

AN OVERSET TIME-SPECTRAL METHOD
FOR RELATIVE MOTION

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

Joshua Leffell

June 2014

© Copyright by Joshua Leffell 2015
All Rights Reserved
Last Modified July 2015

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 J. 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.

(Thomas H. Pulliam)

Approved for the University Committee on Graduate Studies

Abstract

Periodic flow arises in a broad range of aerodynamic applications, including turbomachinery, rotorcraft and flapping-wing configurations. The standard procedure for simulating such flows involves integrating the governing equations forward in time until a statistically-stationary flow is achieved. Such flows, while unsteady in the time domain, are invariant in the frequency domain. As such, the governing equations can be transformed into a steady set of equations by expanding the temporal variation at every point in space as a truncated Fourier series. The Time-Spectral method is a Fourier pseudospectral scheme developed to exploit this fact, while maintaining the ability to resolve the nonlinear dynamics of the unsteady governing equations. Fourier collocation techniques have been shown to reduce the computational costs associated with periodic steady-state flows by up to an order of magnitude by obviating the need to simulate through initial transients in physical time and because of the spectral convergence properties of the Fourier series.

This dissertation extends the Time-Spectral method to overset flow solvers in a general manner. Overset grid technology provides the ability to resolve geometry of arbitrary complexity with structured grids that offer enhanced boundary-layer resolution and efficient data structures. However, relative motion between components results in spatial nodes that lack complete time histories as they are dynamically removed from the computational domain when located within the impermeable boundaries of solid bodies. Since the infinite support of the complex exponential basis functions prevents direct application of the standard Time-Spectral method, a novel hybrid Time-Spectral method has been developed that expands the temporal variation at dynamically-blanked nodes in an alternative manner. Investigation of a

number of strategies to handle dynamically-blanked nodes is described. However, the proposed scheme is extensible by construction to enable the incorporation of future improvements with minimal effort.

The hybrid Time-Spectral discretization has been incorporated within NASA's well-established three-dimensional implicit Reynolds-averaged Navier-Stokes (RANS) solver OVERFLOW. Details concerning the implicit approximate factorization scheme, turbulence modeling, multigrid acceleration and dealiasing are addressed with respect to the Time-Spectral implementation. The augmented Time-Spectral OVERFLOW solver is validated by direct comparison with the existing time-accurate solver for a variety of numerical experiments. A series of two-dimensional oscillating airfoil cases are included to assess the ability of the hybrid Time-Spectral scheme to handle high-frequency, transonic and turbulent flows. The relatively inexpensive two-dimensional test cases serve as meaningful model problems that are used to uncover solution strategies for more expensive realistic three-dimensional flows. The thesis concludes with a pair of three-dimensional calculations of the quarter-scale V-22 Tilt Rotor Aeroacoustic Model (TRAM) in hover and forward (edgewise) flight. The Time-Spectral hover simulation matches the time-accurate calculation using a single harmonic. Significantly more temporal modes and dealiasing are required to accurately compute the forward flight case because of its more active, high-frequency wake.

Acknowledgements

This process has been a wonderful learning experience and I will be forever indebted to the host of people who have helped me along the way.

First and foremost, I would like to thank both of my primary advisors: Professor Antony Jameson of Stanford University and Thomas Pulliam of NASA Ames Research Center. I greatly appreciate the advice Professor Jameson has offered over the years, and I hope that we can continue working together in the future. I would like to thank Tom for inviting me to visit NASA during his introductory CFD course. Working with Tom has been a tremendous experience and I will be forever grateful for the opportunities he has provided and for his mentorship over the more than five years since I first joined him at NASA. The opportunity to work with both of them has been a privilege.

I would like to thank Professor Juan Alonso for his participation on my reading committee. I am fortunate to have taken several of his classes and have greatly appreciated his interest and input during my time at Stanford. I would like to acknowledge Professor Robert MacCormack for serving on my oral exam committee. His CFD classes were an integral part of my education and I look back fondly to his enlightened digressions on the wonders of *Nonlinearland*. I would also like to thank Professor Michael Saunders for serving as the chair of my oral exam committee and providing some helpful typesetting advice. I must also express my appreciation to the Late Entrance Accelerated Program (LEAP) at Boston University for preparing me for graduate study and the Department of Aeronautics and Astronautics at Stanford for taking a chance on me, despite my limited background.

I owe much gratitude to the people at the NASA Advanced Supercomputing division, both for their camaraderie and for sharing with me their technical expertise. In particular, I would like to acknowledge Scott Murman for his invaluable contributions to this research. I will be forever grateful for the time he has invested and for helping to *keep the blinders on* to focus my attention on the relevant topics. I would also like to thank the members of the Stanford outpost at Ames during my time there: Mathias Wintzer, Noël Bakhtian, George Anderson, Samira Motiwala and Jonathan Chiew.

