the numerical inviscid results. Second, the Euler numerical results are consistently closer to the adaptive-wall data than the slotted-wall data. However, Navier-Stokes calculations presented in future sections exhibit a reduced lift-curve slope in good agreement with the slotted-wall data (especially run 134). Based on the steady inviscid results, one can conclude that the resolution of the grids is adequate to predict unsteady lift coefficients using the Euler equations under the assumption that the flow remains attached.

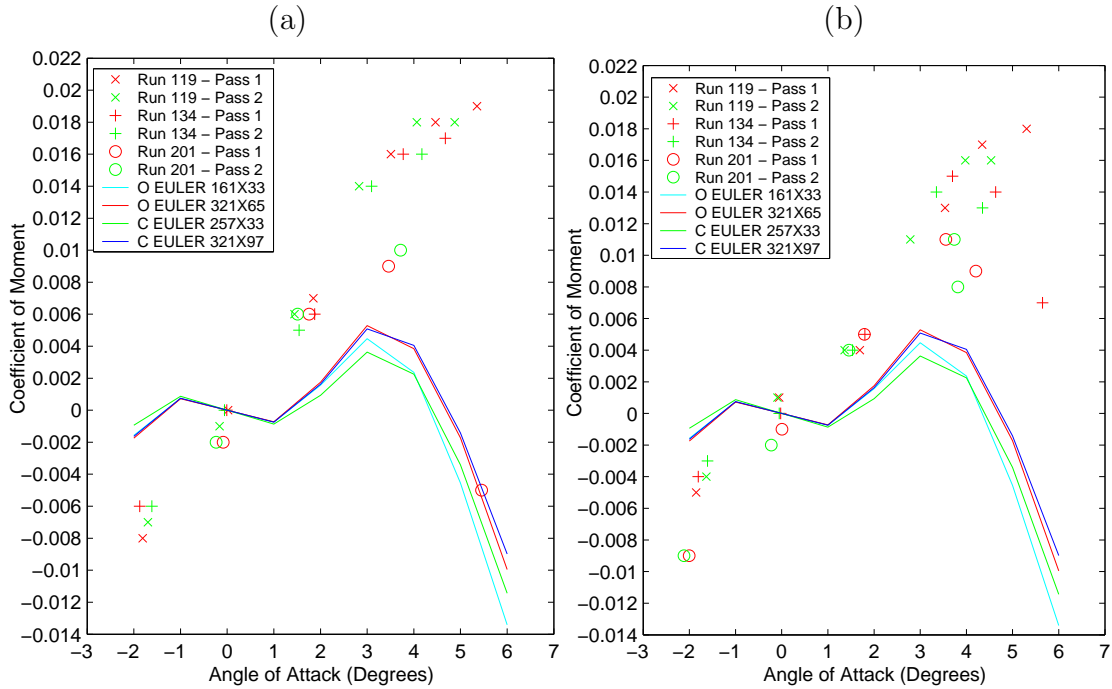


Figure 4.7: Moment coefficient as a function of angle of attack at a free-stream Mach number of 0.7 including both Euler NLFD and experimental data. (a) Experimental Reynolds number equal to 3.0 million. (b) Experimental Reynolds number equal to 9.0 million.

Coefficient of Drag Results

A slight variation in C_d results appears over the different spatial resolutions employed. This variation is attributable to both the different mesh topologies and the difference in artificial dissipation schemes employed in the O-mesh and C-mesh solvers. As expected, the numerical C_d results based on the Euler equations differ from the viscous experimental results in the absence of an estimate of the drag due to friction. However, both results consistently predict the onset of drag divergence as the angle of attack is increased. Since the unsteady experimental results do not differentiate between the viscous and shock wave components of drag it will be difficult to make useful comparisons of numerical and experimental results for any test cases except those dominated by strong shocks.

Results Type	Reynolds Number (mil)	Lift Curve Slope (per degree)
Run 119	3	0.157
Run 119	9	0.163
Run 134	3	0.157
Run 134	9	0.162
Run 201	3	0.172
Run 201	9	0.181
O-mesh 161x33	∞	0.193
O-mesh 321x65	∞	0.191
C-mesh 257x33	∞	0.190
C-mesh 321x97	∞	0.191

Table 4.7: Steady lift curve slope based on angles of attack less than 3° for all numerical and experimental results.

Coefficient of Moment Results

At higher angles of attack, there is a significant variation between the numerical C_m predictions over the range of spatial resolutions. The finer discretizations tend toward the experimental results, but the variation between the datasets remains large. Further grid refinements produce no substantive effects. Even in areas where the flow remained fully attached, the numerical C_m predictions approximate the experimental data, yet show the opposite slope with respect to change in angle of attack. However, when the magnitude of the C_m results are scaled by the maximum C_l (the maximum C_m is roughly forty times less than the maximum C_l), the numerical and experimental results approximate the prediction of thin airfoil theory that the coefficient of moment for a symmetric airfoil is zero. Despite the scaling arguments, the results of the steady calculations force the expectation that Euler calculations will be inaccurate in the prediction of the unsteady moment coefficients.

4.5 Steady Viscous Flow

This section compares the experimental results of Green and Newman [21] to steady Navier-Stokes solutions produced by the NLFD code on the two grids introduced in the section 4.3. The farfield Mach number and the range in angle of attack are identical to those used in the Euler surveys presented in section 4.4. Figure 4.8 compares numerical results from the NLFD code to the experimental data collected at a Reynolds number of 3.0 million.

Coefficient of Lift Results

A description of the six different sets of experimental data is provided in section 4.4. As with the Euler results, the numerical C_l results show little variation between the coarse and fine grids for angles of attack less than 3° . The numerical solution at 3° angle of attack contains a small zone of separation immediately following the shock. This area extends roughly 4 percent downstream of the shock where the flow reattaches. As angle of attack increases, the length of the separated region continues to increase. Calculations on the coarse grid at 6° angle of attack predict separated flow from the shock position to the trailing edge. Calculations on the fine grid at the same angle of attack would not accurately converge to a steady state. At issue is the ability of the turbulence model to accurately capture the physics of the flow in the separated region ultimately predicting the location of reattachment. Given the simplicity of the Baldwin-Lomax turbulence model, any results showing regions of separated flow are suspect.

Table 4.8 provides the lift curve slope for the experimental results and the numerical Navier-Stokes results based on angles of attack less than 3° . Increasing the spatial resolution has a slight tendency to increase the lift curve slope moving the numerical predictions toward the adaptive-wall data.

Coefficient of Drag Results

Even with significant regions of separated flow, the numerical C_d results agree well with the experimental data. The adaptive-wall data shows remarkable agreement for

Results Type	Lift Curve Slope (per degree)
Run 119	0.157
Run 134	0.157
Run 201	0.172
C-mesh 257x49	0.165
C-mesh 513x97	0.167

Table 4.8: Steady lift curve slope based on angles of attack less than 3° for Navier-Stokes calculations (Reynolds number of 3.0 million) and experimental results.

the lower angles of attack. Pass two of the slotted-wall data exhibits slightly poorer agreement for these lower angles of attack, but improves at angles above 2° . As with the numerical C_l results, little is gained in the prediction of C_d by increasing the spatial resolution.

Coefficient of Moment Results

In contrast to the numerical Euler results, the viscous C_m results show excellent agreement with the experimental data when the flow remains attached. Agreement with the adaptive-wall data set is poor above 3° angle of attack, but both sets of slotted-wall data show decent agreement through the entire range. The data shows that for symmetric airfoils which induce relatively negligible moment coefficients, the addition of the viscous terms in the numerical calculations is required to accurately capture the experimental C_m data. Except in the separated flow region, little difference is exhibited between the two numerical results.

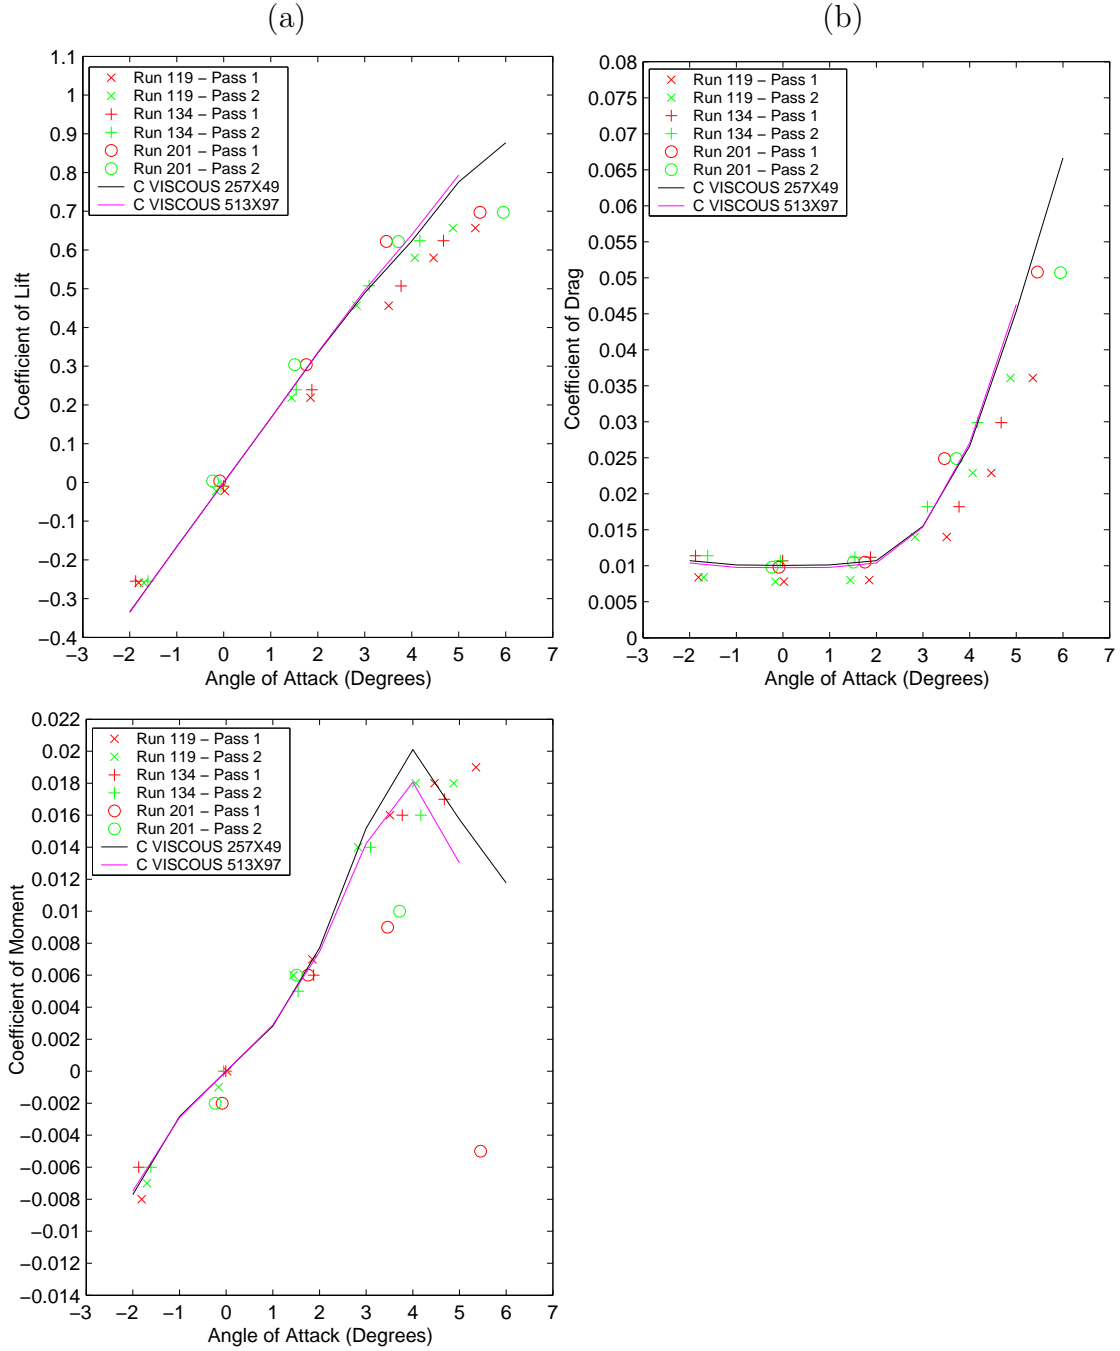


Figure 4.8: Force coefficient as a function of angle of attack at a free-stream Mach number of 0.7 and a Reynolds number of 3.0 million including both Navier-Stokes NLFD and experimental data. (a) Coefficient of lift. (b) Coefficient of drag. (c) Coefficient of moment.

4.6 Unsteady Inviscid Flow

This section compares the experimental results of Davis and Landon with numerical calculations of the unsteady Euler equations by the NLFD code. The parameters of each numerical simulation match the description of the experiments provided in section 4.2. For each Euler mesh presented in section 4.3 three separate solutions are calculated using one, two and three time varying modes. Both the 0012 and 64A010 airfoils are included in this section. Overall, 12 different cases have been calculated.

In the case of the 0012 airfoil the unsteady data was obtained at a lower Mach number than the steady data with a dynamic angle of attack above the stall angle deduced from the steady results presented in section 4.4. Inspection of the viscous solutions presented later in this chapter show a small region of separated flow appearing briefly behind the shock. The flow reattaches only to re-separate at the trailing edge. This latter separation zone is larger (roughly 1 to 2 percent of chord) and exists on the suction side of the airfoil for half the period of oscillation. Given these viscous effects the accuracy of the 0012 inviscid solutions is suspect.

Coefficient of Lift Results

Figures 4.9 and 4.10 show both the numerical and experimental C_l results as a function of the instantaneous angle of attack for the 0012 and 64A010 airfoils. For each grid, a subfigure shows several ellipses each computed using a different number of time varying modes. These plots show that the variation in time varying C_l as a function of the temporal resolution is negligible, and that results convergent to plotting accuracy can be obtained using one mode.

Figures 4.11 and 4.12 shows the same force coefficient data, except that each subfigure now plots the results as a function of the spatial resolution. Three subfigures, one for each temporal resolution, contain multiple ellipses representing the solutions calculated on different grids. There are some variations in the results as a function of spatial resolution, but the effects are negligible and provide no visible trend toward improving the comparison between experimental and numerical results.

Transfer functions are introduced to quantify the time varying results presented

in previous paragraphs. The magnitude of this transfer function is computed by dividing the magnitude of the Fourier coefficient for C_l at a given wavenumber by the magnitude of the Fourier coefficient for α at the same wavenumber

$$L_\alpha = \frac{\hat{C}_{l_k}}{\hat{\alpha}_k} = \frac{\|\hat{C}_{l_k}\|}{\|\hat{\alpha}_k\|} e^{i(\angle \hat{C}_{l_k} - \angle \hat{\alpha}_k)} \quad (4.1)$$

This is analogous to the lift curve slope commonly used in the analysis of steady state airfoils. The phase lag of the transfer function is defined as the difference in phase angles between the Fourier coefficients for C_l and α quantified in terms of degrees. The phase angle represents the fraction of the time period that the coefficient of lift lags behind the angle of attack.

Nonlinear systems transfer energy between modes providing non-zero contributions to C_l and C_m in temporal modes that are not being forced by the boundary conditions. Transfer functions will not be quantified for these unforced modes, since they are generally smaller than the fundamental harmonic and $\hat{\alpha}_k$ is zero for these cases. Tables 4.9 to 4.11 provide the transfer functions in terms of magnitude and phase lag for the experimental data and all spatial and temporal resolutions in the NLFD numerical datasets.

For the zeroth harmonic (time average), the numerical estimate of the magnitude of the 0012 lift transfer function consistently over-predicts the experimental data. On average the difference between the numerical and experimental predictions is 11 percent, with negligible difference existing between the various spatial and temporal resolutions. The steady numerical results in sections 4.4 and 4.5, and the unsteady viscous results that will be introduced in section 4.7 show that this over-prediction can be attributed to the lack of viscosity in these simulations.

However for the fundamental harmonic, there is no consistent trend in the magnitude of the C_l transfer function. The coarse grid 0012 results provide excellent agreement with the experimental data. Unfortunately, increasing the spatial resolution lowers these estimates, and increases the distance between the experimental and numerical results. A negligible difference is exhibited as additional temporal harmonics are added to the solution. The agreement between the experimental and

Mesh Type	1 Harmonic	2 Harmonic	3 Harmonic
Experimental	0.143		
O-mesh 161x33	0.161	0.161	0.161
O-mesh 321x65	0.160	0.161	0.161
C-mesh 257x33	0.160	0.160	0.160
C-mesh 321x97	0.160	0.160	0.160

Table 4.9: Magnitude of zeroth harmonic in coefficient of lift divided by the magnitude of the zeroth harmonic in angle of attack $\left(\frac{\|\hat{C}_{l_0}\|}{\|\hat{\alpha}_0\|}\right)$ for various temporal resolutions employed by NLFD simulations on the 0012 airfoil.

average numerical results is excellent with a variation of approximately 1 percent. In regards to the 64A010 dataset, little variation is observed in the magnitude of the C_l transfer function throughout the range of spatial and temporal resolutions. Contrary to the 0012 case, the 64A010 results consistently over-predict the experimental data on average by 9 percent.

Airfoil	Mesh Type	1 Harmonic	2 Harmonic	3 Harmonic
0012	Experimental	0.121		
0012	O-mesh 161x33	0.122	0.121	0.121
0012	O-mesh 321x65	0.119	0.118	0.118
0012	C-mesh 257x33	0.123	0.122	0.122
0012	C-mesh 321x97	0.117	0.116	0.116
64A010	Experimental	0.095		
64A010	O-mesh 81x33	0.104	0.104	0.103
64A010	O-mesh 161x33	0.104	0.104	0.104
64A010	C-mesh 129x33	0.104	0.104	0.104
64A010	C-mesh 193x49	0.106	0.105	0.105

Table 4.10: Magnitude of fundamental harmonic in coefficient of lift divided by the magnitude of the fundamental harmonic in angle of attack $\left(\frac{\|\hat{C}_{l_1}\|}{\|\hat{\alpha}_1\|}\right)$ for various temporal resolutions employed by NLFD simulations.

The numerical results for the phase lag between C_l and α for the 0012 and 64A010 airfoils respectively over-predict and under-predict the experimental data. In all but one case, increasing the spatial or temporal resolutions moves the numerical results

away from the experimental data. However, the variation between datasets remains small; on average -2.1° and 4.2° for the 0012 and 64A010 airfoils respectively.

Airfoil	Mesh Type	1 Harmonic	2 Harmonic	3 Harmonic
0012	Experimental	-10.9		
0012	O-mesh 161x33	-13.1	-13.8	-13.7
0012	O-mesh 321x65	-13.9	-14.7	-14.7
0012	C-mesh 257x33	-11.2	-11.8	-11.8
0012	C-mesh 321x97	-11.8	-12.5	-12.5
64A010	Experimental	-25.6		
64A010	O-mesh 81x33	-22.3	-22.1	-22.2
64A010	O-mesh 161x33	-21.3	-21.2	-21.2
64A010	C-mesh 129x33	-20.7	-20.6	-21.0
64A010	C-mesh 193x49	-21.5	-21.3	-21.3

Table 4.11: Phase lag between the fundamental harmonic in lift coefficient and angle of attack $(\angle \hat{C}_{l_1} - \angle \hat{\alpha}_1)$ for various temporal resolutions employed by NLFD simulations on the 0012 airfoil.

Coefficient of Moment Results

This section presents a corresponding set of figures for the coefficient of moment results. Figures 4.13 and 4.14 show both the numerical and experimental C_m results as a function of the instantaneous angle of attack for the 0012 and 64A010 airfoils. For each grid, a subfigure shows several ellipses computed using different numbers of time varying modes. As could be expected from the steady simulations, the agreement between the experimental and numerical results is poor, and will be improved by the addition of viscosity. For the 0012 airfoil, some variation in C_m results appears to be a function of temporal resolution while the 64A010 results are invariant to the number of modes used in the solution. Although not shown, cases with higher temporal resolution were calculated for the 64A010 airfoil, showing that convergence to plotting accuracy was achieved with three temporal modes.

Figures 4.15 and 4.16 show the same force coefficient data, except that each subfigure now shows the results as a function of the spatial resolution. Three subfigures,

one for each temporal resolution, contain multiple ellipses representing the solutions calculated on different grids. There are some variations in the results as a function of spatial resolution, but the effects are negligible and provide no visible trend in improving the comparison between experimental and numerical results.

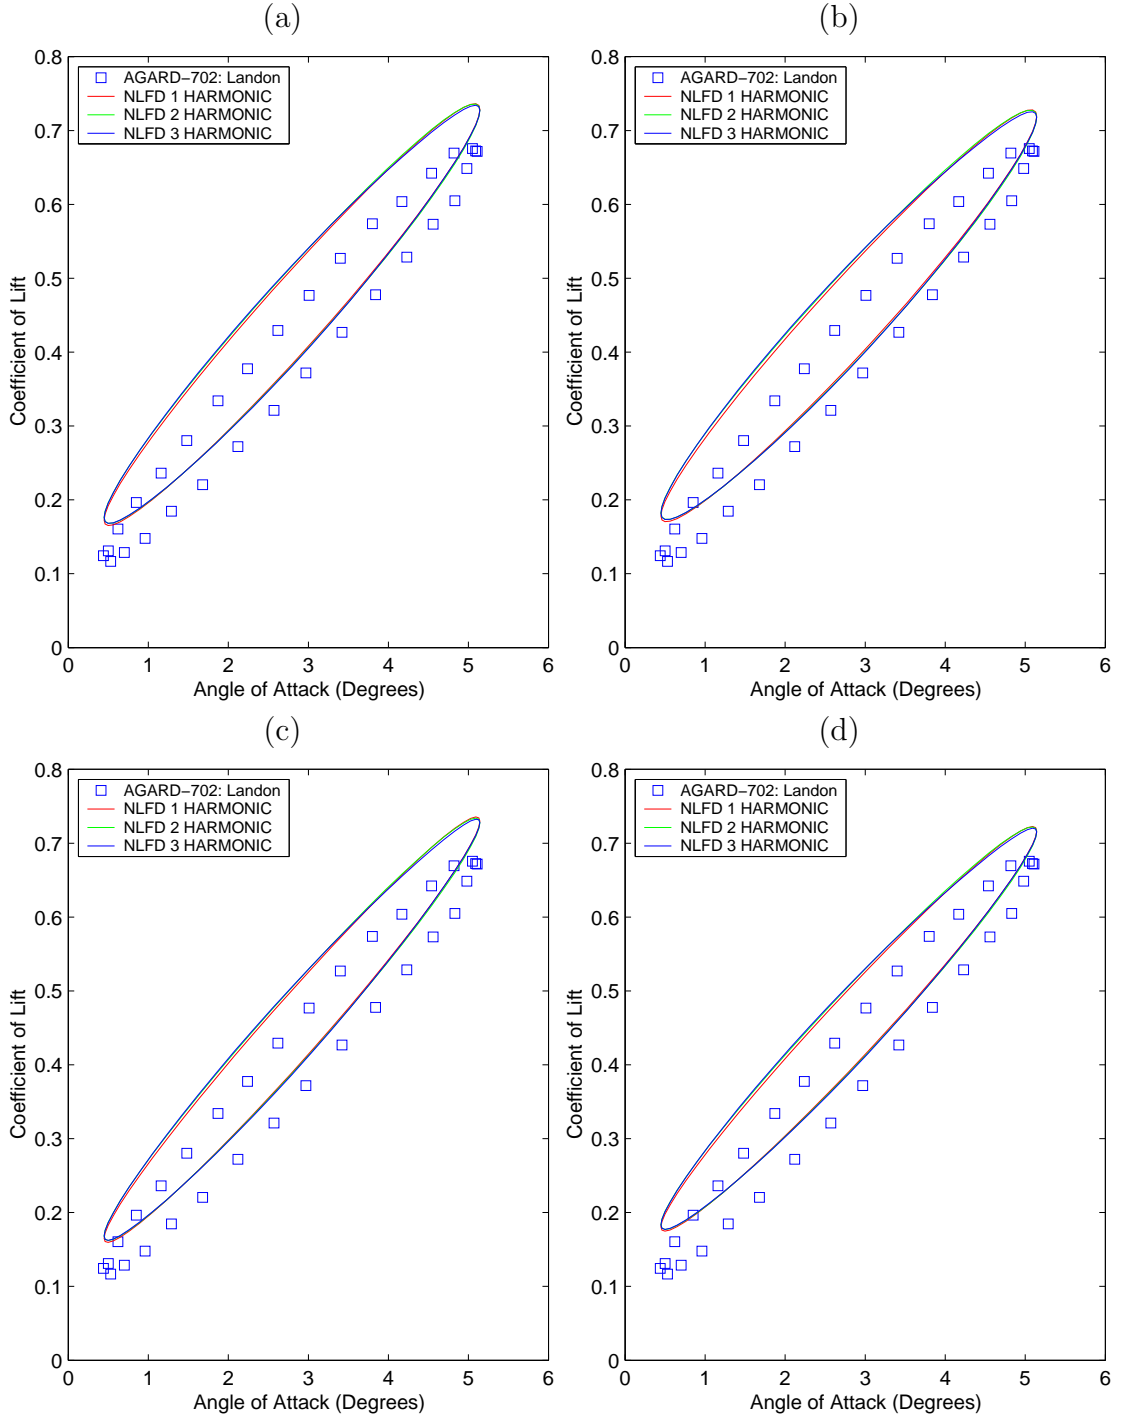


Figure 4.9: Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Landon's 0012 experiment. (a) 161x33 O-mesh. (b) 321x65 O-mesh. (c) 257x33 C-mesh. (d) 321x97 C-mesh.