Additionally, this research would not have been possible without the funding provided by Roger Strawn at the Army Aeroflightdynamics Directorate (AFDD). In addition to the financial support, I would like to acknowledge Roger and the rest of the group for their considerable technical input. I would also like to thank the NASA Education Associates program for facilitating my research position.

Finally, I would like to express my sincere gratitude to my family for their unwavering support during this long and circuitous journey. I would also like to thank my friends and last, but not least, Cici for putting up with my antics over the course of these last few months.

Nomenclature

η	Body-normal coordinate in the computational domain
ξ	Streamwise coordinate in the computational domain
$\Delta\eta$	Grid spacing in the η -direction
$\Delta\xi$	Grid spacing in the ξ -direction
α	Angle of attack
ΔQ	Update vector of conserved quantities in the computational domain
ΔQ	Update vector of conserved quantities in the physical domain
Δx	Grid spacing in the x -direction
Δy	Grid spacing in the y -direction
\Im	Imaginary component
λ	Eigenvalue
$\rho\mathbf{u}$	Fluid momentum vector
\mathbf{u}	Fluid velocity vector
\mathcal{A}	Flux Jacobian in the ξ -direction
\mathcal{B}	Flux Jacobian in the η -direction
\mathcal{F}	Flux vector in the ξ -direction
\mathcal{G}	Flux vector in the η -direction
Q	Vector of conserved quantities in the computational domain
Ω	Spatial domain
ω	Fundamental frequency
$\partial\Omega$	Spatial domain boundary
ϕ_k	Basis function corresponding to the k^{th} mode or node
Ψ	Azimuthal angle

ψ_k	Test function corresponding to the k^{th} mode or node
\Re	Real component
ρ	Fluid density
σ	Amplification factor
τ	Pseudotime
θ_0	Collective pitch amplitude
θ_{C1}	Longitudinal cyclic-pitch amplitude
θ_{S1}	Lateral cyclic-pitch amplitude
$\tilde{\nu}$	Undamped eddy viscosity
a	Plunging amplitude
b	Fourier continuation parameter
c	Chord length
d	Barycentric rational interpolant approximation order
d_{\max}	Barycentric rational interpolant approximation order limit
e	Total energy
F	Flux vector in the x -direction
G	Flux vector in the y -direction
h	Nondimensional plunging amplitude, $h = a/c$
i	Imaginary unit, $\sqrt{-1}$
J	Mesh Jacobian, V^{-1}
K	Number of temporal modes, $K = (N - 1) / 2$
k	Reduced frequency, $\omega c / V_\infty$
M	Mach number
N	Number of temporal degrees of freedom
N_Q	Number of conserved quantities
N_{sd}	Number of spatial dimensions
Q	Vector of conserved quantities in the physical domain
R	Rotor radius
Re	Reynolds number
St	Strouhal number, kh
T	Temporal period

t	Physical time
V	Reference velocity or volume, J^{-1}
w	Weight function

Superscripts

*	Conjugate transpose
n	Iteration index in physical time
s	Iteration index in pseudotime

Operators

δ_x	General finite-difference operator in the x -direction
\mathcal{D}_N	Differentiation operator in the time domain, $\mathbf{u}'_N = \mathcal{D}_N \mathbf{u}_N$
\mathcal{L}	Left-hand side operator
\mathcal{R}	Generalized residual
Φ	Transformation operator to time domain
Φ^{-1}	Transformation operator from time domain
D_N	Differentiation operator in transformed domain
I_N	Interpolation
P_N	Projection
FFT	Fast Fourier Transform
IFFT	Inverse Fast Fourier Transform

Subscripts

∞	Freestream value
tip	Tip value

Contents

Abstract	v
Acknowledgements	vii
Nomenclature	ix
1 Introduction	1
1.1 Previous Research	3
1.2 Research Contributions	6
1.3 Dissertation Outline	9
2 The Time-Spectral Method	11
2.1 Fourier Collocation in Time	12
2.2 Aliasing & Spectral Vanishing Viscosity	19
3 Overset Grid Methodology	24
3.1 Overview	25
3.2 Representative One-Dimensional Grid System	27
4 Dynamic Hole Cutting	34
4.1 Spatial Smoothing	37
4.2 Global Expansion of the Solution	38
4.3 Local Expansion of the Solution	45
4.3.1 Periodization	46
4.3.2 Bounded Interval	49