Figure 4.10: Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Davis's 64A010 experiment. (a) 81x33 O-mesh. (b) 161x33 O-mesh. (c) 129x33 C-mesh. (d) 193x49 C-mesh.

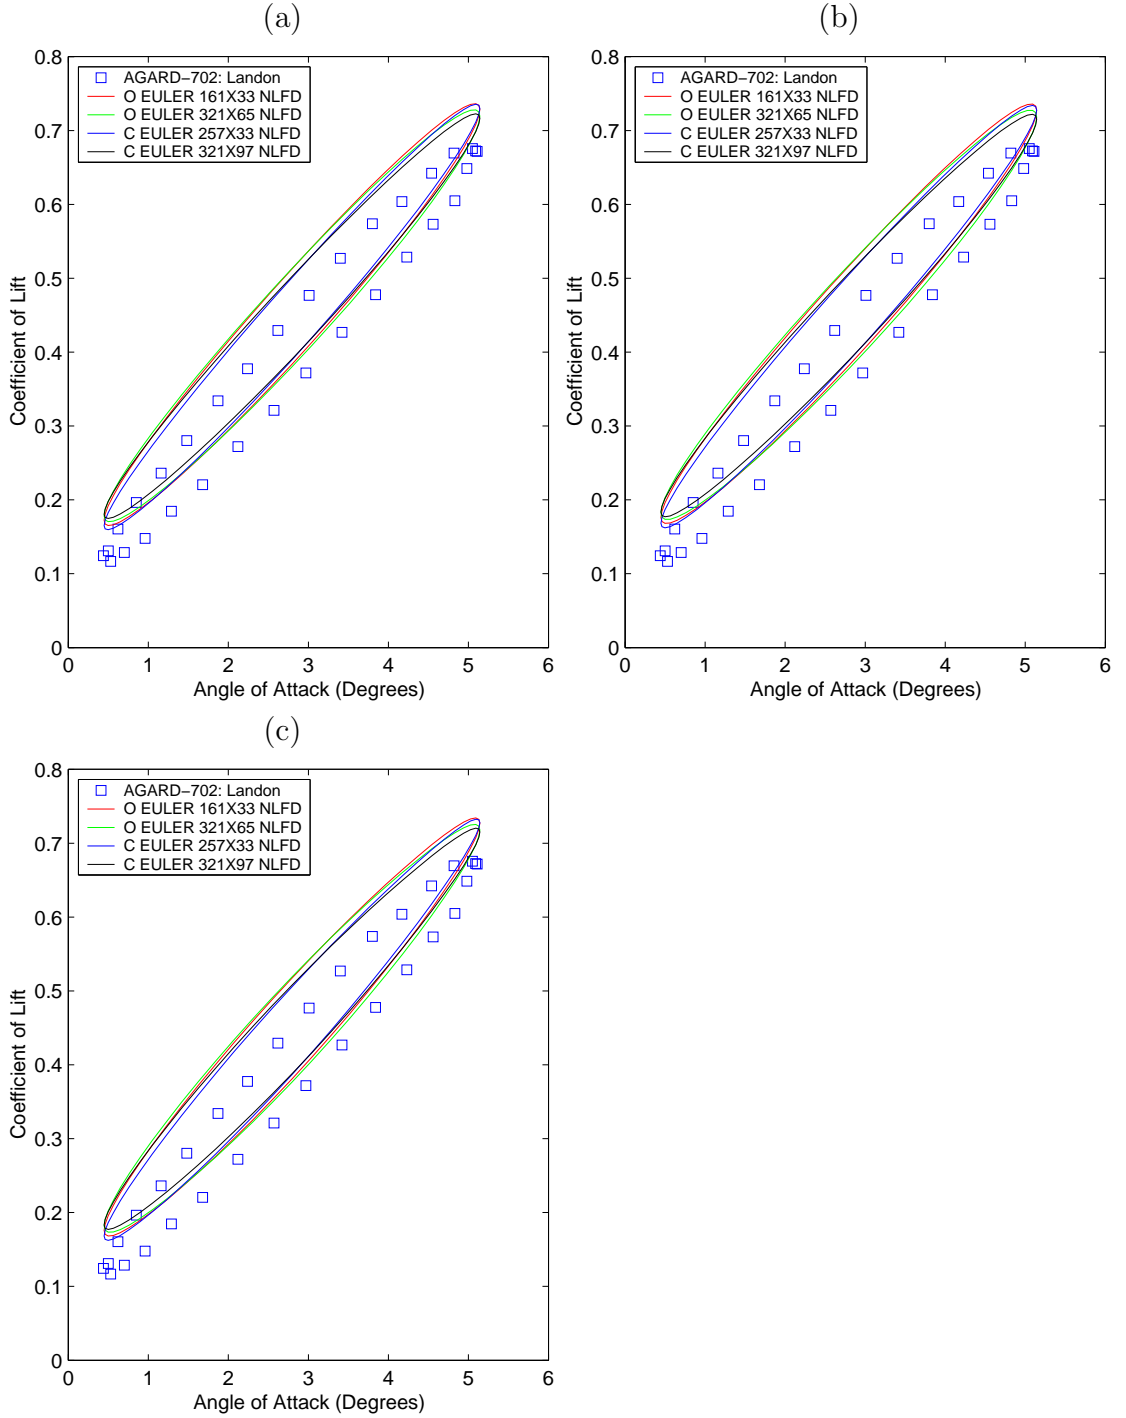


Figure 4.11: Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Landon's 0012 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics.

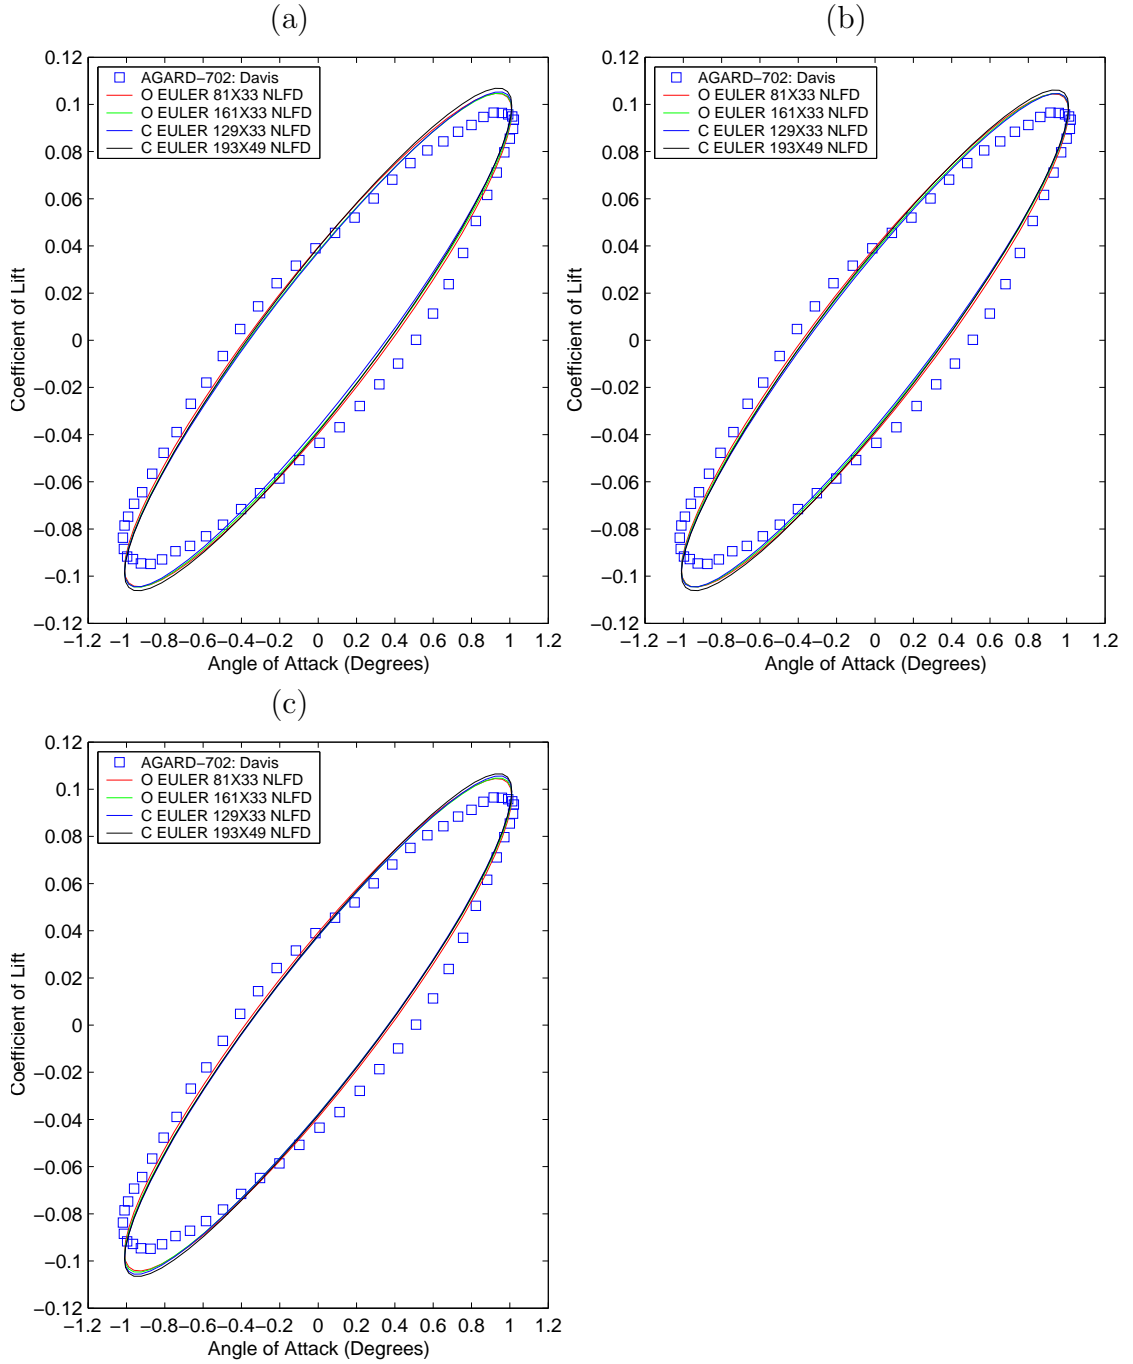


Figure 4.12: Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Davis's 64A010 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics.

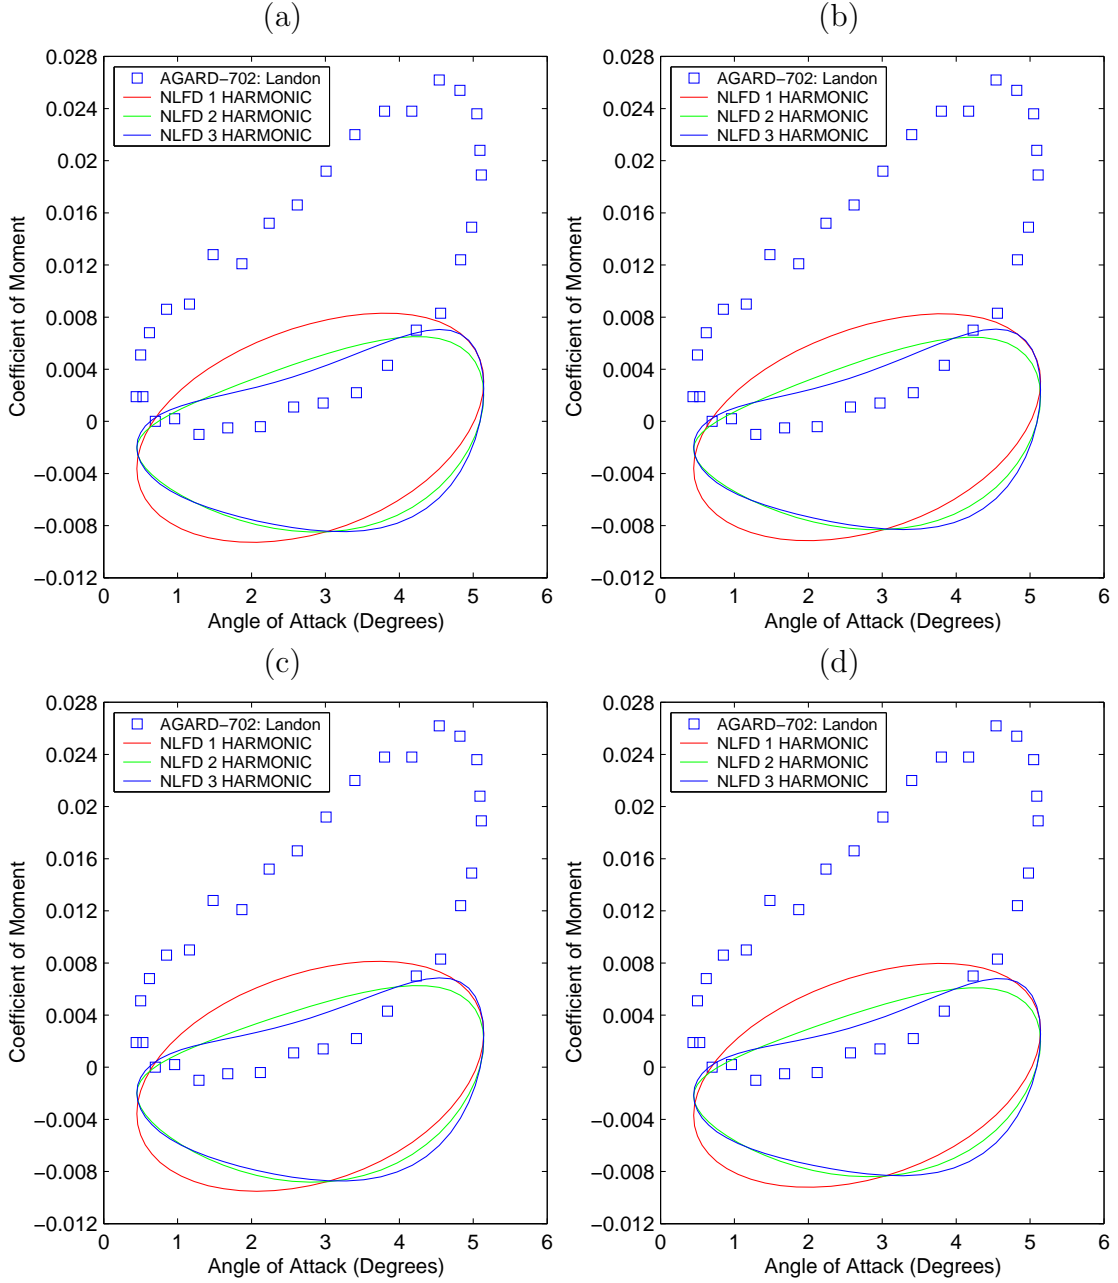


Figure 4.13: Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Landon's 0012 experiment. (a) 161x33 O-mesh. (b) 321x65 O-mesh. (c) 257x33 C-mesh. (d) 321x97 C-mesh.

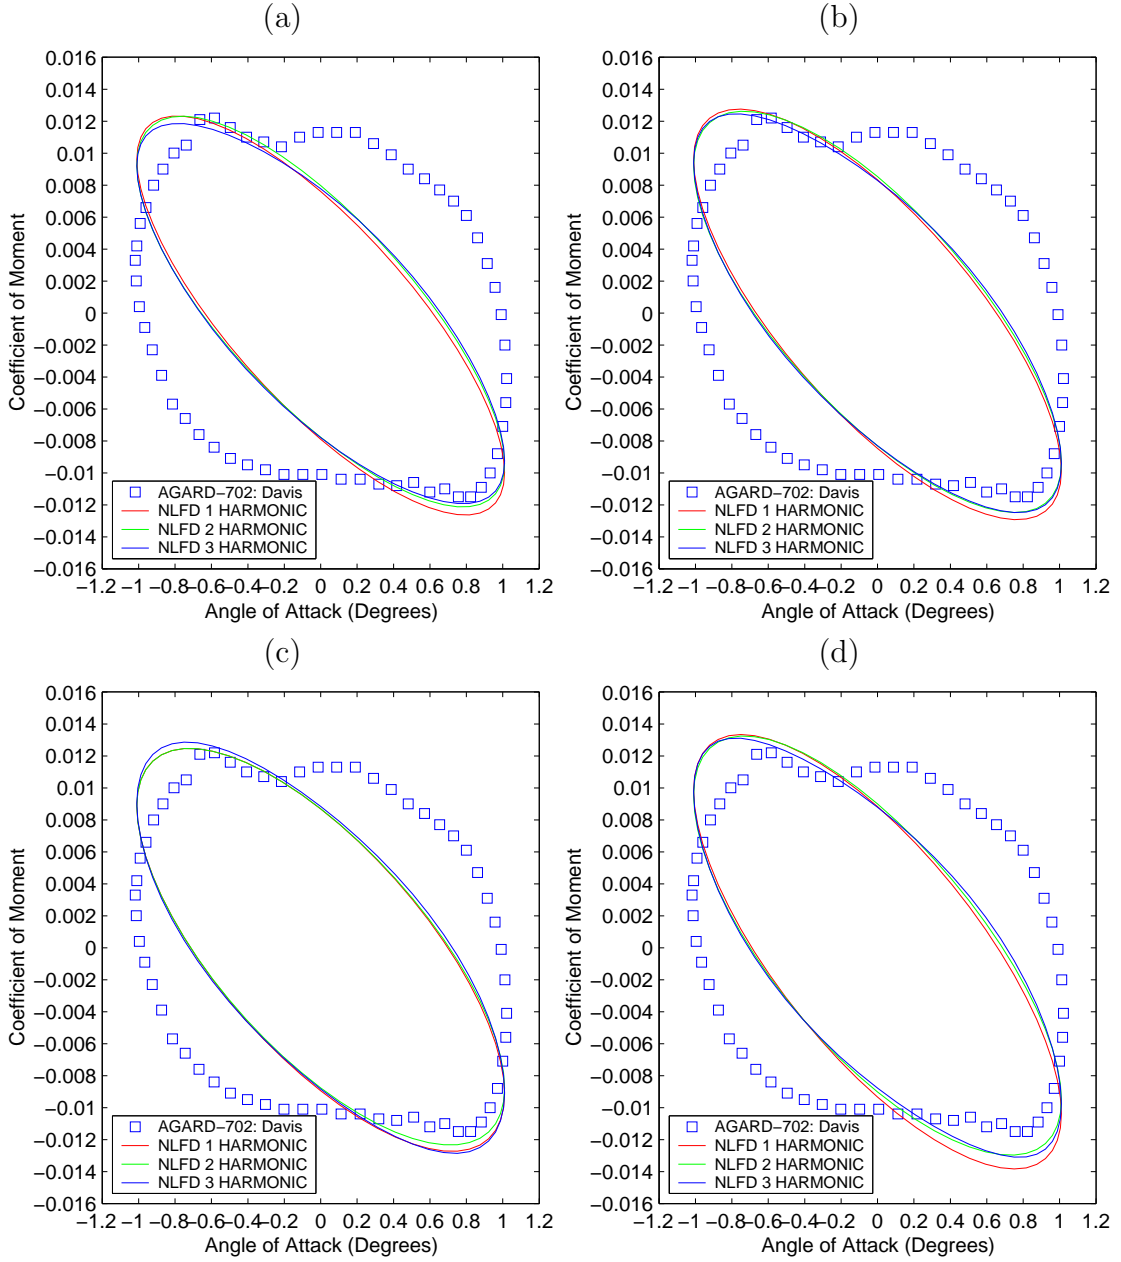


Figure 4.14: Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Davis's 64A010 experiment. (a) 81x33 O-mesh. (b) 161x33 O-mesh. (c) 129x33 C-mesh. (d) 193x49 C-mesh.

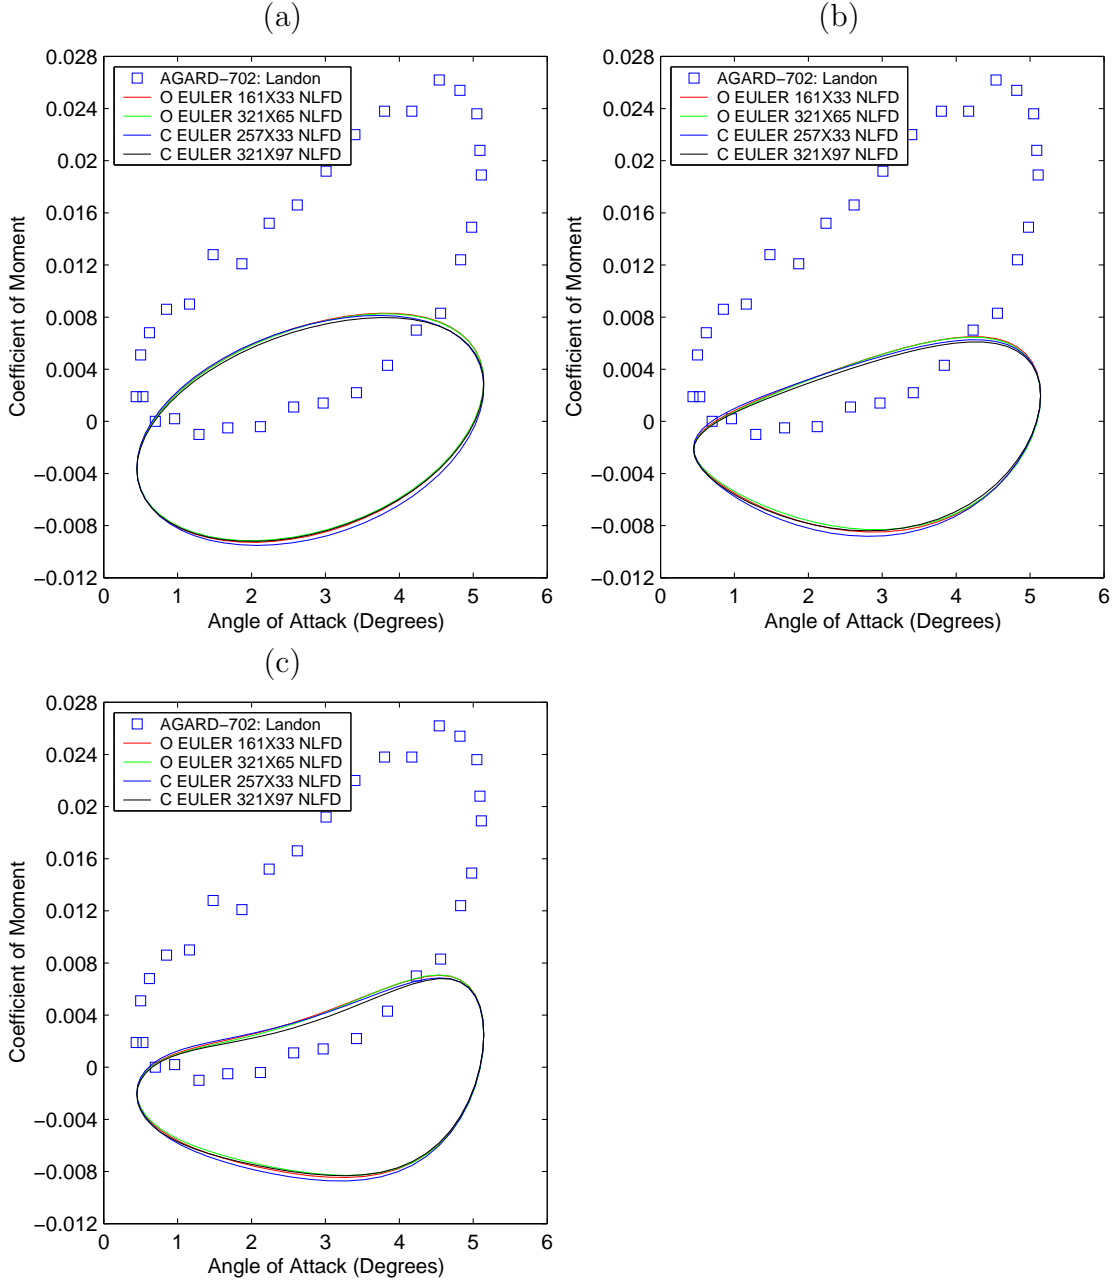


Figure 4.15: Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Landon's 0012 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics.

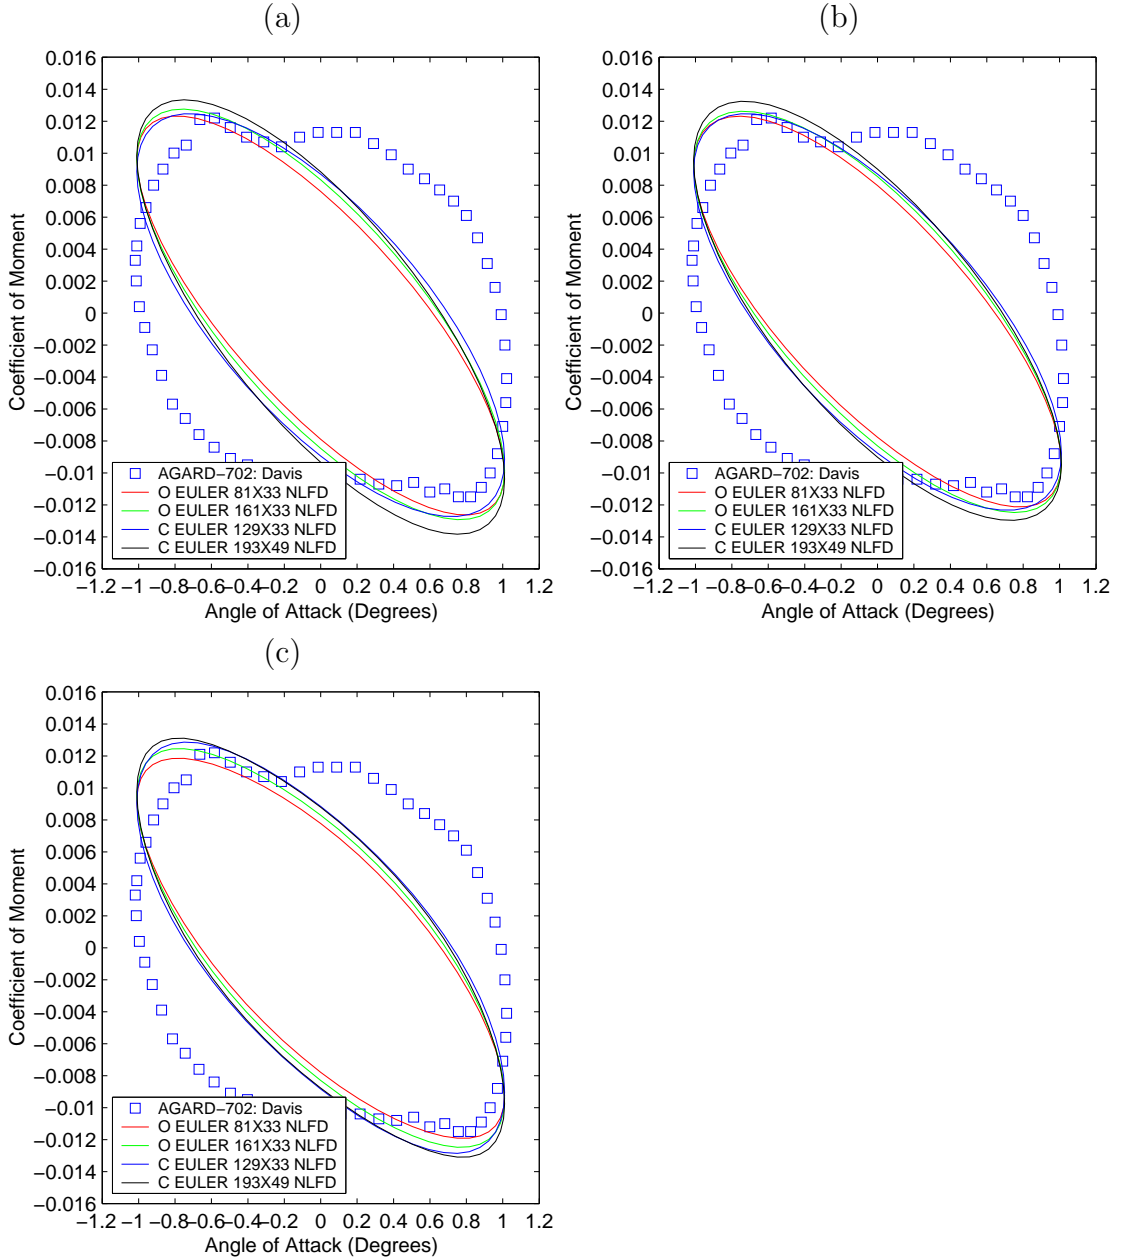


Figure 4.16: Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Euler NLFD calculations and Davis's 64A010 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics.

4.7 Unsteady Viscous Flow

This section compares the experimental results of Davis and Landon to numerical calculations by the NLFD code using the unsteady Navier-Stokes equations. The format of this section closely mirrors the format for the Euler results presented in section 4.6. For each viscous mesh presented in section 4.3 three separate solutions were calculated using one, two and three time-varying modes. Six cases, comprising two grids and three temporal resolutions, are provided for the 0012 airfoil. Nine cases, comprising three grids and three temporal resolutions, are supplied for the 64A010 airfoil.

Coefficient of Lift Results

Figures 4.17 and 4.18 provide C_l predictions as a function of temporal resolution, while figures 4.19 and 4.20 provide these quantities in terms of spatial resolution. To ease the comparison, the experimental results are overplotted for all the cases. The most striking feature of the results is that they exhibit little variation over the range of temporal resolutions. The viscous and inviscid results are consistent in that one temporal mode provides a solution convergent to plotting accuracy. While there are some variations in the results as a function of spatial resolution for the 64A010 airfoil, these effects are small and provide no visible trend toward improving the comparison between experimental and numerical results.

This section also provides the magnitudes and phase lags of the lift transfer functions in tables 4.12 through 4.14. For the zeroth harmonic, numerical predictions of the lift transfer function show excellent agreement with experimental results with little difference observed as time varying modes are added. The improvement of these unsteady viscous calculations over the unsteady inviscid calculations, resembles a similar trend between the steady inviscid results in section 4.4 and the steady viscous results in section 4.5. This confirms that good approximations to the mean lift transfer function require viscous simulations.

For the fundamental harmonic, the magnitude of the 0012 lift transfer functions vary on average only 2.2 percent from the experimental results. Using coarser grids

Mesh Type	1 Harmonic	2 Harmonic	3 Harmonic
Experimental	0.143		
C-mesh 257X49	0.142	0.143	0.143
C-mesh 513X97	0.144	0.144	0.144

Table 4.12: Magnitude of zeroth harmonic in coefficient of lift divided by the magnitude of the zeroth harmonic in angle of attack for various temporal resolutions employed by the NLFD method $\left(\frac{\|\hat{C}_{l_0}\|}{\|\hat{\alpha}_0\|}\right)$.

for the 64A010 airfoil this deviation is on average 5.5 percent. No distinct trends are found in the data as spatial resolution is varied. Adding time varying modes into the calculation of the solution provides no significant difference.

To determine the impact of viscosity, these results can be compared to the equivalent inviscid values in table 4.10. For the 0012 cases, the error approximately doubles as one adds viscosity while the opposite is true for the 64A010 cases where the error is halved. Unlike the time-average, the data for the fundamental harmonic is inconclusive with respect to the fundamental harmonics dependency on viscous effects.

Airfoil	Mesh Type	1 Harmonic	2 Harmonic	3 Harmonic
0012	Experimental	0.121		
0012	C-mesh 257X49	0.119	0.119	0.119
0012	C-mesh 513X97	0.118	0.118	0.117
64A010	Experimental	0.095		
64A010	C-mesh 129X33	0.104	0.103	0.103
64A010	C-mesh 193X49	0.097	0.097	0.097
64A010	C-mesh 257X65	0.101	0.101	0.101

Table 4.13: Magnitude of fundamental harmonic in coefficient of lift divided by the magnitude of the fundamental harmonic in angle of attack for various temporal resolutions employed by the NLFD method $\left(\frac{\|\hat{C}_{l_1}\|}{\|\hat{\alpha}_1\|}\right)$.

In contrast to the magnitude predictions, the addition of the viscous terms moves the phase lag estimates farther from the experimental results than similar estimates

provided by Euler calculations. The average deviation between the NLFD predictions and the experimental results is 5.2° for the 64A010 airfoil and -4.1° for the 0012 airfoil. These estimates represent an increase in the average deviation from the Euler predictions. Little variation in the phase lag is exhibited as a function of the temporal resolution, and no consistent trend is found as a function of spatial resolution.

Airfoil	Mesh Type	1 Harmonic	2 Harmonic	3 Harmonic
0012	Experimental	-10.9		
0012	C-mesh 257X49	-14.7	-15.4	-15.2
0012	C-mesh 513X97	-14.6	-15.0	-14.9
64A010	Experimental	-25.6		
64A010	C-mesh 129X33	-20.3	-20.0	-20.1
64A010	C-mesh 193X49	-20.9	-20.9	-20.9
64A010	C-mesh 257X65	-20.0	-20.0	-20.0

Table 4.14: Magnitude of phase shift between the fundamental harmonic in lift coefficient and angle of attack for various temporal resolutions employed by the NLFD method $(\angle \hat{C}_{l_1} - \angle \hat{\alpha}_1)$.

Coefficient of Moment Results

Figures 4.21 and 4.22 provide C_m predictions as a function of temporal resolution, while figures 4.23 and 4.24 provide this quantity in terms of spatial resolution. The experimental results are overplotted for all the cases to ease the comparison with the numerical results.

As in the case of the Euler calculations, significant variation is exhibited between the temporal resolutions in the results for the 0012 airfoil. In comparison to the Euler results, the addition of viscous terms improves the agreement between the numerical and experimental results. However, the 64A010 results show little variation as a function of temporal resolution indicating that one mode is adequate. In either case, the variation between the turbulent Navier-Stokes calculations and experimental results is still significant and should be a topic of further investigation.

Little conclusive evidence is offered by these figures in regards to the variation

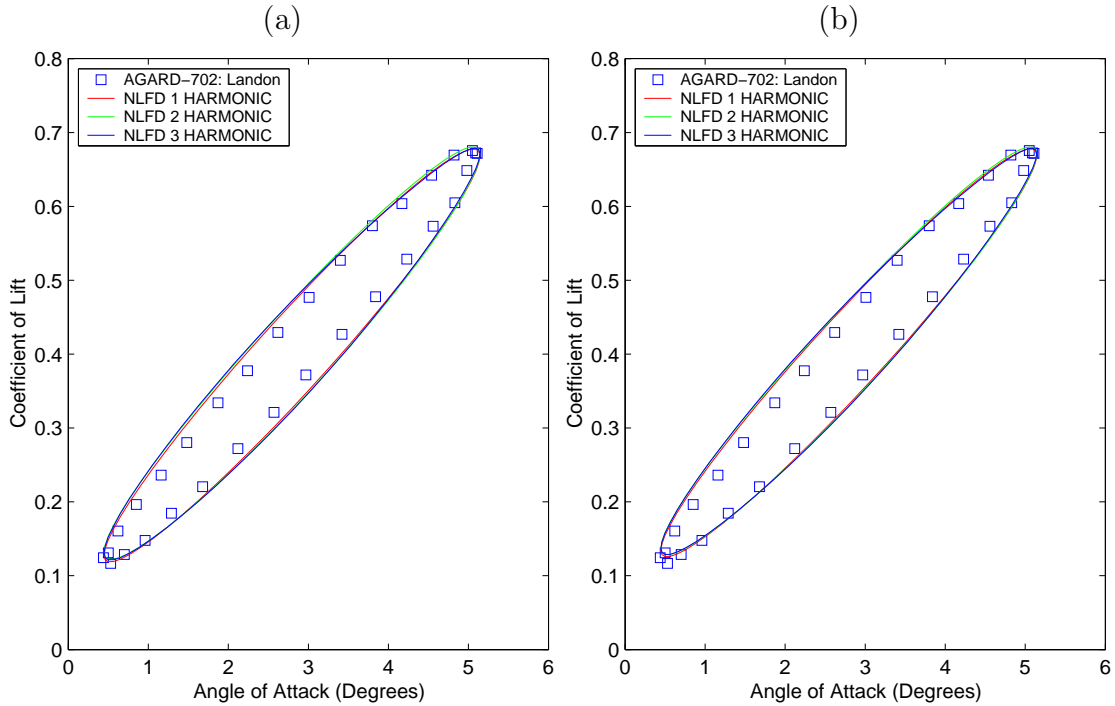


Figure 4.17: Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Landon’s 0012 experiment. (a) 257x49 C-mesh. (b) 513x97 C-mesh.

in C_m results as a function of spatial resolution. There is some variation in the results between the different meshes, but no consistent trend is observed between grid refinement and improving the agreement with experimental results.

4.8 Summary

The NLFD method has been used to solve both the Euler and Reynolds averaged Navier-Stokes equations for flow around a pitching airfoil. The numerical results have been compared with experimental data for both steady and unsteady flow. The accuracy of the results is obviously a function of the resolution of the discretization and the choice of the governing equations, but some generalizations can be made. Increasing the spatial resolution had little effect on the results suggesting that the

grids are sufficiently fine. Excluding the C_m results for the 0012, in most of these cases there is no discernible benefit in using more than one temporal mode. In the case of coefficient of moment, the comparison between the numerical and experimental results is poor. Results presented in chapter 5 of this thesis show that this is no different than time-accurate solvers with similar spatial discretizations. Consequently, it may be concluded that the discretization of the temporal derivative is not the source of the discrepancy, which may partly be due to turbulence-modeling errors.




Figure 4.18: Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Davis's 64A010 experiment. (a) 129x33 C-mesh. (b) 193x49 C-mesh. (c) 257x65 C-mesh.

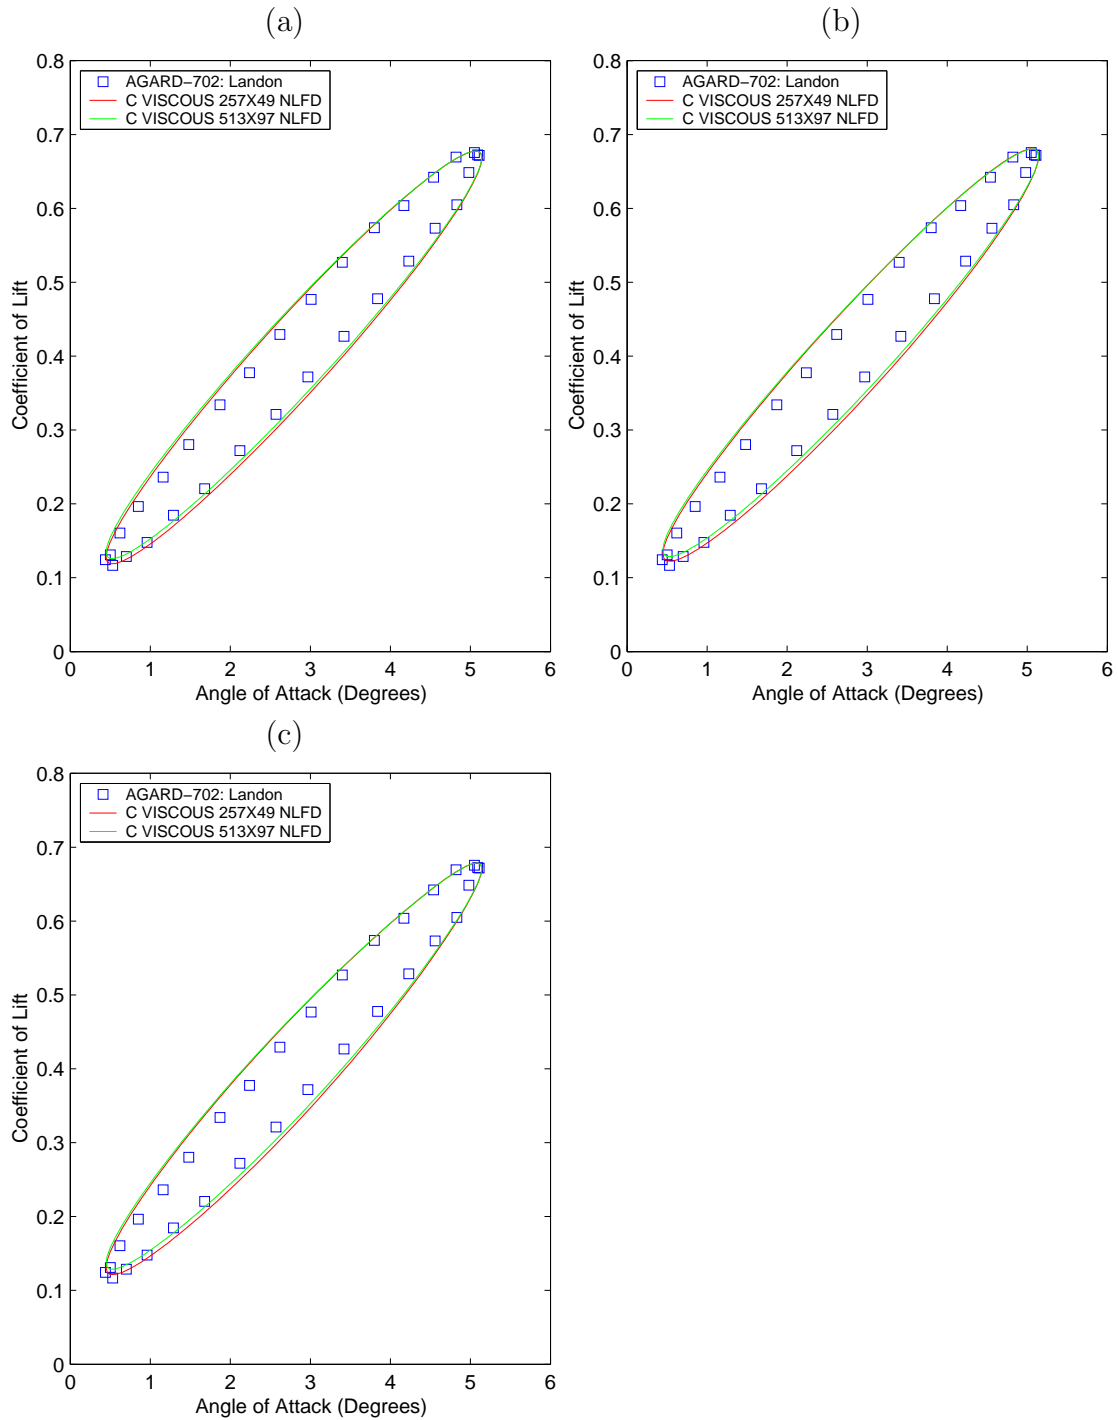


Figure 4.19: Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Landon's 0012 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics.

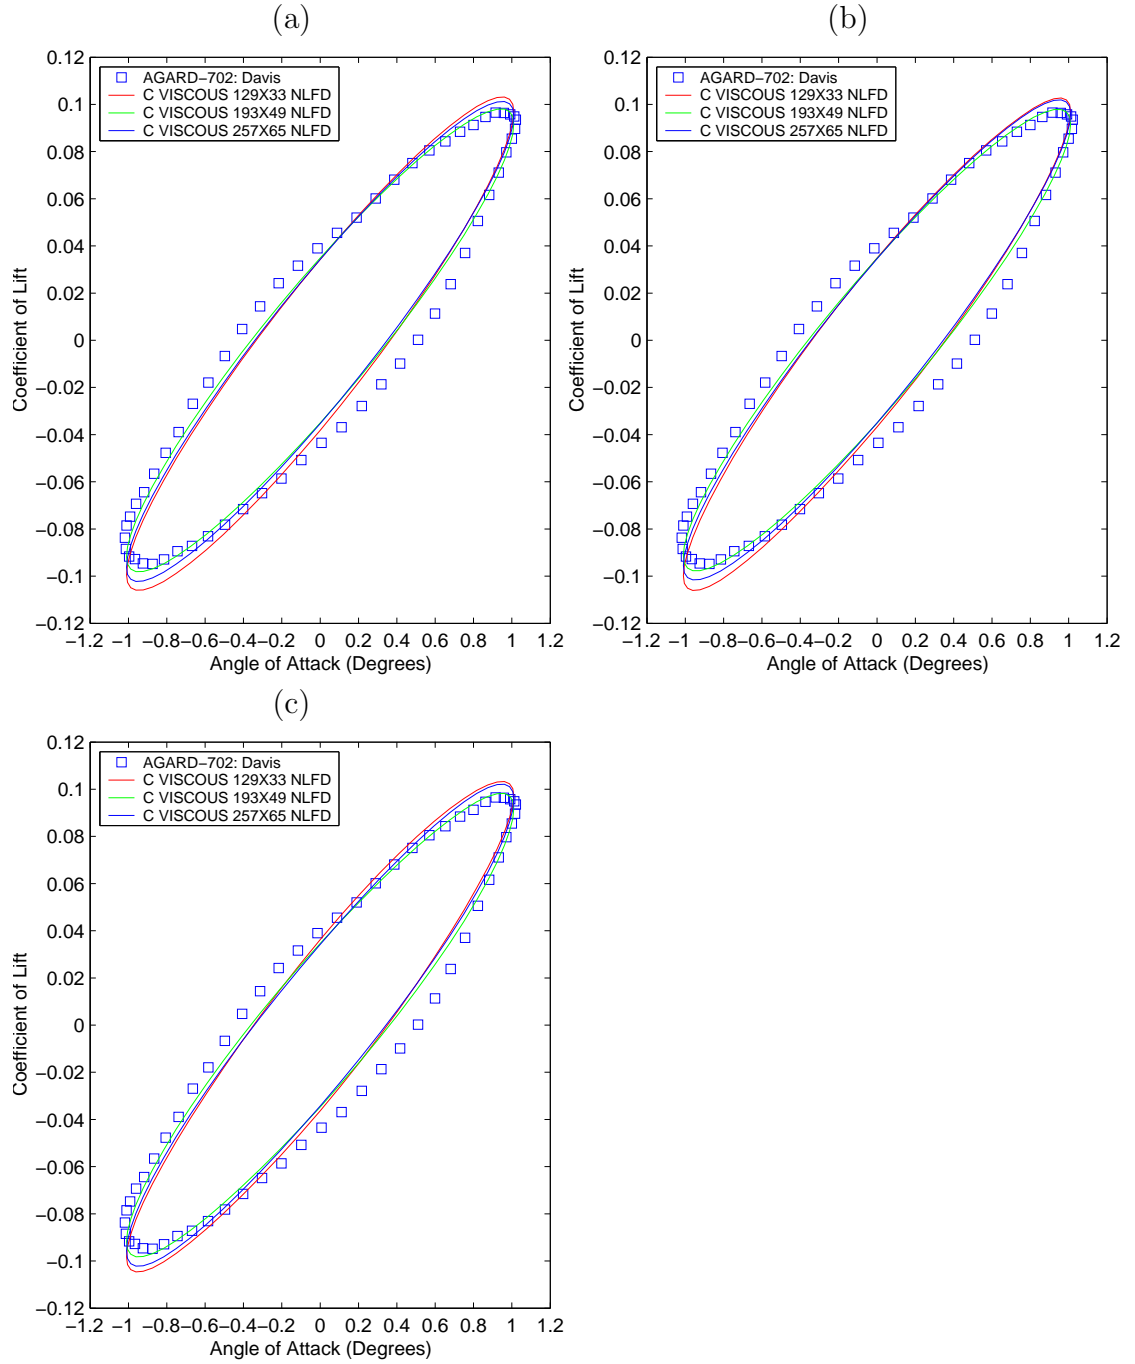


Figure 4.20: Lift coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Davis's 64A010 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics.

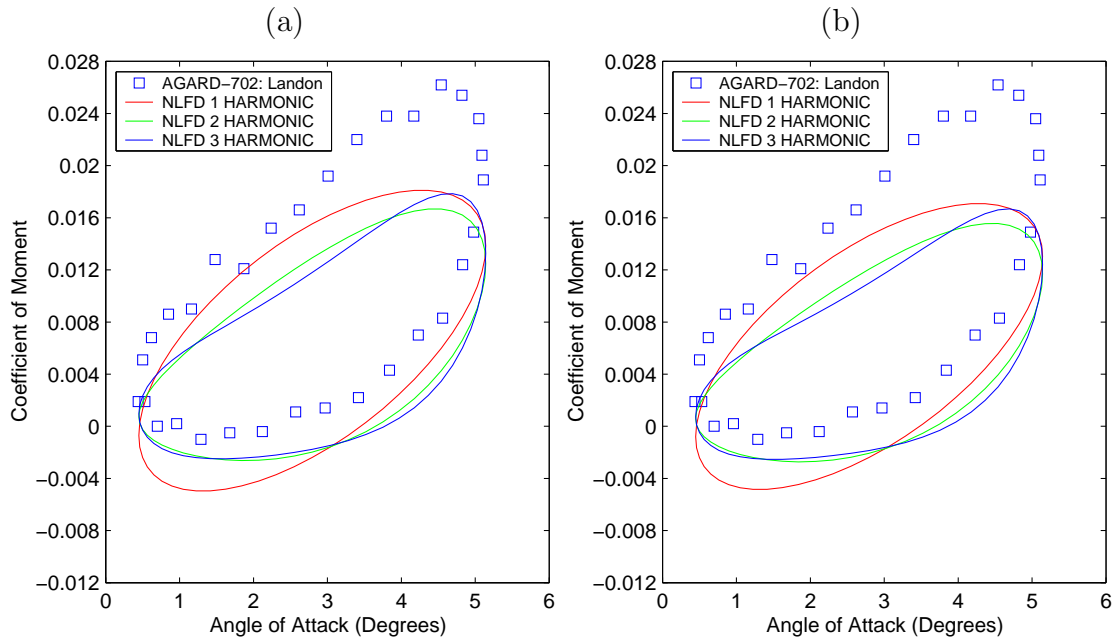


Figure 4.21: Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Landon's 0012 experiment. (a) 257x49 C-mesh. (b) 513x97 C-mesh.

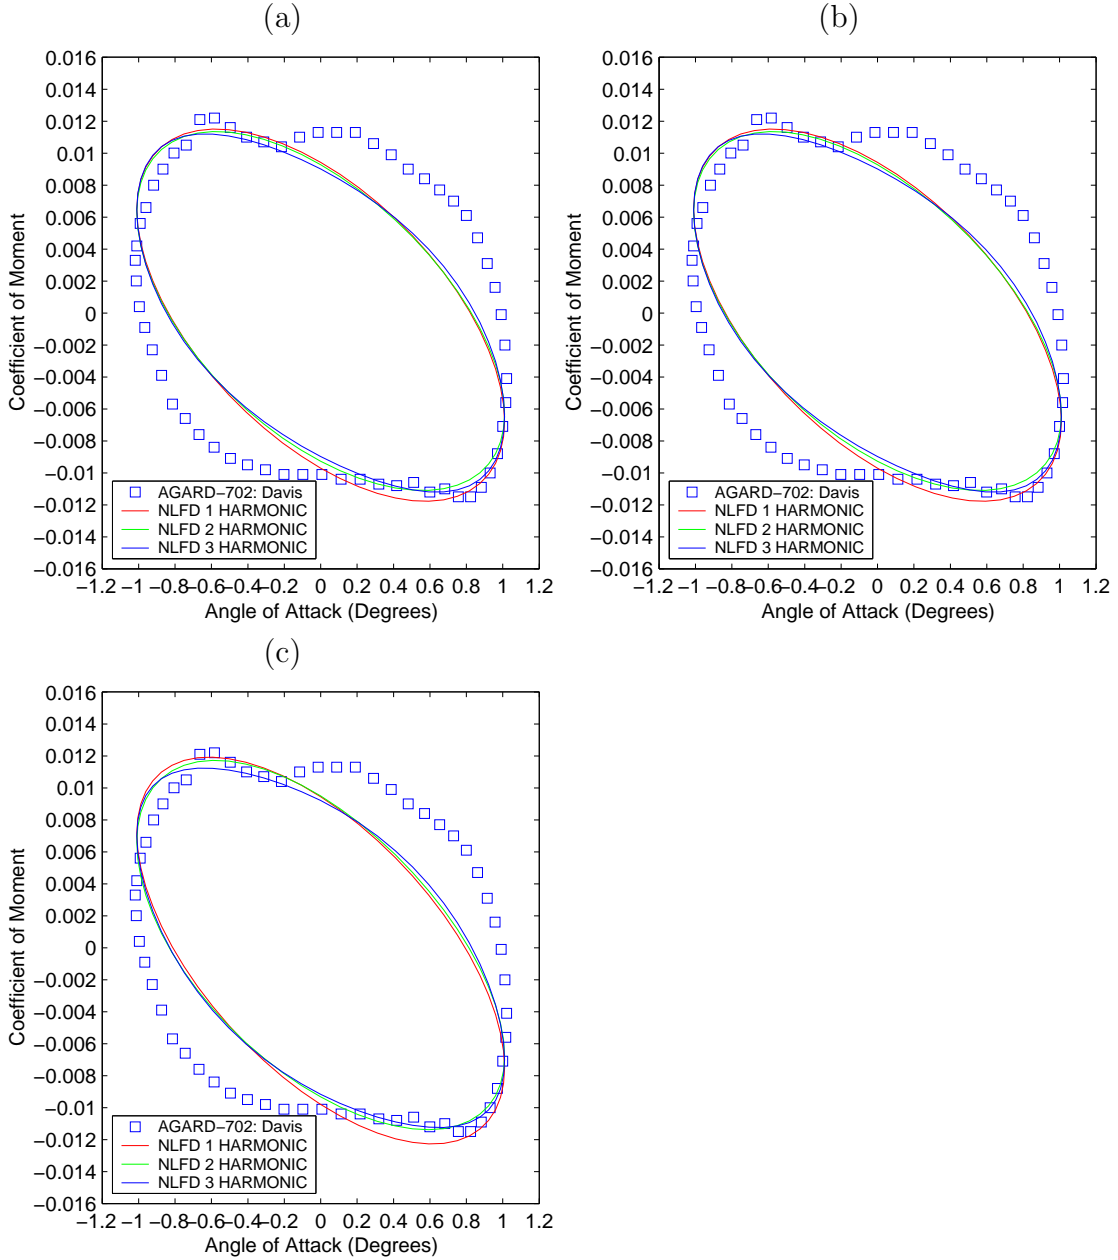


Figure 4.22: Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Davis's 64A010 experiment. (a) 257x49 C-mesh. (b) 513x97 C-mesh. (c) 257x65 C-mesh.

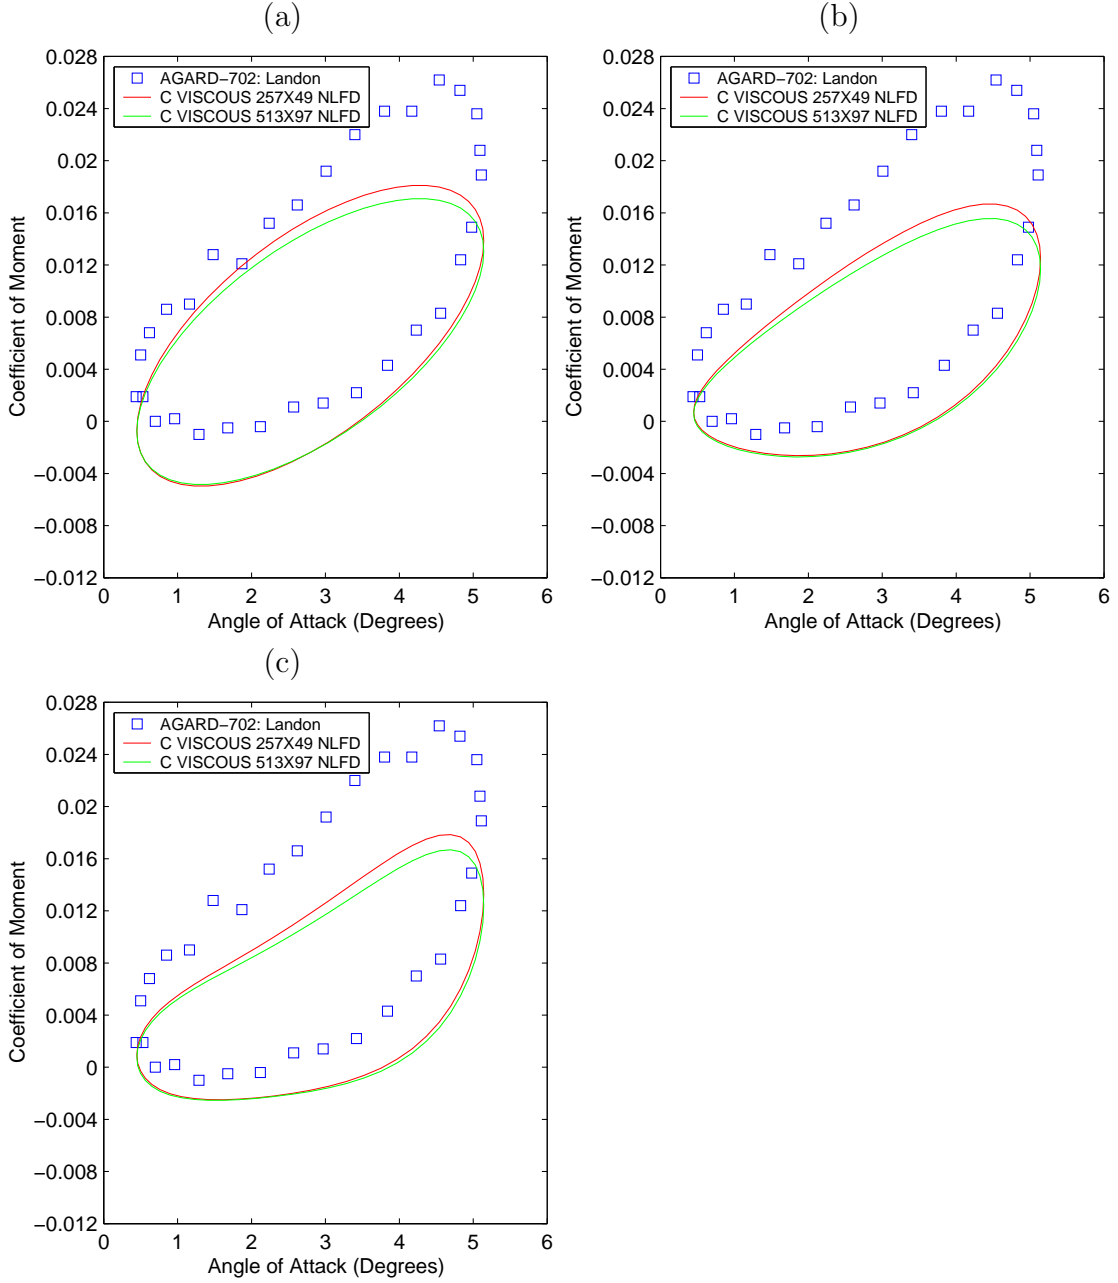


Figure 4.23: Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Landon's 0012 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics.

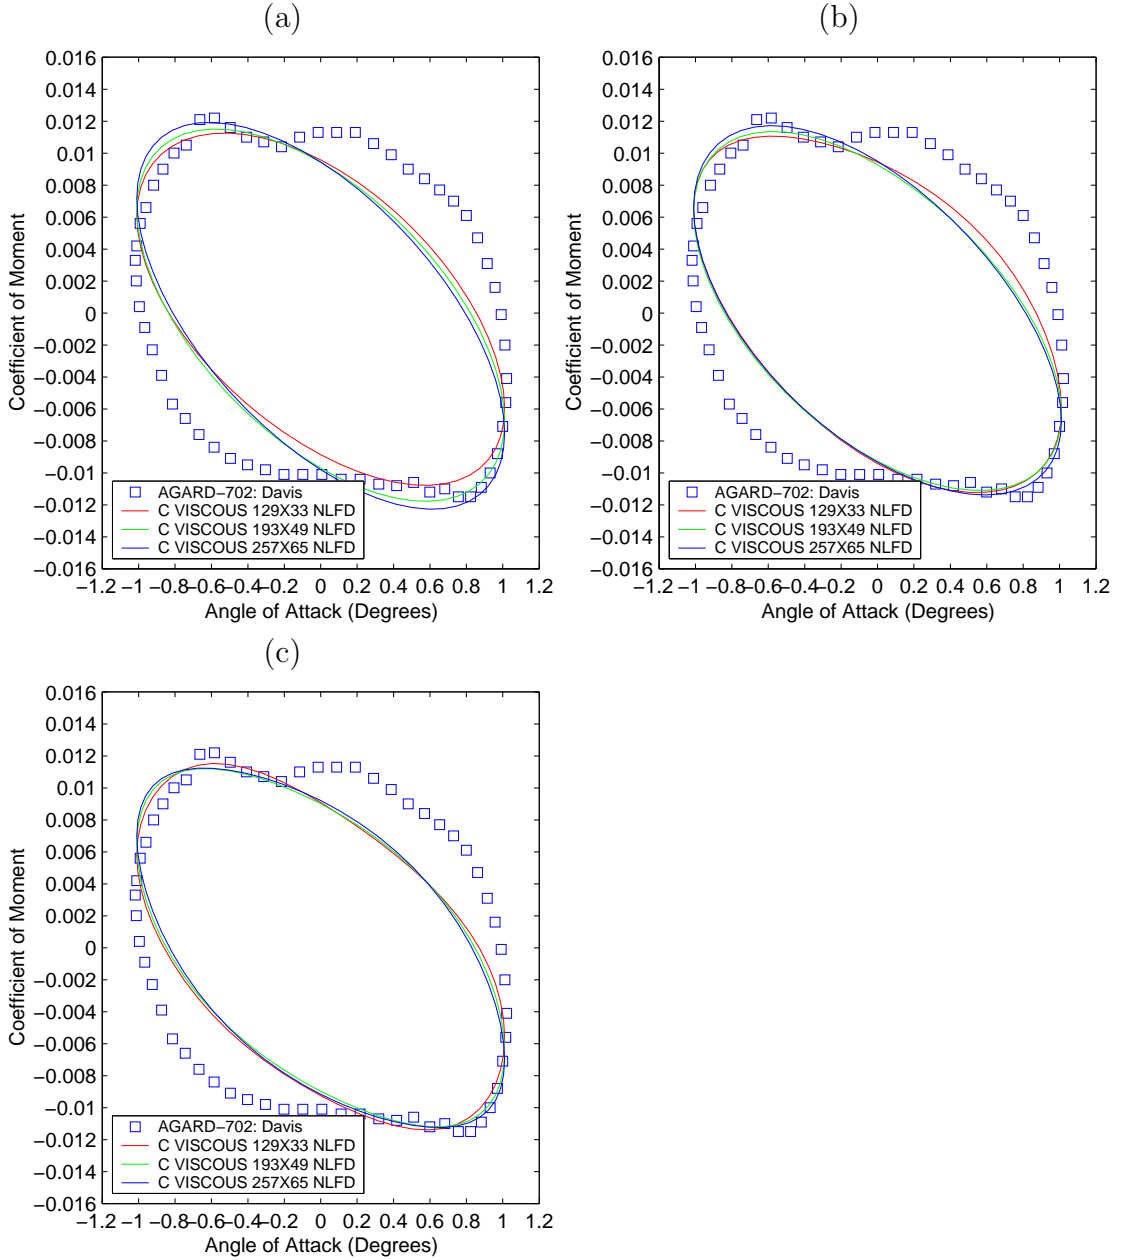


Figure 4.24: Moment coefficients as a function of the instantaneous angle of attack provided from unsteady Navier-Stokes NLFD calculations and Davis's 64A010 experiment. (a) 1 Harmonic. (b) 2 Harmonics. (c) 3 Harmonics.

Chapter 5

Pitching Airfoil Numerical Validation

A crucial question is whether the Non-Linear Frequency Domain method offers a major reduction in computational cost over established time-accurate methods. This chapter addresses this question by presenting a comparison of a state of the art time-accurate code UFLO82 [34, 58], with an NLFD code for Euler simulations of a pitching airfoil. In order to make a fair comparison the boundary conditions and spatial operators of the NLFD code have been replaced with exactly those of the UFLO82 code. With equivalent spatial errors the two codes can then be compared solely on their approximation of the temporal derivative in the conservation equations.

A convergence study is first performed to quantify the error as a function of temporal resolution via time steps of modes. Using equivalent overall error levels the two codes are then compared on the basis of computational cost. This is the inverse of Pulliam's approach outlined in section 1.2 where he fixed the cost and compared methods on basis of error.

5.1 Choice of a Test Case

In the comparison of experimental and numerical results in the previous sections the flows contained only moderately strong shocks. In order to evaluate the NLFD

method in the presence of strong nonlinear effects the case used in this section is more challenging. The reduced frequency and the dynamic angle of attack have been chosen to promote nonlinearities that result in a coupling of temporal modes within the flow field. Table 5.1 provides a list of the parameters defining this case.

Parameter	Value
Airfoil	NACA 64A010
Mean Angle of Attack	0.0°
Angle of Attack Variation	$\pm 2.0^\circ$
Mach Number	0.8
Reduced Frequency	0.05

Table 5.1: Description of the test case used in the comparison of UFLO82 and the NLFD codes.

5.2 Computational Grid

Due to the mesh topology requirements of the UFLO82 code the grid used in this section is a 161x33 O-mesh. This mesh is substantially different from the 161x33 grid used in the experimental validation. Specifically the average height of the cell adjacent to the airfoil surface is 0.0048 chords, roughly half of the height of the cell contained in grids used in the experimental comparisons. Also the average farfield boundary distance for this grid is 129 chords approximately six times the farfield distance in the previous grids. Figure 5.1 provides a picture of the nearfield distribution of cells within this mesh. The mesh was created using the algebraic grid generation procedures of UFLO82.

5.3 Comparison With UFLO82 For Steady Flow

The first test is to verify the exact equivalence of the spatial discretization and boundary conditions of the modified NLFD code with those contained in the UFLO82 code. Both UFLO82 and the modified NLFD code employ a cell centered finite volume




Figure 5.1: Nearfield resolution of 161x33 O-mesh used in the numerical comparison of UFLO82 and NLFD codes.

scheme with JST artificial dissipation, a linear pressure extrapolation at the wall, and Riemann invariants in the farfield boundary condition [28, 29, 31, 32, 33].

By default, the NLFD code has the capability to do steady calculations. However UFLO82 required a slight modification; the coefficients used in the temporal derivative were set to zero to nullify any unsteady effects. For both codes, the magnitude of the dynamic angle of attack was set to zero to nullify any grid rotation terms. One set of multigrid cycles was used to drive the steady residual to a negligible value. The output from both codes was a steady solution at the time-averaged angle of attack and farfield boundary conditions.

Figure 5.2 plots the difference between steady solutions produced by these two codes. The four components of the different solutions match to machine accuracy, verifying the equivalence of the spatial operators and boundary conditions. Obviously, any derived quantity like C_l or C_m would match to machine accuracy as well.

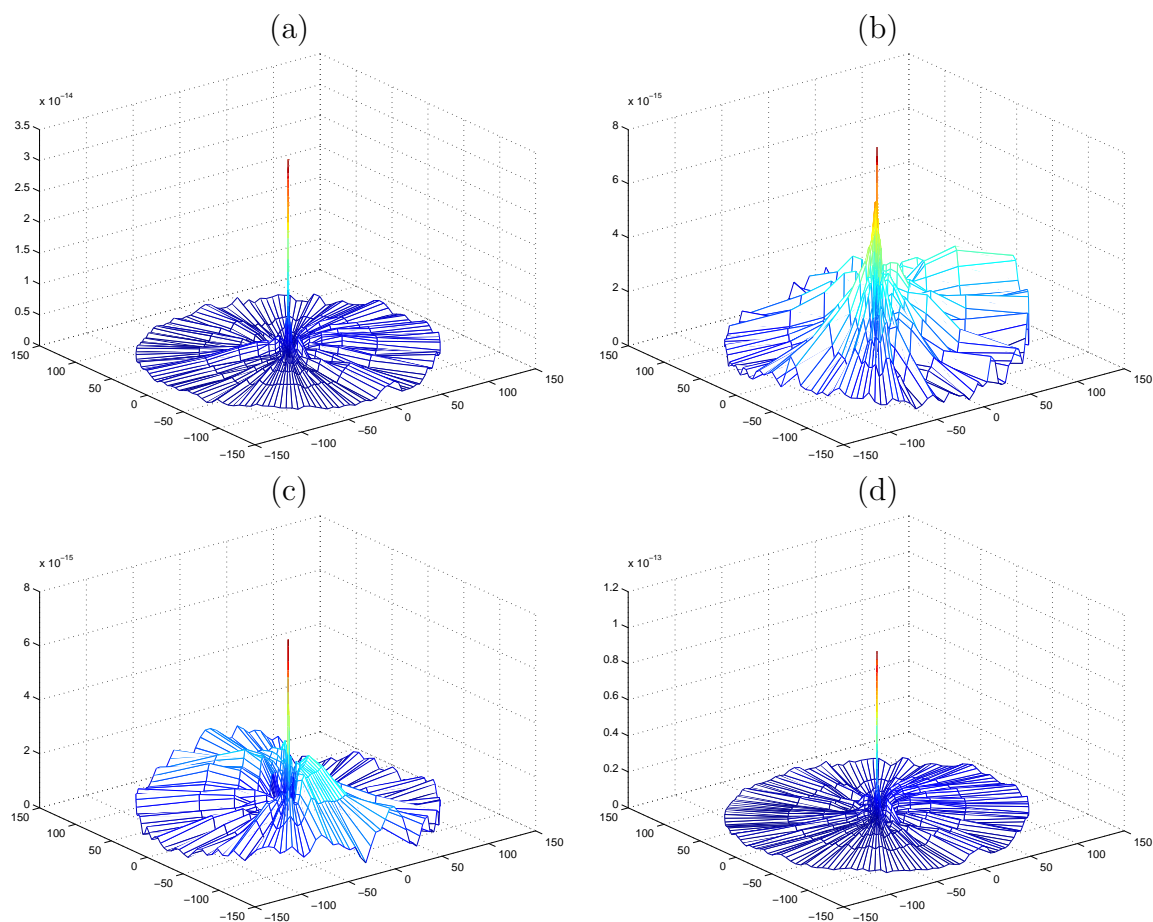


Figure 5.2: The absolute value of the difference between a UFLO82 and NLFD steady solution for Euler calculation on a 161x33 O-mesh. The different figures represent different components of the solution. (a) ρ (b) ρu (c) ρv (d) ρe .

5.4 Comparison With UFLO82 For Unsteady Flow

In order to substantiate the cost analysis, this section highlights the differences between UFLO82 and the modified NLFD code. UFLO82 uses a dual time stepping technique implementing a series of nested loops. The inner loop is a set of multigrid cycles used to drive the unsteady residual for that time step to a negligible value. This unsteady residual is written as the combination of an implicitly evaluated spatial residual and a backward difference formula approximating the temporal derivative. As presented earlier, UFLO82 uses the second order accurate A-stable scheme

$$\frac{3W^{n+1} - 4W^n + W^{n-1}}{2\Delta t} + R(W^{n+1}) = 0. \quad (5.1)$$

Typically, the solution at the end of an inner loop is not converged to machine accuracy. In this case, the marginal change in the fluid properties due to one additional multigrid cycle should be less than the accuracy/stability requirements for the calculation. The number of multigrid cycles used per time step ultimately controls the accuracy with which the discretized equations of motion are being solved. For frequency domain solvers, the magnitude of the unsteady residual for all the wavenumbers (or time instances in the temporal domain) is minimized by a series of multigrid cycles. The equations are simultaneously solved to a given accuracy level in a manner analogous to dual time stepping techniques.

The final residual levels are functions of the value of the initial residual level used to start the multigrid cycle and the rate of decay in the residual per multigrid cycle. The initial residual for dual time stepping techniques will become smaller as the number of points per wavelength is increased. In the limit as the time step goes to zero, simple extrapolation techniques provide excellent first guesses to the solution. For most practical cases, the initial residual level for the UFLO82 solver is one or two orders of magnitude less than the initial unsteady residual for NLFD solvers. Depending on the implementation, the decay rate of the residual per multigrid cycle may or may not be equivalent between the solvers.

The outer loop of the dual time stepping technique provides the time history of the discrete solution. A complete discussion of the error associated with the NLFD

and UFLO82 temporal discretizations is provided in section 2.2.3.

The final parameter affecting a time-accurate solution is its proximity to a periodic steady state. Time-accurate solvers capture the decay of initial transients. As these diminish the solution asymptotically approaches a periodic steady state. Eventually the marginal decay of these transients during each additional time step is within the accuracy requirements for the solution. When this condition has been met, the calculation is typically stopped. The cost of resolving the initial transients is eliminated in NLFD methods since these solvers admit only components of the solution which are periodic over a predefined length of time.

Table 5.2 itemizes the costs associated with the dual time stepping method, and table 5.3 does the same for the NLFD method. The final cost analysis presented at the end of this section quantifies each of the terms presented in these two tables for test cases producing equivalent error.

Description	Variable
Number of multigrid cycles required for a solution at each time step.	N_M
Number of time steps (or solutions) required to resolve an oscillation in the solution.	N_N
Number of periods needed to reach a periodic steady state	N_P
Overall Cost	$Cost \approx N_M N_N N_P$

Table 5.2: Cost estimation for dual time stepping codes like UFLO82.

Description	Variable
Number of multigrid cycles to solve the unsteady residual in the frequency domain.	N_M
The ratio of the cost of an NLFD solution to a steady state solution. This is approximately the number of solution instances per time period.	N_N
Overall Cost	$Cost \approx N_M N_N$

Table 5.3: Cost estimation for frequency domain codes like NLFD.

Temporal Resolution

In order to quantify the required number of time steps per oscillation N_N , a convergence study in temporal accuracy was performed for both codes running in unsteady mode according to the parameters listed in table 5.1. For UFLO82 the residual at the end of each time step was driven to machine zero. This code was run at 8 different temporal resolutions including 8, 16, 24, 32, 64, 128, 256 and 512 points per wavelength. To eliminate any error associated with initial transients, 48 periods of the airfoil oscillation were simulated. As the solution evolved in time, Fourier transforms of the latest period of C_l and C_m data were computed. The magnitude of the fundamental harmonic of these data sets is used to quantify the error. The control solution was selected from the last oscillation of the most accurate solution (512 points per wave), and is subtracted from all other solutions to quantify the error due to the temporal resolution E_N

$$\begin{aligned} E_{N_l} &= \left| |\hat{C}_{l_1}| - |\hat{C}_{l_{512}}| \right| \\ E_{N_m} &= \left| |\hat{C}_{m_1}| - |\hat{C}_{m_{512}}| \right|. \end{aligned} \quad (5.2)$$

Similar runs were made with the NLFD code using one, two and three time varying modes. All components of the unsteady residual were driven to zero, and Fourier transforms of C_l and C_m were computed from the periodic steady state solution. The magnitude of the fundamental harmonic was then subtracted from the control solution in a manner consistent with the UFLO82 results.

Figures 5.3 and 5.4 shows the error for both lift and moment coefficients. The plots include the time history of error for the calculation producing the control solution, showing the difference in the figure of merit between the converged solution and the solution at any point in time. This is included to show that the error associated with the initial transients for the control calculation is significantly less than the overall error for any other calculation, and hence can be considered as insignificant.

Each solution approaches an error level representing the minimum error achievable in realistic environments. These minimum errors are replotted on figures 5.5 and 5.6 as

a function of the temporal resolution used by each UFLO82 calculation. Overplotted on these figures are horizontal lines from NLFD calculations with varying numbers of temporal modes. The subfigures provide the same data in different logarithmic formats while quantifying the horizontal axis differently. As expected the loglog plot shows that the error decay rate is proportional to Δt^2 .

The intersection of the NLFD lines with the UFLO82 curve shows the temporal resolutions each code must use to achieve equivalent error levels. This data is summarized in table 5.4. Using C_l as a figure of merit, the NLFD code needs only 1 temporal mode to obtain similar error levels as a UFLO82 code using 45 points per wavelength. Using C_m in the same capacity, the UFLO82 code needs only 18 points for one mode in the NLFD code. Regardless of the figure of merit, for every additional temporal mode added in the NLFD calculation, the temporal resolution of UFLO82 needs to be increased by roughly a factor of 2.5 to 3.

NLFD Modes	C_l UFLO82 SPW	C_m UFLO82 SPW
1	45	18
2	125	45
3	244	123

Table 5.4: The number of solutions per wavelength required by UFLO82 to reach error levels equivalent to results produced by an NLFD code using one, two, and three time varying harmonics.

Multigrid Cycles

In an attempt to quantify N_M for UFLO82, a convergence study was performed to identify the number of multigrid cycles per time step required to reach the asymptotic error levels identified in figures 5.3 and 5.4. These error levels were functions of UFLO82 calculations where the residual was driven to machine zero at each time step. However, the residual error needs only to be small relative to the error due to the temporal discretization. Consequently an unbiased cost comparison of the methods should quantify the minimum number of multigrid cycles per step required

by UFLO82 to reach the equivalent error levels. Figure 5.7 provides this data, for both C_l and C_m , for each temporal resolution identified in table 5.4. In most cases, only six multigrid cycles per step were required.

This number is surprisingly small, but may be due to the properties of the UFLO82 prediction method and the concave curvature of the C_l data. Table 5.5 shows that the effect of temporal dissipation is to lower the magnitude of the unsteady lift coefficient away from the more accurate results. UFLO82 uses a linear extrapolation method to predict the solution at each new time step. Based on the curvature of the C_l time history, the initial prediction will be closer to the more accurate solution. Additional multigrid cycle will lower the residual, heighten the effect of the dissipation and move the solution away from the more accurate result. This theory is reinforced by figure 5.9 which shows that increasing the number of multigrid cycles per time step has in some cases a detrimental effect.

Solutions Per Wave	8	16	24	32	64	128	256	512
$ C_l $	0.1657	0.1740	0.1759	0.1765	0.1771	0.1773	0.1773	0.1773

Table 5.5: Variation in UFLO82 unsteady coefficient of lift as a function of temporal resolution.

A similar convergence study was performed for the NLFD code. Figure 5.8 plots the error level as a function of the number of multigrid cycles. The intersection of the dashed line with the individual curves approximately identifies the minimum number of cycles required to reach the asymptotic error levels. These intersection points are quantified in table 5.6.

The UFLO82 results are quantified in terms of multigrid cycles per step, while the NLFD results are quantified in the overall number of cycles used. For consistent comparisons, this latter result needs to be multiplied by an additional factor to account for the work performed for each additional temporal mode. This factor will be identified in the section that immediately follows.

NLFD Modes	C_l MG Cycles	C_m MG Cycles
1	44	37
2	60	50
3	72	65

Table 5.6: The number of multigrid cycles required by the NLFD solver to obtain a solution at equivalent error levels.

Cost per Multigrid Cycle

NLFD solvers calculate instances of the spatial residual along the time period of the solution. Consequently the work per multigrid cycle scales with the number of time instances used to represent the solution. If N is the number of time varying modes, then the work per multigrid cycle should grow like $2N - 1$. The NLFD solver incurs the additional cost of the transformation of the solution and residual between the temporal and frequency domains by Fast Fourier Transform. Table 5.4 provides measured data for the execution time per NLFD multigrid cycle divided by the time required for a similar steady state cycle. The data approximates the growth rate cited above; substantiating the claim that the cost of the FFT is small in comparison to the total work done in the multigrid cycle.

	0 th mode	1 st mode	2 nd mode	3 rd mode
Normalized execution speed (time per cycle)	1.0	3.06	5.13	7.30

Table 5.7: NLFD execution time per multigrid cycle as a function of temporal resolution. Results are normalized by the execution time required by the steady state or zeroth mode case.

Decay of Initial Transients

The final factor in the cost of the UFLO82 solution is the time rate of decay of the initial transients. Figures 5.9 and 5.10 plot the decay in the error over time for

each temporal resolution identified in table 5.4 and each multigrid cycle permutation provided in figure 5.7. Based on this data, table 5.8 provides estimates of the number of periods required for the error to approach its asymptotic value. Given the curves plotted in figures 5.9 and 5.10, these numbers were selected conservatively to the benefit of the cost estimation for UFLO82.

C_l		C_m	
SPW	Periods	SPW	Periods
45	4	18	3
125	5	45	4
244	7	123	6

Table 5.8: The number of time periods required by the UFLO82 solver to reach convergence. Separate data is provided for each temporal resolution listed as SPW or Solutions Per Wavelength.

Cost Comparison

The cost of the two solvers can now be compared using the data provided in the previous sections. This comparison assumes that the work associated with a single multigrid cycle is equivalent between the two codes. This is an accurate approximation, given that UFLO82 and the NLFD code use the same pseudo-time advancement, residual averaging and multigrid aggregation and prolongation operators. This approach provides a relative comparison independent of the machine or compiler optimization algorithms. These latter parameters affect the cost of each method, but code implementation and profiling details are not the focus of this research.

Tables 5.9 and 5.10 provide the relative cost data using C_l and C_m as the figures of merit. If the user requires equivalent error levels for both C_l and C_m data, then the lift results will drive the cost comparison. In this case, the NLFD code is a factor of 8 to 19 times faster than the UFLO82 code depending on the error level. In the worst case, using only C_m error as the basis for comparison, the NLFD code at the lowest temporal resolution is approximately 3 times faster the UFLO82 code. Not surprisingly, the marginal cost of the NLFD solver is also better than the UFLO82

code. For either figure of merit, the cost multiple between the two codes increases in favor of the NLFD method as the number of modes is increased.

NLFD Modes	$N_M \times N_N = Cost$	UFLO82 SPW	$N_M \times N_N \times N_M = Cost$
1	$44 \times 3.06 = 135$	45	$6 \times 45 \times 4 = 1080$
2	$60 \times 5.13 = 308$	125	$6 \times 125 \times 5 = 3750$
3	$72 \times 7.30 = 526$	244	$6 \times 244 \times 7 = 10248$

Table 5.9: Cost comparison between UFLO82 and NLFD codes using the error in coefficient of lift as the figure of merit.

NLFD Modes	Cost Multigrid Cycles	UFLO82 SPW	Cost Multigrid Cycles
1	$3.06 \times 37 = 113$	18	$18 \times 6 \times 3 = 324$
2	$5.13 \times 50 = 257$	45	$45 \times 6 \times 4 = 1080$
3	$7.30 \times 65 = 475$	123	$123 \times 6 \times 6 = 4428$

Table 5.10: Cost comparison between UFLO82 and NLFD codes using the error in coefficient of moment as the figure of merit.

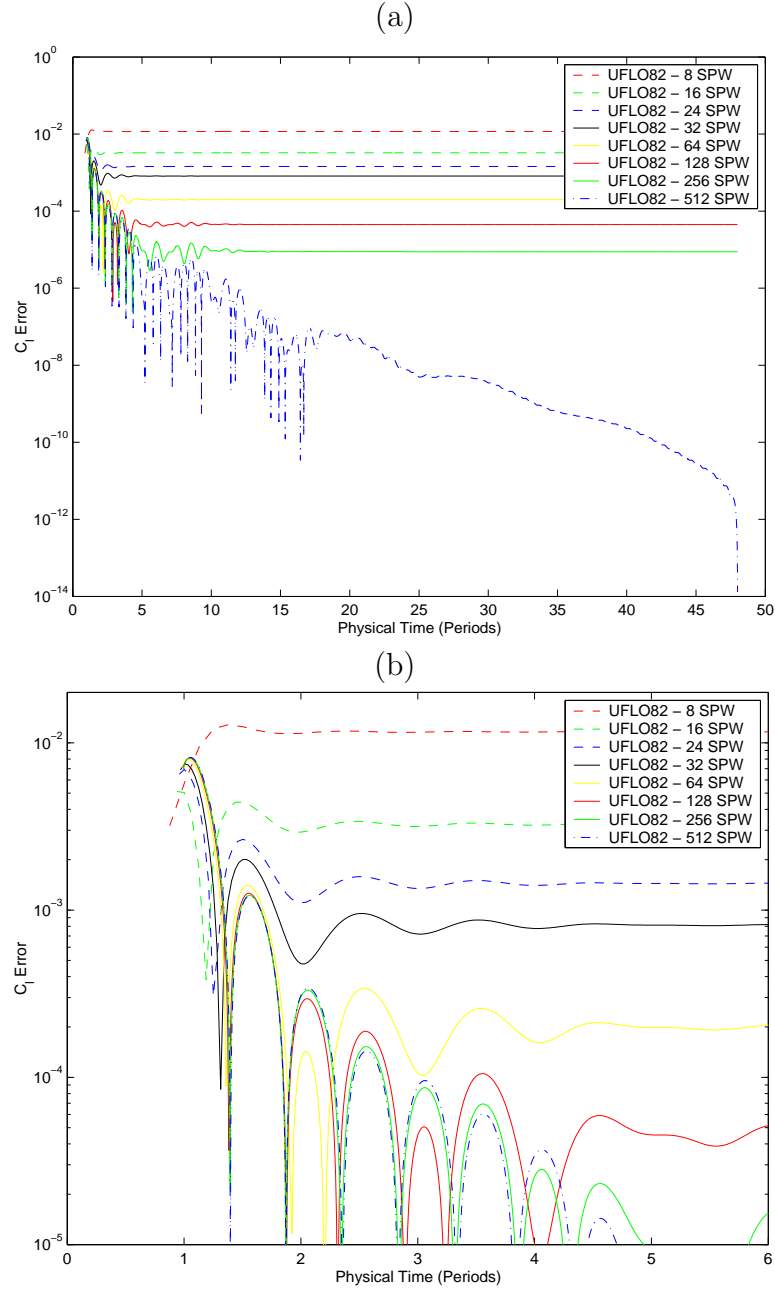


Figure 5.3: C_{l_1} error as defined by equation 5.2 between various temporal resolutions of the UFLO82 code and the control solution. The control solution is based on the coefficient of lift extracted from the final period of the most accurate solution (512 solutions per wavelength). (a) Entire time history of all data. (b) Magnified version of the same data focusing on less accurate solutions.

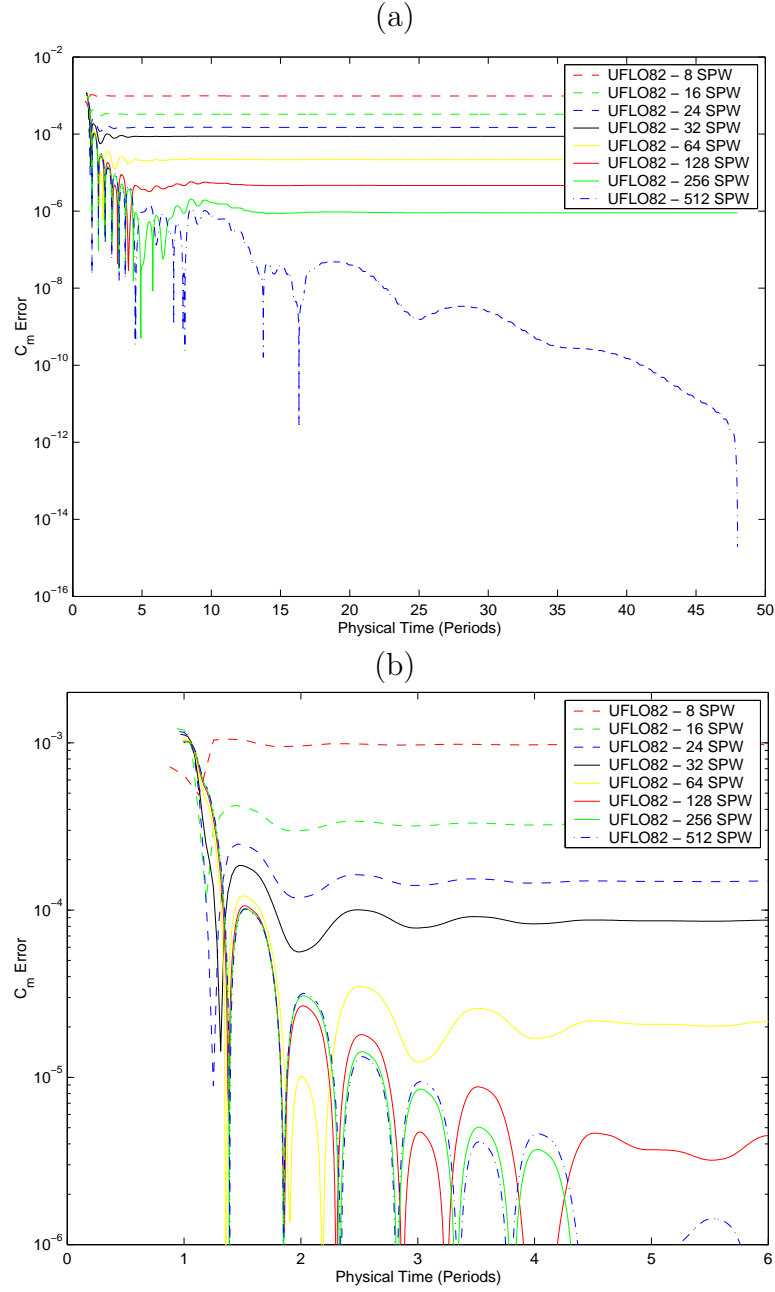


Figure 5.4: C_{m1} error as defined by equation 5.2 between various temporal resolutions of the UFLO82 code and the control solution. The control solution is based on the coefficient of moment extracted from the final period of the most accurate solution (512 solutions per wavelength). (a) Entire time history of all data. (b) Magnified version of the same data focusing on less accurate solutions.




Figure 5.5: Asymptotic error in C_{l_1} computed at various temporal resolutions by the UFLO82 and NLFD codes. (a) C_{l_1} Error as a function of Δt on a fully logarithmic axis. (b) C_{l_1} Error as a function of the number of solutions used to resolve an oscillation in the solution using a logarithmic vertical scale.

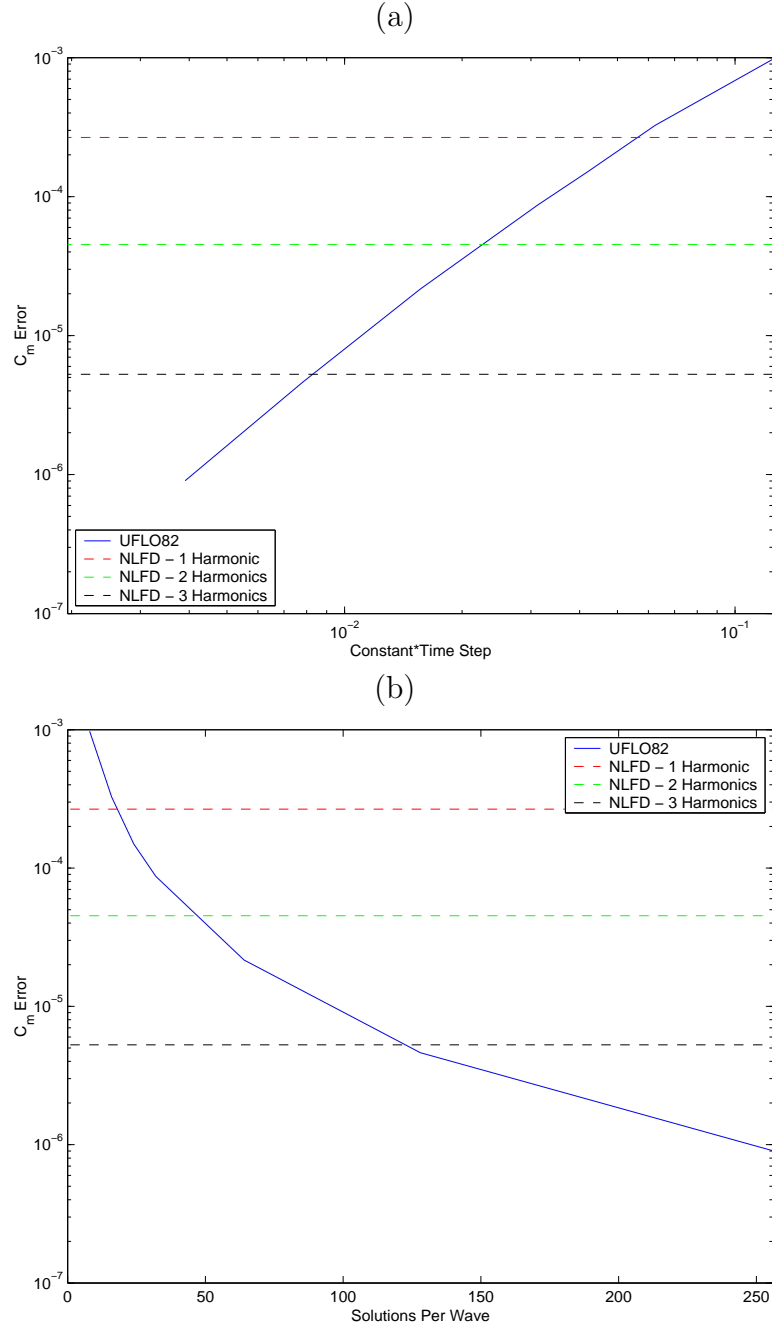


Figure 5.6: Asymptotic error in C_{m_1} computed at various temporal resolutions by the UFLO82 and NLFD codes. (a) C_{m_1} Error as a function of Δt on a fully logarithmic axis. (b) C_{m_1} Error as a function of the number of solutions used to resolve an oscillation in the solution using a logarithmic vertical scale.

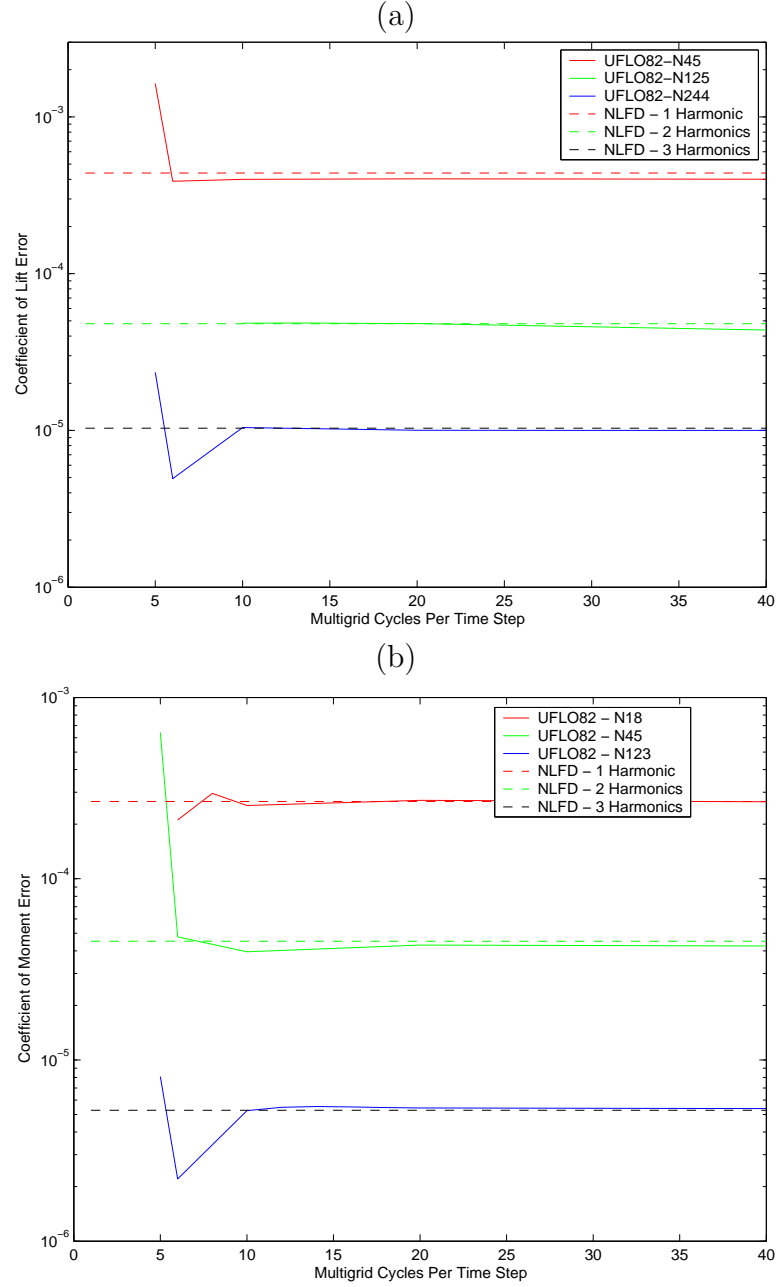


Figure 5.7: The number of multigrid cycles required by the UFLO82 code to reach equivalent error levels. (a) C_{l_1} . (b) C_{m_1} .

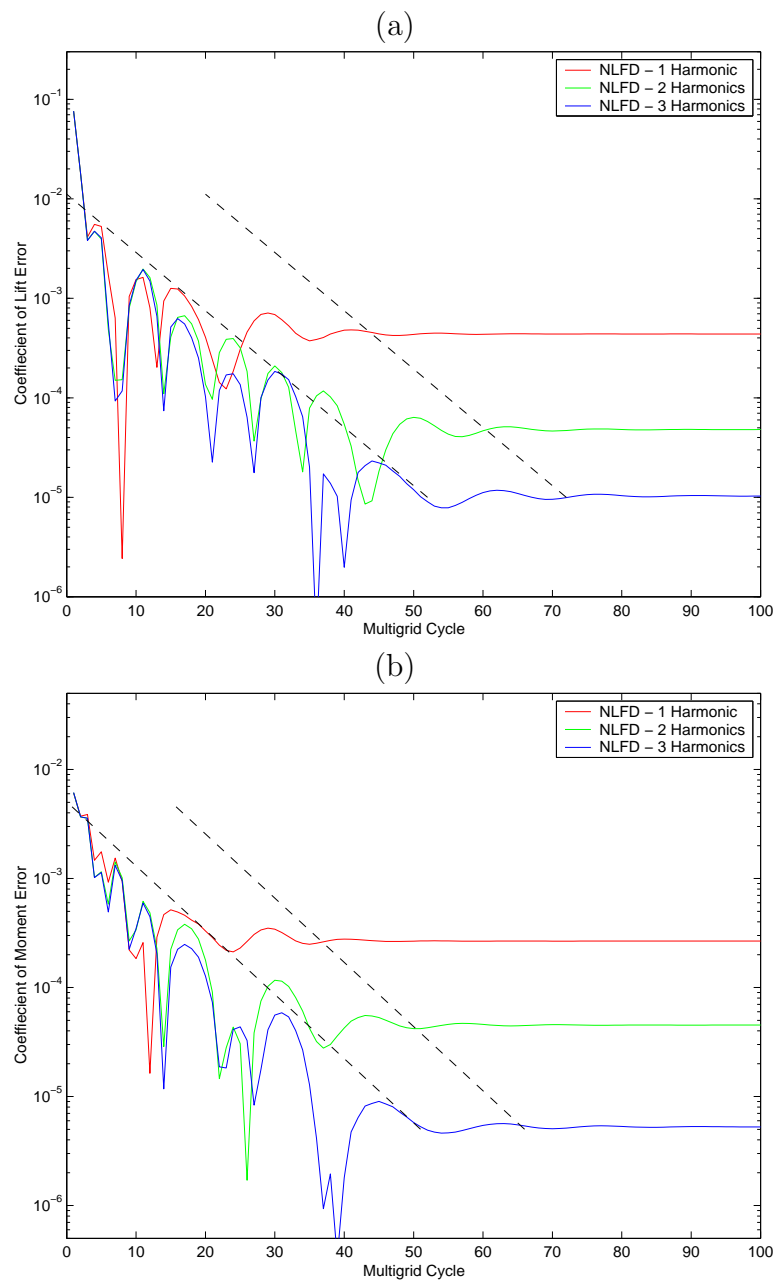


Figure 5.8: The number of multigrid cycles required by the NLFD code to reach equivalent error levels. (a) C_{l_1} . (b) C_{m_1}

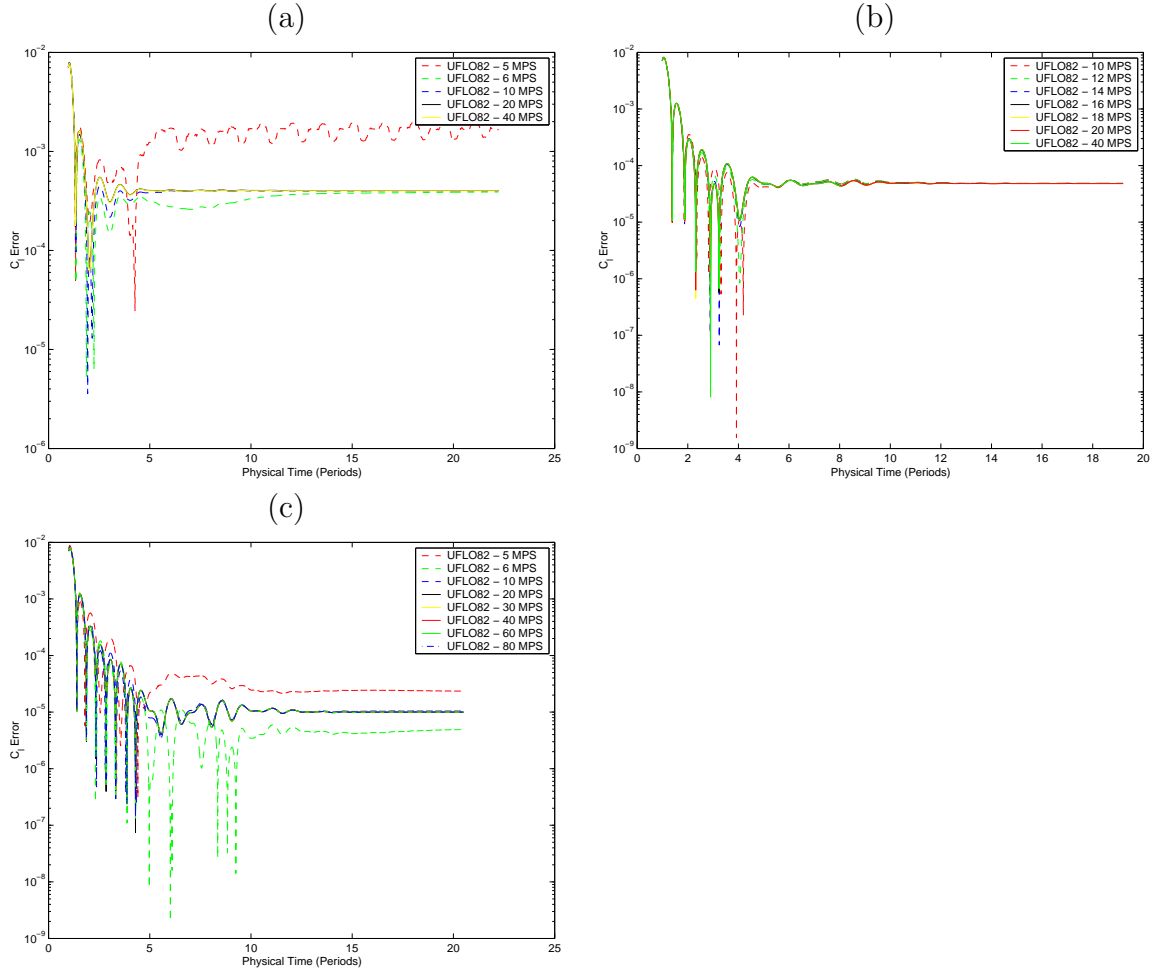


Figure 5.9: C_l error as a function of physical time in units of periods. (a) 45 Steps per wavelength (b) 125 Steps per wavelength (c) 244 Steps per wavelength.

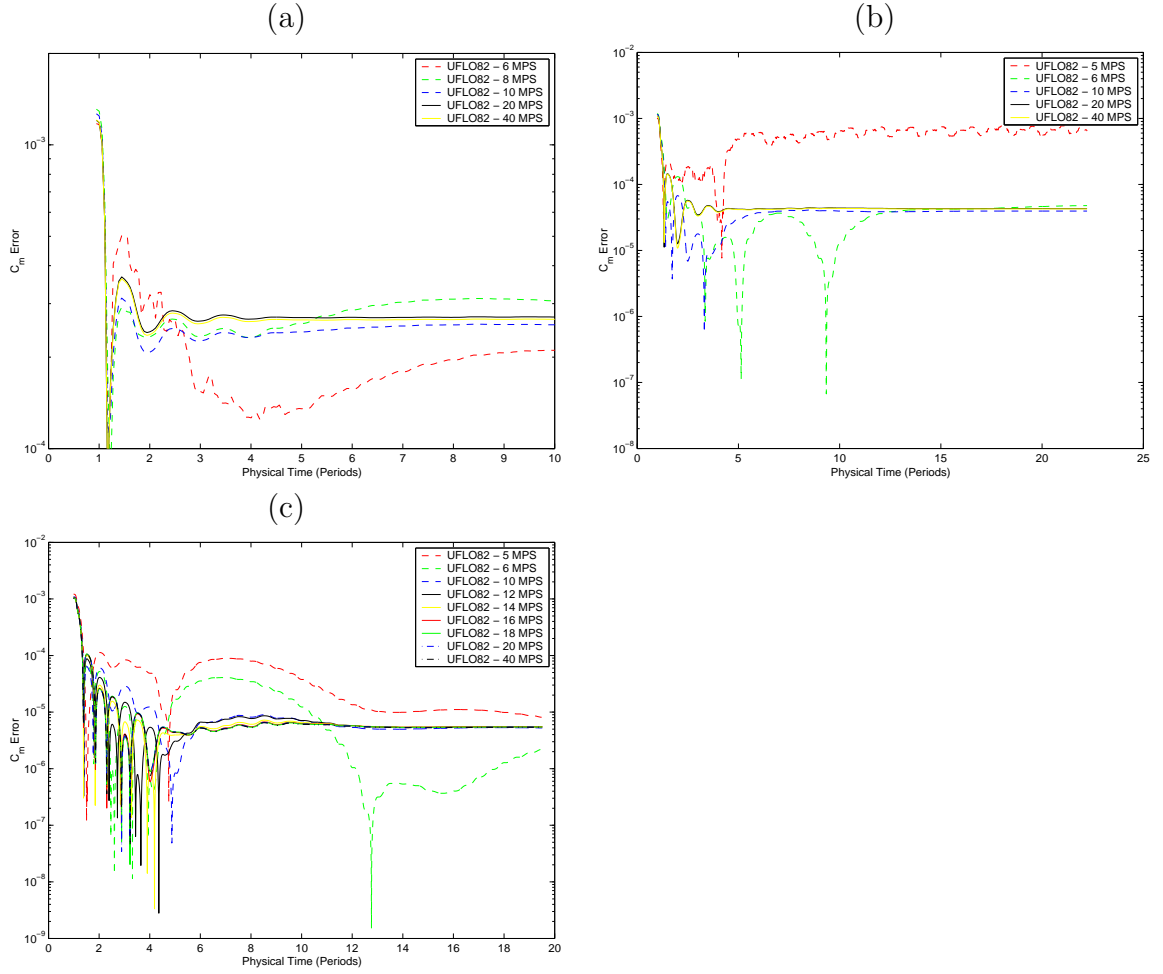


Figure 5.10: C_m error as a function of physical time in units of periods (a) 18 Steps per wavelength (b) 45 Steps per wavelength (c) 123 Steps per wavelength.

5.5 Sensitivity of the NLFD Solver To Flow Conditions

In addition to comparisons with experimental and other numerical data, variations in the efficiency of the NLFD solver with the flow field properties have also been evaluated. Steady parameters such as farfield Mach and Reynolds numbers will not be investigated. Instead the following sections look strictly at unsteady parameters such as pitching frequency and dynamic angle of attack.

5.5.1 Convergence Versus Reduced Frequency

According to equation 2.18, increased pitching frequency obviously moves the unsteady residual away from its steady analogue. A study was performed to determine the effect of this displacement on the residual convergence rate of the solver. Steady convergence rates represent the optimal efficiency, but unsteady conditions will provide additional scales of motion possibly impacting the convergence rate of the solver.

The test case is consistent with Davis's 64A010 experiment, except that the dynamic angle of attack has been increased from $\pm 1.01^\circ$ to $\pm 2.0^\circ$. Solutions were calculated on the 193x49 C-mesh proposed in section 4.3 using three time varying modes such that the highest temporal component is oscillating at three times the fundamental frequency. The calculations were performed at reduced frequencies between 0 and 0.5, holding constant any convergence acceleration parameters and artificial dissipation coefficients. The difference in residual convergence rates over this set of 16 calculations is thus strictly a function of the pitching frequency of the airfoil.

The first subplot within figure 5.11 shows the maximum over all the wavenumbers of the absolute value of the residual as a function of the multigrid cycle. The fastest converging solution was produced from a steady calculation (zero reduced frequency) at the mean angle of attack. The unsteady calculations exhibit similar convergence trends up to 100 multigrid cycles. At this point, the convergence rates begin to vary until the residual reaches machine zero. The second subplot within figure 5.11 shows the convergence rate per multigrid cycle based on the first 140 cycles of each

calculation. The plot shows a distinct difference between the steady and unsteady convergence rates with the latter being slightly less efficient. Another noticeable trend is also observed for the high frequency band, where larger frequencies reduce the efficiency of the solver by a maximum of 2.7 percent from the steady-state value.

5.5.2 Convergence Versus Dynamic Angle of Attack

Unlike the reduced frequency, the dynamic angle of attack affects terms calculated within the spatial operator, and consequently may have a stronger influence on the convergence rate. The effect of increasing the dynamic angle of attack is to increase the flow field nonlinearities and subsequent coupling between temporal modes.

A convergence study was performed with a fixed reduced frequency while the dynamic angle of attack was varied from $\pm 0^\circ$ to $\pm 2.8^\circ$. Again, convergence acceleration parameters and artificial dissipation coefficients are held constant to ensure convergence rates are strict functions of dynamic angle of attack

Subplot (a) within figure 5.12 shows the maximum over all the wavenumbers of the absolute value of the residual as a function of the multigrid cycle. The fastest converging solution was predictably produced by the steady solver (dynamic angle of attack at $\pm 0^\circ$). Overall, little variation is found over the range of dynamic angles of attack, with the poorest converging solution requiring only 10 additional multigrid cycles to reach machine zero. The statement is reinforced by subplot (b) within figure 5.12, which shows the convergence rate per multigrid cycle based on the first 140 cycles of each calculation. The overall variation in convergence rate per cycle is 1.6 percent; with a slight trend toward reducing the efficiency of the solver with increased dynamic angle of attack.

5.6 Summary

The computational efficiency of the NLFD solver has been compared to that of UFLO82, a state-of-the-art time accurate code. Unbiased comparisons were made

between the methods by measuring the amount of work required by each solver to reach equivalent error levels. Using error in coefficient of lift as the figure of merit the NLFD code is 8 to 19 times faster than UFLO82. Using coefficient of moment as the basis for comparison, the NLFD method is roughly three to nine times faster than the time-accurate approach. Numerical experiments confirm that the NLFD solver retains its efficiency over a wide range of unsteady flow conditions. This confirms its robustness as a tool for practical applications.

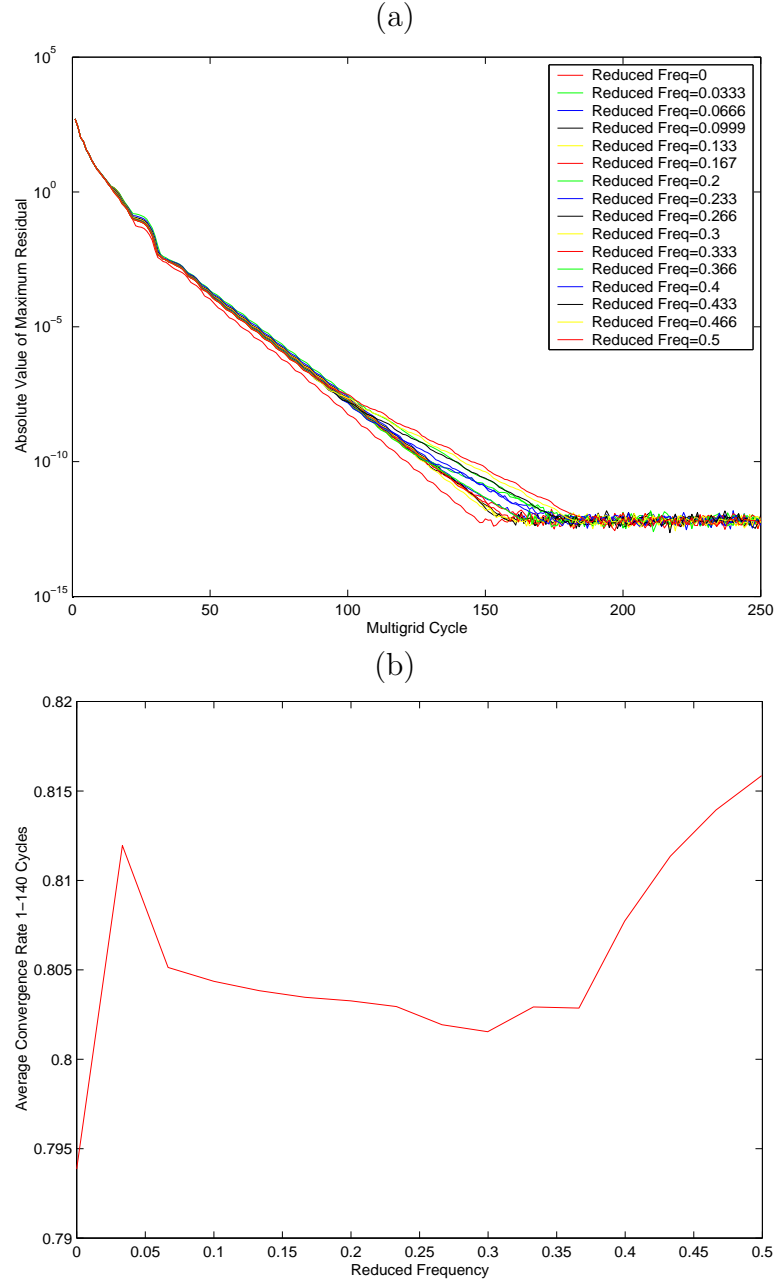


Figure 5.11: The effects of reduced frequency on the convergence of the NLFD unsteady residual. (a) The maximum over all the wavenumbers of the absolute value of the residual as a function of the multigrid cycle. (b) The residual convergence rate per multigrid cycle for the first 140 cycles.

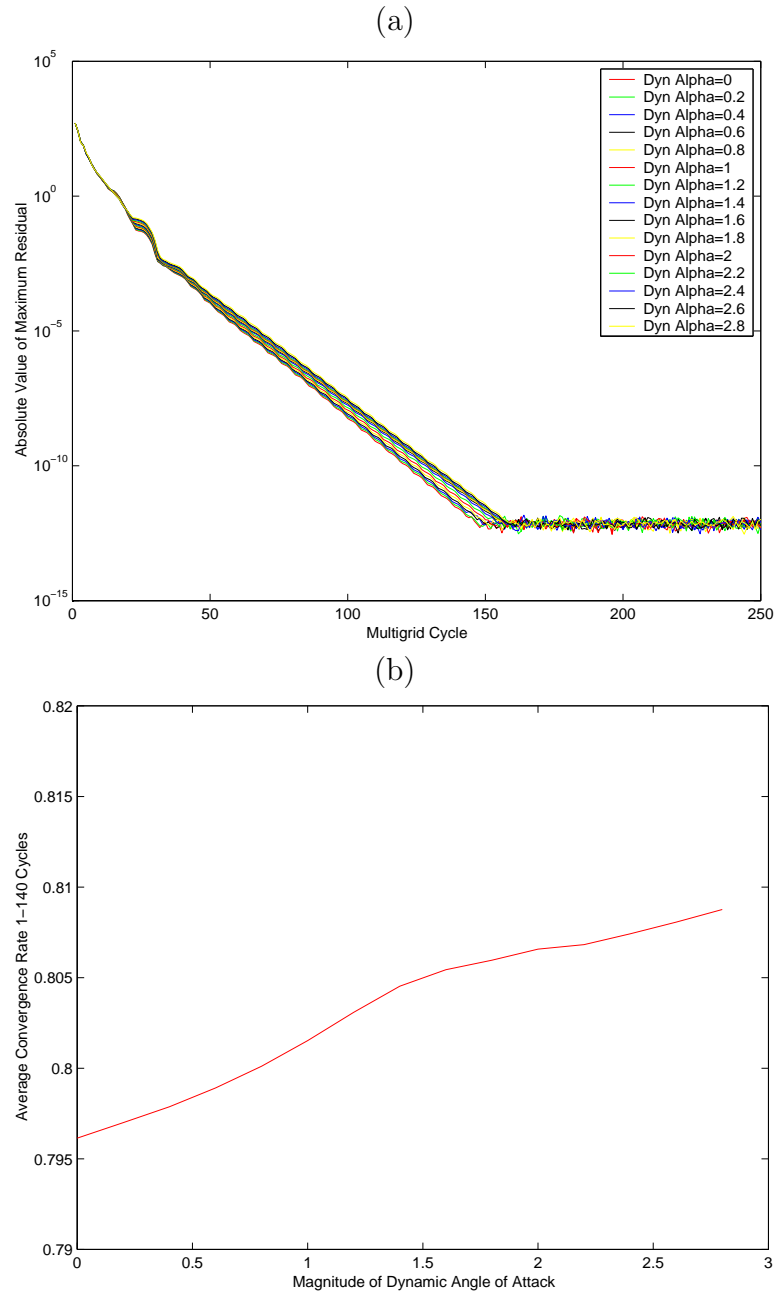


Figure 5.12: The effects of dynamic angle of attack on the convergence of the NLFD unsteady residual. (a) The maximum over all the wavenumbers of the absolute value of the residual as a function of the multigrid cycle. (b) The residual convergence rate per multigrid cycle for the first 140 cycles.

Chapter 6

Conclusions

The main conclusions which can be drawn from this work are presented below in categories corresponding to chapters 2 through 5.

6.1 Solver Implementation

Non-linear frequency domain methods can be integrated into robust solvers exhibiting convergence performance close to that of more established steady-state solvers like FLO82 and FLO103. Convergence acceleration technologies like multigrid need to be modified with coarse grid spectral viscosity to mitigate the dependence of convergence rates on unsteady effects. The point-implicit treatment commonly found in the inner iterations of the dual time stepping codes is not recommended for NLFD solvers, due to the introduction of instabilities associated with the extension to the complex domain. Given that the temporal derivative in pseudo-time is handled explicitly, residual averaging techniques need to be modified with CFL limiters to ensure stability for coarse grids in the multigrid cycle. In contrast to Harmonic Balance techniques, the NLFD method advances the unsteady residual in the frequency domain, allowing the use of a separate pseudo-time step for each wavenumber that includes the effect of the temporal derivative on the stability of the system.

For simple problems, numerical experiments confirm fast convergence to a periodic steady state, independent of the decay rate of the initial transients in physical time.

This uncouples the efficiency of the solver from the physical aspects of the flow field, and encourages the application of this method to more complex problems in turbomachinery, where resolving the transient decay dominates the cost of a conventional time-dependent calculation.

6.2 Cylindrical Vortex Shedding

The motivation behind performing NLFD calculations of vortex shedding behind a cylinder were threefold:

- To demonstrate that the NLFD method using a limited number of time varying modes could accurately resolve complex flow field physics such as the von Kármán vortex street.
- To demonstrate the ability of the GBVTP method to automatically determine the shedding frequency of the discrete equations.
- To show the impact of the GBVTP method on the solution in comparison to fixed-time-period results.

The temporal resolution survey shows that only three time varying modes are required to predict the Strouhal number and base suction coefficient to engineering accuracy for the entire range of laminar Reynolds numbers. However, the energy in the higher harmonics decreases with decreasing Reynolds number. At the lower end of the Reynolds number range, one or two modes can be used to achieve results that are as accurate as three mode calculations at the higher end of the range. Surprisingly, some of the NLFD results are closer to the experimental data than the Henderson results calculated using a very high order method.

Using an initial guess based on experimental data, the GBVTP method can be used to predict the time period of the fundamental harmonic. The iterative adjustment of this parameter results in the prediction of the Strouhal number for the discrete equations. This shedding frequency is unique for a given mesh and boundary

conditions, and is the only frequency that will allow the unsteady residual to decay to zero.

In general, the GBVTP method improved the agreement between the numerical and experimental results for base suction coefficient in comparison to methods where the time-period was fixed *a priori*. The improvement is largely a function of the difference between the shedding frequency value associated with the fixed-time-period calculation and the final value predicted by the GBVTP method. This improvement is more pronounced for coarser grids where the discretization error produces a substantial Reynolds number shift, exacerbating the difference in values of shedding frequency used by the different methods.

6.3 Pitching Airfoil

The motivations for the selection of the pitching airfoil problem are summarized below:

- To demonstrate the ability of the NLFD method to time accurately resolve flows of engineering importance like a transonic pitching airfoil in inviscid and turbulent viscous environments.
- To compare the efficiency of the NLFD solver in comparison to time-accurate codes like UFLO82 which implements a dual time stepping technique.

The numerical experiments confirm the accuracy and efficiency of the NLFD method for this problem. It is remarkable that using just one time varying mode, the NLFD predictions for coefficient of lift provide excellent agreement with experimental results. This applies even to flow fields where mild separation effects are predicted.

In the case of the coefficient of moment, NLFD predictions provide poor agreement with experimental studies. This is no different than time-accurate solvers with similar spatial discretizations. Consequently it may be concluded that the discretization of the temporal derivative is not the source of the discrepancy, which may partly be due to turbulence-modeling errors. Moreover, the experimental results are questionable,

given the magnitude of the corrections applied by the experimentalists during post-test analysis.

6.4 Computational Efficiency of the NLFD Method

Given equivalent spatial discretizations and overall error, the NLFD method is significantly more efficient than dual time stepping codes like UFLO82. Specifically, using coefficient of lift as a figure of merit the NLFD solver is 8 to 19 times faster than the dual time stepping code. Using coefficient of moment as the basis for comparison, the NLFD method is roughly three to nine times faster than the time-accurate approach. The ratio of efficiency between the codes is provided as a range because the NLFD method gains efficiency at higher temporal resolutions.

The NLFD method should in principle provide the very rapid convergence of a spectral method with increasing numbers of modes, and in this sense it is an optimal discretization scheme for time-periodic problems. In practice it turns out that it can also be effectively used as a reduced order method in which users deliberately choose not to resolve the higher temporal modes in the solution. It is particularly efficient in cases where the vast majority of the energy within the solution is contained in the fundamental harmonic. More challenging cases include solutions where energy is distributed among several modes. Consequently, validation of this method should be done on a case by case basis with a focus on quantifying not only the solution error, but also error in the data values useful to engineers, which must be extracted by post-processing the solution.

6.5 Future Work

In this work the NLFD method has been applied only to the solution of the Euler and Navier-Stokes equations. The same efficiency benefits might be realized for other systems of equations. Specifically, it could be used to solve the adjoint system to the Euler and Navier-Stokes equations for shape optimization problems in unsteady flows [55]. This could benefit the design of both turbomachinery and helicopter blades.

In addition, solutions of the adjoint equations could be used to obtain error estimates for functionals such as lift and drag coefficients. This might eliminate the need for refinement studies where solutions with additional modes are calculated to estimate the functional convergence.

The stationary system of equations produced by the NLFD method has been solved using established techniques to accelerate convergence. Specifically these consist of a modified Runge-Kutta scheme, used with pseudo local time stepping, residual averaging and multigrid. Recently, other solvers using Gauss-Seidel relaxation in a multigrid procedure have demonstrated significantly faster convergence rates. The application of these new convergence acceleration techniques to the NLFD method offers the opportunity for a further major improvement in efficiency.

Appendix A

Discretization

A.1 Non-Dimensionalization

The analysis starts by presenting a form of the conservation equations where the flux vector, \vec{F} , has been split into convective, \vec{F}_c , and dissipative, \vec{F}_v , components:

$$\frac{d}{dt} \int_{\Omega} W dV + \oint_{\partial\Omega} \vec{F}_c \cdot \vec{N} ds + \oint_{\partial\Omega} \vec{F}_v \cdot \vec{N} ds = 0. \quad (\text{A.1})$$

Following the work of Martinelli [45], non-dimensionalization of the equation set can be achieved by dividing the dimensional quantities by characteristic properties, denoted by the ∞ symbol, to achieve the dimensionless quantities denoted by the $*$ symbol:

$$\begin{aligned} x_i^* &= \frac{x_i}{x_{\infty}} \\ t^* &= \frac{t}{x_{\infty} \left(\frac{\rho_{\infty}}{p_{\infty}} \right)^{\frac{1}{2}}} \\ \rho^* &= \frac{\rho}{\rho_{\infty}} \\ u_i^* &= \frac{u_i}{\left(\frac{p_{\infty}}{\rho_{\infty}} \right)^{\frac{1}{2}}} \end{aligned}$$

$$\begin{aligned}
b_i^* &= \frac{b_i}{\left(\frac{p_\infty}{\rho_\infty}\right)^{\frac{1}{2}}} \\
E^* &= \frac{E}{\left(\frac{p_\infty}{\rho_\infty}\right)} \\
\mu^* &= \frac{\mu}{\mu_\infty}.
\end{aligned} \tag{A.2}$$

The volumetric and flux integrals of the conservation equations can be rewritten in dimensionless quantities:

$$\frac{d}{dt} \int_{\Omega} W dV = x_\infty^2 \begin{bmatrix} (p_\infty \rho_\infty)^{\frac{1}{2}} \\ p_\infty \\ p_\infty \\ \left(\frac{p_\infty^3}{\rho_\infty}\right)^{\frac{1}{2}} \end{bmatrix} \frac{d}{dt^*} \int_{\Omega} W^* dV^* \tag{A.3}$$

$$\oint_{\partial\Omega} \vec{F}_c \cdot \vec{N} ds = x_\infty^2 \begin{bmatrix} (p_\infty \rho_\infty)^{\frac{1}{2}} \\ p_\infty \\ p_\infty \\ \left(\frac{p_\infty^3}{\rho_\infty}\right)^{\frac{1}{2}} \end{bmatrix} \oint_{\partial\Omega} \vec{F}_c^* \cdot \vec{N} ds^* \tag{A.4}$$

$$\oint_{\partial\Omega} \vec{F}_v \cdot \vec{N} ds = x_\infty^2 \begin{bmatrix} 0 \\ \frac{\mu_\infty}{x_\infty} \left(\frac{p_\infty}{\rho_\infty}\right)^{\frac{1}{2}} \\ \frac{\mu_\infty}{x_\infty} \left(\frac{p_\infty}{\rho_\infty}\right)^{\frac{1}{2}} \\ \frac{\mu_\infty}{x_\infty} \left(\frac{p_\infty}{\rho_\infty}\right) \end{bmatrix} \oint_{\partial\Omega} \vec{F}_v^* \cdot \vec{N} ds^*. \tag{A.5}$$

The conservation equations then simplify into the following dimensionless form:

$$\frac{d}{dt^*} \int_{\Omega} W^* dV^* + \oint_{\partial\Omega} \vec{F}_c^* \cdot \vec{N} ds^* = \sqrt{\gamma} \frac{M_\infty}{Re_\infty} \oint_{\partial\Omega} \vec{F}_v^* \cdot \vec{N} ds^*. \tag{A.6}$$

Ratios of common dimensionless coefficients can be used to scale the viscous stress terms. Here M_∞ and Re_∞ are the Mach number and Reynolds number respectively. For the sake of brevity, the remainder of this paper will drop the $()^*$ notation.

A.2 Spatial Operators

A.2.1 Convective Fluxes

In applying this formula to the convective flux, the following notation is introduced:

$$\begin{aligned}\vec{S} &= |S| (N_1 \mathbf{e}_1 + N_2 \mathbf{e}_2) \\ \vec{F}_c &= F_{c1} \mathbf{e}_1 + F_{c2} \mathbf{e}_2.\end{aligned}\tag{A.7}$$

The outward facing normal can be calculated as the perpendicular vector to the cell face whose length is equal to the length of the face.

Figure A.1 illustrates a typical grid (shown by the solid lines), and the derived control volume (shown as dashed lines). The values of the state vector are held at the center of the cell. The exact flux vectors are approximated by taking the average of the flux based on the cell values. The overall convective flux evaluated at the $i + \frac{1}{2}$ face can be written as:

$$\vec{F}_{c_{i+\frac{1}{2}}} \cdot \vec{S}_{i+\frac{1}{2}} = \frac{1}{2} \left[\vec{F}_c(W_{i+1}) + \vec{F}_c(W_i) \right] \cdot \vec{S}_{i+\frac{1}{2}}.\tag{A.8}$$

On a Cartesian grid this results in a pure central difference scheme that provides no damping to the odd-even modes. Hence upwind biasing by artificial dissipation or directly constructed upwind schemes are employed.

A.2.2 Artificial-Dissipation Fluxes

A variety of artificial-dissipation schemes were used in this work. The cylinder and O-mesh pitching airfoil cases were calculated with a Symmetric Limited Positive (SLIP) [30, 80, 65, 79] scheme based on differences in the characteristic form of the state vector. The C-mesh pitching airfoil cases were calculated with a Jameson-Schmidt-Turkel (JST) [39] scheme that will be the focus of this section. A thorough discussion of the implementation and mathematical properties (positivity, etc.) of both these schemes is provided by Jameson [35, 36]. Using the pitching airfoil results shown in chapter 4 as a reference, little difference is exhibited between solutions using




Figure A.1: Convective Flux Control Volume

different schemes.

For the purposes of this section, the Δ symbol denotes the difference in the solution between two adjacent cells

$$\Delta W_{i+\frac{1}{2}} = W_{i+1} - W_i. \quad (\text{A.9})$$

The JST scheme uses a blended mix of first ($\Delta W_{i+\frac{1}{2}}$) and third order ($\Delta W_{i+\frac{3}{2}} - 2\Delta W_{i+\frac{1}{2}} + \Delta W_{i-\frac{1}{2}}$) dissipation operators evaluated at each cell face to stabilize the solution

$$\begin{aligned} F_{d_{i+\frac{1}{2}}} &= \varepsilon_{i+\frac{1}{2}}^{(2)} \Delta W_{i+\frac{1}{2}} \\ &- \varepsilon_{i+\frac{1}{2}}^{(4)} \left(\Delta W_{i+\frac{3}{2}} - 2\Delta W_{i+\frac{1}{2}} \Delta W_{i-\frac{1}{2}} \right). \end{aligned} \quad (\text{A.10})$$

The coefficients that control which dissipation operator is dominant, $\varepsilon_{i+\frac{1}{2}}^{(2)}$ and $\varepsilon_{i+\frac{1}{2}}^{(4)}$,

are given by:

$$\begin{aligned}\varepsilon_{i+\frac{1}{2}}^{(2)} &= c_2 a_{i+\frac{1}{2}} Q_{i+\frac{1}{2}} \\ \varepsilon_{i+\frac{1}{2}}^{(4)} &= \left(c_4 - c_2 Q_{i+\frac{1}{2}} \right) a_{i+\frac{1}{2}}.\end{aligned}\tag{A.11}$$

The spectral radius of the convective fluxes at the cell face $(a_{i+\frac{1}{2}})$ properly scales the coefficients to ensure stability

$$a_{i+\frac{1}{2}} = \vec{U}_{i+\frac{1}{2}} \cdot \vec{S}_{i+\frac{1}{2}} + c \|\vec{S}_{i+\frac{1}{2}}\|.\tag{A.12}$$

This mixing process is controlled by a switch $(Q_{i+\frac{1}{2}})$ based on differences in the pressure field

$$\begin{aligned}Q_i &= \left| \frac{p_{i+1} - 2p_i + p_{i-1}}{p_{i+1} + 2p_i + p_{i-1} + \epsilon} \right| \\ Q_{i+\frac{1}{2}} &= \min \left(\frac{1}{4}, \max(Q_{i+2}, Q_{i+1}, Q_i, Q_{i-1}) \right).\end{aligned}\tag{A.13}$$

At discontinuities such as a shock wave, the pressure differences between adjacent cells will be large, causing $Q_{i+\frac{1}{2}}$ to saturate at $\frac{1}{4}$. Consequently the low order term dominates the dissipation scheme essentially providing first order upwinding. Away from the shock where the solution is smooth the switch tends toward zero. In this case a Taylor series analysis is valid (under the assumption of a smooth pressure field), and shows that the switch decays like Δx^2 . This decay ensures that the third order term is the dominant contribution to the dissipation scheme in smooth areas of the solution.

A.2.3 Viscous Fluxes

This section describes the spatial discretization used to approximate the viscous stresses associated with the Navier-Stokes equations. In his dissertation Martinelli described two separate discretizations of viscous terms based on different control volumes named CCS-A and CCS-B [45]. This section outlines the mathematics behind

the CCS-A approach, but both schemes were implemented in the NLFD code with negligible difference.

The viscous fluxes require us to estimate the values of σ_{ij} and q_i at the cell faces. These terms are functions of derivatives of temperature and velocity. To estimate these derivatives, a second control volume is introduced and shown as the dashed lines in figure A.2. Through the application of the divergence theorem the derivatives




Figure A.2: Viscous Flux Control Volume

of any flow field quantity at point A can be approximated as:

$$\int_{\Omega} \frac{\partial u}{\partial x_i} dV = \int \nabla \cdot u \mathbf{e}_i dV = \oint_{\partial\Omega} u \mathbf{e}_i \cdot \vec{N} ds. \quad (\text{A.14})$$

The same technique can be shifted to point B, and an estimate can then be made of the derivative along cell face AB. The variable u is used in this section as a sample variable and does not necessarily represent velocity.

On a discrete mesh the continuous integrals can be replaced by their discrete approximations. The product of the volume and the average value of the derivative

at the cell center approximates the volumetric integral. As in previous cases, the surface integral is replaced by the sum of fluxes over each cell face

$$V_{cv} D_{x_i} u = \sum_{cv} u S_i. \quad (\text{A.15})$$

The derivatives of all the quantities are computed consistently. To illustrate the technique, D_x evaluated at point A (or $X_{i,j}$ on the mesh) can be stated as:

$$\begin{aligned} V_{cv} D_x u_{i,j} &= \frac{1}{2} (u_a + u_b) (y_a - y_b) + \frac{1}{2} (u_b + u_c) (y_b - y_c) \\ &+ \frac{1}{2} (u_c + u_d) (y_c - y_d) + \frac{1}{2} (u_d + u_a) (y_d - y_a) \\ &= \frac{1}{2} [(u_a - u_c)(y_d - y_b) + (u_d - u_b)(y_c - y_a)]. \end{aligned} \quad (\text{A.16})$$

The values of the solution at the vertices of the control volume can be approximated by averaging the nearest solution values held at the cell centers:

$$\begin{aligned} u_a &= \frac{1}{2} (u_{i,j+1} + u_{i+1,j+1}) \\ u_b &= \frac{1}{2} (u_{i,j+1} + u_{i,j}) \\ u_c &= \frac{1}{2} (u_{i+1,j} + u_{i,j}) \\ u_d &= \frac{1}{2} (u_{i+1,j} + u_{i+1,j+1}). \end{aligned} \quad (\text{A.17})$$

The position of the vertices defining the control volume are calculated by averaging the end points of the mesh line containing that vertex:

$$\begin{aligned} y_a &= \frac{1}{2} (y_{i,j+1} + y_{i,j}) \\ y_b &= \frac{1}{2} (y_{i-1,j} + y_{i,j}) \\ y_c &= \frac{1}{2} (y_{i,j-1} + y_{i,j}) \\ y_d &= \frac{1}{2} (y_{i+1,j} + y_{i,j}). \end{aligned} \quad (\text{A.18})$$

The averaged derivative notation is introduced to rewrite equation A.16 into a form mirroring the strong conservation form of Viviand [67] and Vinokur [66] (see equation A.33 for the conservative form of the derivative operators):

$$\begin{aligned} u_d - u_b &= \frac{1}{2} (u_{i+1,j+1} - u_{i,j+1} + u_{i+1,j} - u_{i,j}) \\ &= \frac{1}{2} \left(D_i u_{i+\frac{1}{2},j+1} + D_i u_{i+\frac{1}{2},j} \right) = \overline{D_i} u \end{aligned} \quad (\text{A.19})$$

$$\begin{aligned} u_a - u_c &= \frac{1}{2} (u_{i+1,j+1} - u_{i+1,j} + u_{i,j+1} - u_{i,j}) \\ &= \frac{1}{2} \left(D_j u_{i+1,j+\frac{1}{2}} + D_j u_{i,j+\frac{1}{2}} \right) = \overline{D_j} u \end{aligned} \quad (\text{A.20})$$

$$y_a - y_c = \frac{1}{2} (y_{i,j+1} - y_{i,j-1}) = \overline{D_j} y \quad (\text{A.21})$$

$$y_d - y_b = \frac{1}{2} (y_{i+1,j} - y_{i-1,j}) = \overline{D_i} y. \quad (\text{A.22})$$

Substituting equations A.19- A.22 into equation A.16 results in an expression for the spatial derivative in terms of solution differences defined in the computational plane and metric terms:

$$D_x u_A = \frac{(\overline{D_i} u \overline{D_j} y - \overline{D_j} u \overline{D_i} y)}{2V_{cv}}. \quad (\text{A.23})$$

The stress tensor terms σ_{ij} for the AB face can now be approximated by averaging the derivative terms evaluated at points A and B as:

$$D_x u_{AB} = \frac{1}{2} (D_x u_A + D_x u_B). \quad (\text{A.24})$$

The viscous fluxes can then be evaluated at each face and summed over the control volume defined in figure A.1.

A.2.4 Turbulence Model

Several different types of simulations have been used in this research to validate the numerical methods. The first model presented is a low Reynolds laminar simulation of unsteady vortex shedding behind a cylinder. Due to the laminar flow characteristics

no turbulence model was needed for this case. A pitching airfoil with a turbulent boundary layer was used as the second case. A Baldwin-Lomax turbulence model was implemented to model the turbulent motions in the boundary layer which would have gone uncaptured by the relatively coarse mesh. Due to the absence of a boundary layer, the Euler calculations presented for the pitching airfoil obviously did not use any turbulence models.

The Baldwin-Lomax turbulence model is based on a set of algebraic equations that estimate eddy viscosity $\tilde{\mu}$ for a steady flow field. The model was implemented for unsteady flow by calculating $\tilde{\mu}$ at a given instance in time based on the solution and the grid at that same instant. This snapshot approach effectively ignores the unsteady effects on the turbulence characteristics.

A.3 Boundary Conditions

A.3.1 Wall

The NLFD solver has been developed using ghost cells to represent the wall boundary conditions. Lower order dissipation schemes typically only need a pressure extrapolation for the ghost cell values since the mass flux is zero through the wall. However, higher order dissipation schemes with extended stencils have been used throughout this research. These extended stencils use the entire state vector at the ghost cell when calculating the artificial dissipation flux at the top face of the cell adjacent to the wall. Hence the boundary condition needs to not only extrapolate pressure to the ghost cell but all other state variables. Figure A.3 provides a cartoon of a curvilinear mesh at the wall. The subscript notation used throughout this section is based on the labels placed in this figure: g denotes the ghost cell, i or i, j denotes the first interior cell.

The density variable for both Navier-Stokes and Euler simulations is simply continued across the wall

$$\rho_g = \rho_i. \quad (\text{A.25})$$

The velocity components are calculated such that the average of the ghost and interior




Figure A.3: Pictorial of cells at the wall for a curvilinear mesh.

velocity vectors provide a zero mass flux at the wall. For the Euler equations this implies that the averaged vector is perpendicular to the wall vector, and the tangential components of the velocity are equal for both the ghost and interior cells. If \vec{N} is the wall normal then an equation can be written for the ghost cell velocity vector as:

$$\vec{U}_g = \vec{U}_i - 2 \left(\vec{U}_i \cdot \vec{N} \right) \vec{N}. \quad (\text{A.26})$$

For the Navier-Stokes case, the no-slip boundary condition implies that the averaged vector is zero at the wall

$$\vec{U}_g = -\vec{U}_i. \quad (\text{A.27})$$

A pressure extrapolation for these flow simulations can be presented as:

$$p_g = p_i + \alpha (p_i - p_{i+1}). \quad (\text{A.28})$$

The α variable acts as a switch between Navier-Stokes and Euler computations. In the viscous case, α equal to zero provides a simple pressure continuation across the wall. In the inviscid case, α equal to one ensures that the ghost cell pressure is equal

to a linear extrapolation of the pressure using the next two interior points.

Rizzi [59] proposed another boundary condition for pressure applicable to the Euler equations. Following the notes of Alonso [3], the remainder of this section will document the implementation of this boundary condition.

As a particle moves along a streamline the volumetric flux perpendicular to that streamline is always zero; or in mathematical terms the substantial derivative of the volumetric flux is zero

$$\left(\frac{\partial}{\partial t} + \vec{U} \cdot \nabla \right) (\vec{U} \cdot \vec{S}) = 0. \quad (\text{A.29})$$

For the remainder of this subsection the use of vector notation will be dropped; this section opts to individually name the velocity components u and v and the spatial dimensions x and y . Using the definition of the dot product and applying the chain rule to separate any derivatives of products will result in the following expression:

$$\begin{aligned} S_x \frac{\partial u}{\partial t} + u \frac{\partial S_x}{\partial t} + S_y \frac{\partial v}{\partial t} + v \frac{\partial S_y}{\partial t} + u S_x \frac{\partial u}{\partial x} + u^2 \frac{\partial S_x}{\partial x} + \\ u S_y \frac{\partial v}{\partial x} + uv \frac{\partial S_y}{\partial x} + v S_x \frac{\partial u}{\partial y} + uv \frac{\partial S_x}{\partial y} + v S_y \frac{\partial v}{\partial y} + v^2 \frac{\partial S_y}{\partial y} = 0. \end{aligned} \quad (\text{A.30})$$

In non-conservative form the x - and y -components of the momentum equations can be written as follows:

$$\begin{aligned} \frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + v \frac{\partial u}{\partial y} + \frac{1}{\rho} \frac{\partial p}{\partial x} &= 0 \\ \frac{\partial v}{\partial t} + u \frac{\partial v}{\partial x} + v \frac{\partial v}{\partial y} + \frac{1}{\rho} \frac{\partial p}{\partial y} &= 0. \end{aligned} \quad (\text{A.31})$$

The x -component of the momentum equation is now multiplied by S_x while the y -component of the momentum equation is multiplied by S_y . Subtracting the resulting expressions from equation A.30 will result in the following equation:

$$\begin{aligned} \frac{S_x}{\rho} \frac{\partial p}{\partial x} + \frac{S_y}{\rho} \frac{\partial p}{\partial y} = \\ \left(u \frac{\partial S_x}{\partial t} + v \frac{\partial S_y}{\partial t} \right) + \left(u^2 \frac{\partial S_x}{\partial x} + uv \frac{\partial S_y}{\partial x} + uv \frac{\partial S_x}{\partial y} + v^2 \frac{\partial S_y}{\partial y} \right). \end{aligned} \quad (\text{A.32})$$

Using the transforms between a curvilinear and uniform mesh presented by Vi-
viand [67] and Vinokur [66] the directional derivatives can be replaced using a com-
bination of metrics and derivatives evaluated on a uniform mesh

$$\begin{aligned}\frac{\partial}{\partial x} &= \frac{1}{J} \left(\frac{\partial y}{\partial \eta} \frac{\partial}{\partial \xi} - \frac{\partial y}{\partial \xi} \frac{\partial}{\partial \eta} \right) \\ \frac{\partial}{\partial y} &= \frac{1}{J} \left(\frac{\partial x}{\partial \xi} \frac{\partial}{\partial \eta} - \frac{\partial x}{\partial \eta} \frac{\partial}{\partial \xi} \right).\end{aligned}\tag{A.33}$$

The metric terms should be evaluated at the face of the wall, hence the outward wall
normal is equivalent to derivatives in the ξ direction

$$\begin{aligned}\frac{\partial x}{\partial \xi} &= S_y \\ \frac{\partial y}{\partial \xi} &= -S_x.\end{aligned}\tag{A.34}$$

Derivative terms in the η direction are approximated by averaging the cell face normals
in the tangential direction (see figure A.3)

$$\begin{aligned}\frac{\partial y}{\partial \eta} &= \bar{R}_x = \frac{1}{2} (R_{x_{i-1}} + R_{x_i}) \\ \frac{\partial x}{\partial \eta} &= -\bar{R}_y = \frac{-1}{2} (R_{y_{i-1}} + R_{y_i})\end{aligned}\tag{A.35}$$

The cell volume J can be calculated as a product of these wall normals

$$J = S_y \bar{R}_x - S_x \bar{R}_y.\tag{A.36}$$

Using these definitions the Rizzi boundary condition can be written using the face
normals of the cell adjacent to the wall

$$\begin{aligned}(S_x^2 + S_y^2) \frac{\partial p}{\partial \eta} &= J \rho \left(u \frac{\partial S_x}{\partial t} + v \frac{\partial S_y}{\partial t} \right) + \\ &\rho (u \bar{R}_x + v \bar{R}_y) \left(u \frac{\partial S_x}{\partial \xi} + v \frac{\partial S_y}{\partial \xi} \right) - (S_x \bar{R}_x + S_y \bar{R}_y) \frac{\partial p}{\partial \xi}.\end{aligned}\tag{A.37}$$

Due to the magnitude of the terms the derivatives of the grid metrics with respect to time $\left(\frac{\partial S_x}{\partial t}, \frac{\partial S_y}{\partial t}\right)$ can be ignored.

A.3.2 Farfield

Two different farfield boundary conditions were used by the NLFD codes. The cylinder simulations used a boundary condition based on one-dimensional characteristic theory applied to a linearized system of Euler equations. The approach approximates a non-reflecting boundary condition by allowing outgoing characteristics to continue information to the halo cells while nullifying incoming characteristics. The method is derived using a solution represented in the frequency domain that takes advantage of the representation of the state vector by the NLFD code. A full derivation of the boundary condition is provided by McMullen *et al.* [48].

In the case of the pitching airfoil, one-dimensional Riemann invariants are applied to each instance of the solution in its time history. The boundary conditions for the cylinder are based on a form of the governing equations linearized using the assumption that the unsteady terms are small in comparison to the time-averaged values. For the pitching airfoil, the validity of this assumption is a function of the magnitude of the mesh velocity terms in equation 2.4. The magnitude of these terms will increase as the boundary distance or the reduced frequency increases. A point exists where the growth of these terms would be large enough to invalidate the linearization of the equations and the resulting boundary conditions.

In addition, applying Riemann invariants facilitates comparisons with other codes. Chapter 5 of this work compares NLFD solutions to those produced by the UFLO and FLO family of solvers. All of these codes use Riemann invariants in their farfield boundary conditions. Thus inconsistency in the boundary conditions is eliminated as a possible source of error.

Results from the pitching airfoil simulations presented in chapters 4 and 5 use farfield boundaries that are 20 chords from the wall. The magnitude of an unsteady wave naturally diminishes as it travels in multi-dimensional space. The effect is

heightened by numerical dissipation and a coarse farfield grid that artificially dampens the waves as they travel outward. The very small magnitude of the unsteady disturbances in the farfield allow for the application of a steady boundary condition, like Riemann invariants, without significant corruption of the solution in the interior.

Using the homentropic assumption, conservation of mass and momentum equations for one-dimensional flow can be simplified into a system of material derivatives

$$\begin{aligned}\frac{D^+}{Dt}(u + f) &= 0 \\ \frac{D^-}{Dt}(u - f) &= 0.\end{aligned}\tag{A.38}$$

The material derivative operators and the f function can be simply defined as follows:

$$\begin{aligned}\frac{D^+}{Dt} &= \frac{\partial}{\partial t} + (u + c)\frac{\partial}{\partial x} \\ \frac{D^-}{Dt} &= \frac{\partial}{\partial t} + (u - c)\frac{\partial}{\partial x} \\ f &= \frac{2}{\gamma - 1}c.\end{aligned}\tag{A.39}$$

Figure A.4 provides a diagram of the cells and the notation used in the implementation of the boundary condition in a two-dimensional environment. Farfield and ghost cell quantities will be denoted with ∞ and g subscript respectively. The coordinate frame for the Riemann invariants will be along the outward pointing normal to the cell face. All velocities will be projected in this direction using the dot product operator

$$q = \vec{U} \cdot \vec{N}.\tag{A.40}$$

The right running characteristics associated with D^+ are a function of the interior properties while the left running characteristics associated with D^- are a function of farfield values. The normal component of the velocity and speed of sound at the ghost cell are simple linear combinations of the invariants

$$J^+ = q_{i,j} + f_{i,j}$$




Figure A.4: Pictorial of cells in the farfield for a curvilinear mesh.

$$\begin{aligned}
 J^- &= q_\infty - f_\infty \\
 q_g &= \frac{1}{2}(J^+ + J^-) \\
 c_g &= \frac{\gamma - 1}{4}(J^+ - J^-).
 \end{aligned} \tag{A.41}$$

Consistent with the direction of the entropic and vortical characteristics, the entropy (s) and tangential component of velocity (r) are continued along the particle path. Hence for outgoing flows where $q_g > 0$ the ghost cell values are equivalent to the interior values

$$\begin{aligned}
 r_g &= r_{i,j} \\
 s_g &= s_{i,j}.
 \end{aligned} \tag{A.42}$$

For incoming flows where $q_g < 0$, the farfield values are simply continued to the ghost

cell

$$\begin{aligned} r_g &= r_\infty \\ s_g &= s_\infty. \end{aligned} \tag{A.43}$$

The velocity at the ghost cell in the X,Y frame can be constructed from the velocity components in the normal frame, and all the thermodynamic properties can be calculated using the values of entropy and speed of sound.

Appendix B

Unsteady Channel Flow

B.1 Small Perturbation Theory

The first model used in the verification of this method was the simulation of incompressible inviscid flow through a one-dimensional channel. A pictorial representation of this problem is provided in figure B.1. This model was selected because an analytic solution to the flow field was derived using small perturbation theory by Merkle and Athavale [51]. This derivation assumes that the magnitude of the unsteadiness is small in comparison with the mean flow values, allowing the problem to be approximated by linear theory where all the wavenumbers of the problem uncouple. As such, the model does not test any of the non-linear capabilities of the NLFD scheme and is provided as basic verification of the overall numerical scheme and boundary conditions.

A derivation of the analytic solution begins with the incompressible form of the mass and momentum equations in one dimension:

$$\begin{aligned}\frac{\partial u}{\partial x} &= 0 \\ \frac{\partial u}{\partial t} + \frac{1}{\rho} \frac{\partial p}{\partial x} &= 0.\end{aligned}\tag{B.1}$$

Differentiating the mass equation with respect to time, the momentum equation with




Figure B.1: Pictorial of one-dimensional channel flow model problem.

respect to space, and subtracting the resulting equations provides the following:

$$\frac{\partial^2 p}{\partial x^2} = 0. \quad (\text{B.2})$$

The solution can be decomposed into two parts; the time-average of the solution over a given period and an unsteady component which is assumed to be significantly smaller than the average

$$\begin{aligned} p &= \bar{p} + \acute{p} \\ u &= \bar{u} + \acute{u}. \end{aligned} \quad (\text{B.3})$$

Given an infinite period of time, the following integral relations are assumed to hold for both pressure and velocity:

$$\begin{aligned} \frac{1}{T} \int_0^{\infty} p \, dt &= \bar{p} \\ \frac{1}{T} \int_0^{\infty} \acute{p} \, dt &= 0. \end{aligned} \quad (\text{B.4})$$

By substituting the above definitions into the equations B.1 and B.2 and integrating

over a time period one can take advantage of equation B.4 and write the following:

$$\begin{aligned}\frac{\partial^2 \acute{p}}{\partial x^2} &= 0 \\ \frac{\partial \acute{u}}{\partial x} &= 0.\end{aligned}\tag{B.5}$$

These partial differential equations can be integrated with respect to x and written as a function of the perturbed variables in space and time:

$$\begin{aligned}\acute{u} &= f(t) \\ \acute{p} &= g(t)x + h(t).\end{aligned}\tag{B.6}$$

Boundary conditions applied at the extents of the channel dictate that inlet stagnation pressure is held constant in time while a sinusoidal variation in exit pressure is specified as:

$$\begin{aligned}p^0(x=0) &= p(0,t) + \frac{1}{2}\rho u(0,t)^2 \\ p(L,t) &= \bar{p} + \acute{p}_e \sin(wt).\end{aligned}\tag{B.7}$$

Substituting the definitions provided in equations B.3 the boundary conditions can be linearized by dropping any of the higher order unsteady terms. Repeating the technique introduced above, one can integrate over a time period and take advantage of the definitions in equation B.4. This provides a linearized set of equations for the perturbed variable

$$\begin{aligned}\acute{p}(0,t) + \rho \bar{u} \acute{u}(0,t) &= 0 \\ \acute{p}(L,t) &= \acute{p}_e \sin(wt).\end{aligned}\tag{B.8}$$

Using the boundary conditions and the momentum equation the temporal functions $f(t)$, $g(t)$, and $h(t)$ can be derived as

$$g(t)L + h(t) = \acute{p}_e \sin(wt)$$

$$\begin{aligned} h(t) + \rho \bar{u} f(t) &= 0 \\ \frac{df(t)}{dt} + \frac{1}{\rho} g(t) &= 0. \end{aligned} \quad (\text{B.9})$$

The above three equations can be combined to form an ordinary differential equation

$$\frac{df(t)}{dt} + \frac{\bar{u}}{L} f(t) = -\frac{\dot{p}_e}{\rho L} \sin(wt). \quad (\text{B.10})$$

This ODE can be simplified into a form which can be directly integrated

$$\begin{aligned} \frac{d\left(e^{\frac{\bar{u}t}{L}} f(t)\right)}{dt} &= -e^{\frac{\bar{u}t}{L}} \frac{\dot{p}_e}{\rho L} \sin(wt) \\ f(t) &= -\frac{\dot{p}_e}{\rho L} e^{-\frac{\bar{u}t}{L}} \left[\int^t e^{\frac{\bar{u}s}{L}} \sin(ws) ds + c \right]. \end{aligned} \quad (\text{B.11})$$

This is integrated over an arbitrary length of time with the constant of integration determined by the initial values of the solution. For an initially steady flow the velocity perturbations are quantified as:

$$\begin{aligned} f(t) &= \dot{u} = -\frac{\dot{p}_e}{\rho \bar{u}} \frac{1}{1 + \Omega^2} \left[\sin(wt) - \Omega \cos(wt) + \Omega e^{-\frac{\bar{u}t}{L}} \right] \\ \Omega &= \frac{wL}{\bar{u}}. \end{aligned} \quad (\text{B.12})$$

Having solved for $f(t)$, solutions for the temporal functions $g(t)$ and $h(t)$ can be solved by employing equation B.9. Combinations of these functions result in an equation for the pressure perturbation as originally provided by Merkle and Athavale:

$$\dot{p}(x, t) = \dot{p}_e \sin(wt) + \dot{p}_e \left(\frac{x - L}{L} \right) \frac{\Omega}{1 + \Omega^2} \left[\Omega \sin(wt) + \cos(wt) - e^{-\frac{\bar{u}t}{L}} \right]. \quad (\text{B.13})$$

B.2 Channel Flow Results

The solution in equations B.13 and B.12 includes terms representing the exponential decay of the initial transients. The NLFD solver does not admit these terms given that they are zero at a periodic steady state. This periodic solution can be restated

as follows:

$$\begin{aligned}\dot{u}(t) &= -\frac{\dot{p}_e}{\rho \bar{u}} \frac{1}{1 + \Omega^2} [\sin(\omega t) - \Omega \cos(\omega t)] \\ \dot{p}(x, t) &= \dot{p}_e \sin(\omega t) + \dot{p}_e \left[\frac{x - L}{L} \right] \frac{\Omega}{1 + \Omega^2} [\Omega \sin(\omega t) + \cos(\omega t)].\end{aligned}\quad (\text{B.14})$$

A one-dimensional incompressible NLFD flow solver was developed based on the theory of artificial compressibility. Because the analytic solution of pressure and velocity is linear in space, a first order spatial differencing scheme was used which should represent this analytic solution without error. The solver used convergence acceleration techniques similar to those documented in previous chapters of this thesis; a simpler multigrid V-cycle was employed with a multistage time advancement scheme.

Results from an arbitrary test case are provided to illustrate the solution. Table B.1 displays the boundary condition parameters that define this case while figures B.2 and B.3 plot the velocity and pressure in space and time. As the analytic solution would suggest both plots produce ruled surfaces where the solution varies linearly in space and sinusoidally in time.

Nomenclature	Symbol	Value
Exit pressure perturbation	$\frac{\dot{p}_e}{\bar{p}}$	1e-4
Reduced frequency	Ω	1.0
Inlet stagnation pressure	p_o	constant

Table B.1: Boundary conditions used for test case.

A parameter survey was designed to verify the accuracy of the solver. Since the spatial and temporal discretization error should be zero, the only error contribution should be from the second order terms dropped in the derivation of the linearized solution. Consequently, the difference between the analytic and NLFD solutions should decay by the square of the exhaust pressure perturbation. A set of test cases was run with the NLFD solver over a range of 12 orders of magnitude in the exhaust pressure perturbation. The infinity norm of the difference between the analytic and numerical solution is provided in figure B.4.




Figure B.2: Three-dimensional plot of channel velocity as a function of streamwise coordinate and time.

The figure shows that the error decays to a minimum which corresponds to the numerical precision of the computer. For all solutions, calculated within the precision of the machine, the error predictably decays like the square of the exhaust pressure perturbation.

B.3 Channel Flow Convergence

This section investigates the convergence rates of the solver as a function of unsteady parameters like exhaust pressure perturbation and the decay rate of the initial transients. Given the the accuracy of the linearization is a function of the magnitude of the unsteadiness, then the exhaust pressure perturbation is an indicator of the nonlinearity inherent in the solution. Plotting convergence rates as a function of this parameter shows the sensitivity of the solver to nonlinearities in the solution. Figure B.5 shows the residual in the fundamental harmonic as a function of the iteration number in the solver. Each line on the figure represents a separate case with a unique exhaust pressure perturbation. The residuals all begin at different levels because the




Figure B.3: Three-dimensional plot of channel pressure as a function of streamwise coordinate and time.

uniform initial solution becomes closer to the unsteady solution as the unsteady perturbations decrease. However, the slope of all the lines on this graph are the same indicating that the convergence rate of the solver is independent of the magnitude of the exhaust pressure perturbation.

In contrast to time-accurate solvers, a frequency domain solver attempts to directly solve for the periodic steady state. The decay rate of the initial transients, which adversely affect time-accurate solvers, should not affect the performance of the NLFD solver. This decay rate for the channel flow problem, is quantified by the exponential term in equations B.12 and B.13. Table B.2 documents the results of a parametric survey that varied the $\frac{\bar{u}}{L}$ term by 10 orders of magnitude while varying the frequency such that the term $\frac{T\bar{u}}{\Delta x}$ remained constant. The data shows that the same number of multigrid cycles are required to obtain a fully converged solution over a range in transient decay rates. This illustrates that for a fixed spatial resolution, the convergence rates of this solver are independent of the physical eigenvalues of the problem.




Figure B.4: The infinity norm of the difference between analytic and NLFD solutions as a function of exhaust pressure perturbation.

$\frac{T\bar{u}}{\Delta x}$	Decay Rate of Initial Transients	Multigrid Cycles
10	$10^{-10} \rightarrow 1$	206
100	$10^{-10} \rightarrow 1$	98

Table B.2: The number of multigrid cycles required to achieve machine zero convergence. Two separate cases are provided at different values of $\frac{T\bar{u}}{\Delta x}$. Within each case, the initial transients decay rate was varied by ten orders of magnitude.




Figure B.5: Absolute value of the residual as a function of the multigrid cycle for a number of exhaust pressure perturbations.

Bibliography

- [1] J.J. Adamczyk. Model Equation for Simulating Flows in Multistage Turbomachinery. *NASA Technical Memorandum 86869*, NASA, November 1984.
- [2] K. Akselvoll and P. Moin. *Large Eddy Simulation of Turbulent Confined Jets and Turbulent Flow Over a Backward Facing Step. Dissertation*, Stanford University, February 1995.
- [3] J.J. Alonso. Personal Notes on Rizzi Boundary Condition Implementation, December 2000.
- [4] J.J. Alonso, T.J. Mitty, L. Martinelli, and A. Jameson. A Two-Dimensional Multigrid Navier-Stokes Solver for Multiprocessor Architectures. In *Proceedings of Parallel CFD*, Kyoto, Japan, May 1994.
- [5] A. Arnone, M. Liou, and L.A. Povinelli. Integration of Navier-Stokes Equations using Dual Time Stepping and a Multigrid Method. *AIAA Journal*, 33(6):985–990, June 1995.
- [6] H. Ashley. Role of Shocks in the "Sub-Transonic" Flutter Phenomenon. *Journal of Aircraft*, 17(3):187–197, 1980.
- [7] R.M. Beam and R.F. Warming. An implicit finite-difference algorithm for hyperbolic systems in conservation law form. *Journal of Computational Physics*, 22:87–110, 1976.

- [8] P. Beran and W. Silva. Reduced-Order Modeling: New Approaches for Computational Physics. *AIAA paper 01-0853*, AIAA 39th Aerospace Sciences Meeting, Reno, NV, January 2001.
- [9] A. Brandt. Multi-Level Adaptive Solutions to Boundary Value Problems. *Math. Comp.*, 31:333–390, 1977.
- [10] W.R. Briley and H. McDonald. Solution of the Three-Dimensional Compressible Navier-Stokes Equations by an Implicit Technique. In *Proc. Fourth International Conference on Numerical Methods in Fluid Dynamics*, Lecture Notes in Physics, pages 105–110. Springer-Verlag, 1974.
- [11] T. Chen, P. Vasanthakumar, and L. He. Analysis of Unsteady Blade Row Interaction Using Nonlinear Harmonic Approach. *Journal of Propulsion and Power*, 17(3):651–658, 2001.
- [12] R. Davis. Personal Communications on TFLO, January-December 2001.
- [13] R.L. Davis, T. Shang, J. Buteau, and R. Ni. Prediction of 3-(d) Unsteady Flow in Multi-Stage Turbomachinery using an Implicit Dual Time-Step Approach. Technical Report 96-2565, AIAA, ASME, SAE, and ASEE 32nd Joint Propulsion Conference and Exhibit, Lake Buena Vista, FL, July 1996.
- [14] S.S. Davis. NACA 64A010 (NASA Ames Model) Oscillatory Pitching. *AGARD Report 702*, AGARD, January 1982. Dataset 2.
- [15] E. Dowell. Unsteady Transonic Aerodynamics and Aeroelasticity. Technical report, Recent Advances in Aerodynamics: Proceedings of an International Symposium held at Stanford University, Palo Alto, California, August 1983.
- [16] E. Dowell, S.R. Bland, and M.H. Williams. Linear/Nonlinear Behavior in Unsteady Transonic Aerodynamics. *AIAA paper 81-0643*, AIAA 22nd Structures, Structural Dynamics and Materials Conference, April 1981.
- [17] M. Van Dyke. *An Album of Fluid Motion*. Parabolic Press, Stanford, CA, 1982.

- [18] R.P. Fedorenko. The Speed of Convergence of One Iterative Process. *USSR Comp. Math and Math Phys.*, 4:227–235, 1964.
- [19] A. Gelb and E. Tadmor. Enhanced Spectral Viscosity Approximations for Conservation Laws. *Applied Numerical Mathematics*, 33:3–21, 2000.
- [20] L. Green, Q. Zhang, J. Garriz, S. Wang, V. Vatsa, K. Haigler, and P. Newman. NASA/CAE Wind Tunnel Interference Cooperative Program—Status and Sample Results. *ICAW Paper-W1*, International Conference on Adaptive Wall Wind Tunnel Research and Wall Interference Correction, 1991.
- [21] L.L. Green and P.A. Newman. Wall-Interference Assessment and Corrections for Transonic NACA 0012 Airfoil Data from Various Wind Tunnels. *NASA Technical Paper 3070*, NASA, April 1991.
- [22] B. Gustafsson, H. Kreiss, and J. Oliger. *Time Dependent Problems and Difference Methods*. John Wiley and Sons, New York, NY, 1995.
- [23] K.C. Hall. Seminar on Harmonic Balance Techniques, January 2000. Presented to the Aerospace Computing Lab, Stanford University.
- [24] K.C. Hall, J.P. Thomas, and W.S. Clark. Computation of Unsteady Nonlinear Flows in Cascades using a Harmonic Balance Technique. Technical report, 9th International Symposium on Unsteady Aerodynamics, Aeroacoustics and Aeroelasticity of Turbomachines, Lyon, France, September 2000.
- [25] L. He and W. Ning. Nonlinear Harmonic Analysis of Unsteady Transonic Inviscid and Viscous Flows. In *Proceedings of the 8th International Symposium on Unsteady Aerodynamics and Aeroelasticity of Turbomachines*, Stockholm, Sweden, September 1998.
- [26] R.D. Henderson. Details of the Drag Curve Near the Onset of Vortex Shedding. *Physics of Fluids*, 1995.
- [27] A. Iserles. *A First Course in the Numerical Analysis of Differential Equations*. Cambridge University Press, Cambridge, United Kingdom, 1996.

- [28] A. Jameson. Solution of the Euler Equations for Two Dimensional Transonic Flow by a Multigrid Method. *Applied Mathematics and Computation*, 13:327–355, 1983.
- [29] A. Jameson. Transonic Flow Calculations. Princeton University Report MAE 1651, Princeton University, 1984. Published in Numerical Methods in Fluid Dynamics Edited by F. Brezzi, Lecture Notes in Mathematics Vol 1127, Springer Verlag, 1985, pp 156-242.
- [30] A. Jameson. Non-Oscillatory Shock Capturing Scheme Using Flux Limited Dissipation. In B.E. Engquist, S. Osher, and R.C.J. Sommerville, editors, *Large-Scale Computations in Fluid Mechanics*, volume 22 of *Lectures in Applied Mathematics*, page 345. 1985.
- [31] A. Jameson. Multigrid Methods for Compressible Flow Calculations. In *Proc. Second European Conference on Multigrid Methods*, Lecture Notes in Mathematics, pages 166–201. Springer-Verlag, 1986.
- [32] A. Jameson. Computational Transonics. *Comm. Pure Applied Mathematics*, 16:507–549, 1988.
- [33] A. Jameson. Computational Aerodynamics for Aircraft Design. *Science*, 245:361–371, 1989.
- [34] A. Jameson. Time Dependent Calculations Using Multigrid, with Applications to Unsteady Flows Past Airfoils and Wings. Technical Report 91-1596, AIAA 10th Computational Fluid Dynamics Conference, June 1991.
- [35] A. Jameson. Analysis and Design of Numerical Schemes for Gas Dynamics I Artificial Diffusion, Upwind Biasing, Limiters and their Effect on Accuracy and Multigrid Convergence, RIACS Technical Report 94.15. *International Journal of Computational Fluid Dynamics*, 4:171–218, 1995.
- [36] A. Jameson. Analysis and Design of Numerical Schemes for Gas Dynamics II Artificial Diffusion and Discrete Shock Structure, RIACS Report No. 94.16. *International Journal of Computational Fluid Dynamics*, 5:1–38, 1995.

- [37] A. Jameson. A Perspective on Computational Algorithms for Aerodynamic Analysis and Design. Technical report, Sixth National Conference on Computational Fluid Dynamics, Taitung, Taiwan ROC, August 1999.
- [38] A. Jameson, R. E. Melnik, and H. R. Mead. A multigrid method for the computation of viscid/inviscid interactions in airfoils. *AIAA paper 83-0234*, AIAA 21st Aerospace Sciences Meeting, Reno, NV, January 1983.
- [39] A. Jameson, W. Schmidt, and E. Turkel. Numerical Solutions of the Euler Equations by Finite Volume Methods Using Runge-Kutta Time-Stepping Schemes. *AIAA paper 81-1259*, AIAA 14th Fluid and Plasma Dynamics Conference, June 1981.
- [40] D.A. Johnson and W.D. Bachalo. Transonic Flow Past a Symmetrical Airfoil - Inviscid and Turbulent Flow Properties. *AIAA Journal*, 18(1):16–24, 1980.
- [41] D. Kwak, W.C. Reynolds, and J.H. Ferziger. *Three-Dimensional Time Dependent Computation of Turbulent Flow. Dissertation*, Stanford University, November 1975.
- [42] R.H. Landon. NACA 0012 Oscillatory and Transient Pitching. *AGARD Report 702*, AGARD, January 1982. Dataset 3.
- [43] H. Lomax, T. Pulliam, and D. Zingg. *Fundamentals of Computational Fluid Dynamics*. Pre-Release Version provided as part of Stanford’s AA-214 Numerical Methods in Fluid Mechanics, 1997.
- [44] T. Manning and S. Lele. *A Numerical Investigation of Sound Generation in Supersonic Jet Screech. Dissertation*, Stanford University, November 1999.
- [45] L. Martinelli and A. Jameson. *Calculations Of Viscous Flows With A Multigrid Method. Dissertation*, Princeton University, October 1987.
- [46] W.J. McCroskey. Unsteady Airfoils. *Annual Review of Fluid Mechanics*, 14:285–311, 1982.

- [47] W.J. McCroskey. A Critical Assessment of Wind Tunnel Results for the NACA 0012 Airfoil. *NASA Technical Memorandum 100019*, NASA, October 1987.
- [48] M. McMullen, A. Jameson, and J. Alonso. Acceleration of Convergence to a Periodic Steady State in Turbomachinery Flows. *AIAA paper 01-0152*, AIAA 39th Aerospace Sciences Meeting, Reno, NV, January 2001.
- [49] M. McMullen, A. Jameson, and J. Alonso. Application of a Non-Linear Frequency Domain Solver to the Euler and Navier-Stokes Equations. *AIAA paper 02-0120*, AIAA 40th Aerospace Sciences Meeting, Reno, NV, January 2002.
- [50] N. D. Melson, M. D. Sanetrik, and H. L. Atkins. Time-Accurate Navier-Stokes Calculations with Multigrid Acceleration. In *Proceedings of the Sixth Copper Mountain Conference on Multigrid Methods*, Copper Mountain, April 1993.
- [51] C.L. Merkle and M. Athavale. Time-Accurate Unsteady Incompressible Flow Algorithms Based on Artificial Compressibility. *AIAA paper 87-1137*, AIAA 8th Computational Fluid Dynamics Conference, 1987.
- [52] R.E. Mineck and L.L. Green. Wall Interference Assessment/Correction for Transonic Airfoil Data from Porous and Shaped Wall Test Sections. *AIAA paper 90-1406*, AIAA 16th Aerodynamic Ground Testing Conference, June 1990.
- [53] B.L. Mitchell, S.K. Lele, and P. Moin. *Direct Computation of the Sound Generated by Subsonic and Supersonic Axisymmetric Jets. Dissertation*, Stanford University, November 1995.
- [54] P. Moin. *ME308 Spectral Methods in Computational Physics*. Class notes provided as part of Stanford's ME308 Course, 1997.
- [55] S. Nadarajah. Optimal Control of Unsteady Flows Using Time Accurate and Non-Linear Frequency Domain Methods. *AIAA paper 02-5436*, 9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization Conference, Atlanta, GA, September 2002.

- [56] D. Nixon. Prediction of Aeroelasticity and Unsteady Aerodynamic Phenomena in Transonic Flow. Technical Report Paper 134, Nielson Engineering and Research, Mountain View, CA, 1981.
- [57] C. Norberg. An Experimental Investigation of the Flow Around a Circular Cylinder: Influence of Aspect Ratio. *Journal of Fluid Mechanics*, 258:287–316, 1994.
- [58] N.A. Pierce and J.J. Alonso. Efficient Computation of Unsteady Viscous Flow by an Implicit Preconditioned Multigrid Method. *AIAA Journal*, 36:401–408, 1998.
- [59] A. Rizzi. Numerical Implementation of Solid-Body Boundary Conditions for the Euler Equations. *Zeitschrift für Angewandte Mathematik und Mechanik*, 7:T301–4, 1978.
- [60] A. Snider. An Improved Estimate of the Accuracy of Trigonometric Interpolation. *SIAM Journal on Numerical Analysis*, 9(3):505–508, 1972.
- [61] G.R. Srinivasan, J.A. Ekaterinaris, and W.J. McCroskey. Evaluation of Turbulence Models for Unsteady Flows of an Oscillating Airfoil. *Computers & Fluids*, 24(7):833–861, 1995.
- [62] E. Tadmor. Convergence of Spectral Methods for Nonlinear Conservation Laws. *SIAM Journal on Numerical Analysis*, 26(1):30–44, 1989.
- [63] S. Taneda. Experimental Investigation of the Wakes Behind Cylinders and Plates at Low Reynolds Number. *Jour. Phys. Soc. Japan*, 11:302–307, 1104–1108, 1956.
- [64] H. Tijdeman. Transonic Flow Past Oscillating Airfoils. *Annual Review of Fluid Mechanics*, 12:181–222, 1980.
- [65] V. Venkatakrishnan and A. Jameson. Computation of Unsteady Transonic Flows by the Solution of the Euler Equations. *AIAA Journal*, 26:974–981, 1988.
- [66] M. Vinokur. Conservation equations of gas dynamics in curvilinear coordinate systems. *Journal of Computational Physics*, 14:105–125, 1974.

- [67] H. Viviand. Conservative Forms of Gas Dynamic Equations. *Rech. Aerosp.*, 1971-1:65–68, 1974.
- [68] C.H.K. Williamson. Fluid Dynamics Research Laboratories of Prof. Charles HK Williamson.
- [69] C.H.K. Williamson. Defining a Universal and Continuous Strouhal-Reynolds Number Relationship for the Laminar Vortex Shedding of a Circular Cylinder. *Phys. Fluids*, 31:2742, 1988a.
- [70] C.H.K. Williamson. The Existence of Two Stages in the Transition to Three-Dimensionality of a Cylinder Wake. *Phys. Fluids*, 31:3165, 1988b.
- [71] C.H.K. Williamson. Oblique and Parallel Modes of Vortex Shedding in the Wake of a Circular Cylinder at Low Reynolds Numbers. *Journal of Fluid Mechanics*, 206:579–627, 1989.
- [72] C.H.K. Williamson. Three-Dimensional Wake Transition. *J. Fluid Mech.*, 328:345–407, 1996.
- [73] C.H.K. Williamson. Vortex Dynamics in the Cylinder Wake. *Annual Review Fluid Mech.*, 28:477–539, 1996.
- [74] C.H.K. Williamson and G.L. Brown. A Series in $(1/\sqrt{\text{Re}})$ to Represent the Strouhal-Reynolds Number Relationship of the Cylinder Wake. *J. Fluids and Struc.*, 1998.
- [75] C.H.K. Williamson and A. Roshko. Measurements of Base Pressure in the Wake of a Cylinder at Low Reynolds Numbers. *Z. Flugwiss. Weltraumforsch.*, 1990.
- [76] A. A. Wray. Very Low Storage Time-Advancement Schemes. *Internal Technical Report*, NASA, Ames Research Center, Moffett Field, California, 1986.
- [77] J. Yao, R. Davis, J.J. Alonso, and A. Jameson. Unsteady Flow Investigations in an Axial Turbine Using the Massively Parallel Flow Solver TFLO. Technical Report 01-0529, AIAA 39th Aerospace Sciences Meeting and Exhibit, January 2001.

- [78] J. Yao, A. Jameson, J.J. Alonso, and F. Liu. Development and Validation of a Massively Parallel Flow Solver for Turbomachinery Flows. Technical Report 00-0882, AIAA 38th Aerospace Sciences Meeting and Exhibit, January 2000.
- [79] H.C. Yee. On Symmetric and Upwind TVD Schemes. In *Proceedings of 6th GAMM Conference on Numerical Methods in Fluid Mechanics*, Gottingen, September 1985.
- [80] S.T. Zalesak. Fully Multi-Dimensional Flux-Corrected Transport Algorithms for Fluids. *Journal of Computational Physics*, 31:335–362, 1979.