4.4	Mixed Expansion of the Solution	62
5	OVERFLOW	68
5.1	Governing Equations	69
5.2	The Standard OVERFLOW Solver	70
5.2.1	Implicit Factored Scheme	70
5.2.2	Multigrid Acceleration	75
5.2.3	Turbulence Model	75
5.2.4	Connectivity	76
5.3	The Time-Spectral OVERFLOW Solver	76
5.3.1	Implicit Factored Scheme	77
5.3.2	Multigrid Acceleration	80
5.3.3	Temporal Vanishing Viscosity	82
5.3.4	Turbulence Model	84
5.3.5	Connectivity	85
5.3.6	Additional Resource Requirements	86
6	Numerical Results	88
6.1	Two-Dimensional Oscillating Airfoils	89
6.1.1	Inviscid Plunging NACA 0012 Airfoil	89
6.1.2	Laminar Plunging NACA 0012 Airfoil	104
6.1.3	Transonic Pitching NACA 0012 Airfoil	113
6.2	Three-Dimensional V-22 Tiltrotor	124
6.2.1	Hover	125
6.2.2	Forward (Edgewise) Flight	130
7	Conclusions	136
7.1	Summary	136
7.2	Limitations	140
7.3	Future Work	142

A	The Navier-Stokes Equations	146
A.1	Cartesian Coordinates	146
A.2	Generalized Curvilinear Coordinates	147
B	Linear Stability Analysis	151
B.1	Standard Time-Spectral Factored Scheme	152
B.2	Hybrid Time-Spectral Factored Scheme	158
	Bibliography	161

List of Tables

6.1	Isolated V-22 Osprey Tiltrotor in Hover. Grid statistics.	126
6.2	Isolated V-22 Osprey Tiltrotor in Hover. Force and moment coefficients and Figure of Merit.	128
6.3	Isolated V-22 Osprey Tiltrotor in Forward Flight. Grid statistics. . .	131

List of Figures

2.1	Aliasing of sine waves.	21
3.1	Two-dimensional overlapping grid-system of a NACA 0012 airfoil. . .	26
3.2	Hole cut at three time instances for a plunging airfoil.	28
3.3	Representative two-grid, one-dimensional overlapping grid-system. . .	29
4.1	Figurative piston trajectory and solution histories.	35
4.2	Spatial smoothing Time-Spectral test case.	39
4.3	Global least-squares Fourier projection and residual.	43
4.4	Fourier continuation of $f(x) = x$ for $x \in [0, 1]$ and its derivative. . . .	50
4.5	Local least-squares Fourier continuation (projection) and residual. . .	51
4.6	Barycentric rational interpolant basis functions.	56
4.7	Barycentric rational interpolation and residual.	57
4.8	Harmonic and Runge's analytic functions and derivatives.	58
4.9	Convergence of differentiated harmonic function.	59
4.10	Convergence of differentiated Runge's function.	60
4.11	Continuous trajectory of moving solid object.	62
4.12	Continuous versus discrete unblanked intervals at a representative node.	63
4.13	Mixed expansion of the solution.	65
5.1	Multigrid acceleration for a representative calculation.	83
5.2	Estimated OVERFLOW Time-Spectral memory usage.	87
6.1	Inviscid Plunging Airfoil. Dynamically-blanked region.	91

6.2	Inviscid Plunging Airfoil. Time-Spectral versus time-accurate values of streamwise momentum, ρu , at node a	93
6.3	Inviscid Plunging Airfoil. Convergence of various hybrid Time-Spectral schemes to the time-accurate solution at node a	94
6.4	Inviscid Plunging Airfoil. Time-Spectral versus time-accurate values of streamwise momentum, ρu , at node b	95
6.5	Inviscid Plunging Airfoil. Convergence of various hybrid Time-Spectral schemes to the time-accurate solution at node b	96
6.6	Inviscid Plunging Airfoil. Time-Spectral versus time-accurate drag coefficients for $N \in \{3, 9, 33\}$	99
6.7	Inviscid Plunging Airfoil. Convergence of rigid- and relative-motion Time-Spectral drag signals.	100
6.8	Inviscid Plunging Airfoil. Time and frequency response of a plunging airfoil for rigid motion.	102
6.9	Inviscid Plunging Airfoil. Time and frequency response of a plunging airfoil for relative motion.	102
6.10	Inviscid Plunging Airfoil. Reconstruction of streamwise momentum at a node on the off-body grid.	103
6.11	Laminar Plunging Airfoil. Grid system.	105
6.12	Laminar Plunging Airfoil. Drag-producing case. Visualization of vorticity magnitude; experimental, time accurate and Time Spectral. . .	107
6.13	Laminar Plunging Airfoil. Drag-producing case. Time-Spectral versus time-accurate drag coefficients for $N \in \{3, 9, 33\}$	108
6.14	Laminar Plunging Airfoil. Drag-producing case. Convergence of rigid- and relative-motion Time-Spectral drag signals.	109
6.15	Laminar Plunging Airfoil. Thrust-producing case. Visualization of vorticity magnitude; experimental, time accurate and Time Spectral. . .	110
6.16	Laminar Plunging Airfoil. Thrust-producing case. Time-Spectral versus time-accurate drag coefficients for $N \in \{3, 9, 33\}$	111
6.17	Laminar Plunging Airfoil. Thrust-producing case. Convergence of rigid- and relative-motion Time-Spectral drag coefficients	112

6.18	Inviscid Transonic Pitching Airfoil. Near-body grid and supersonic region.	113
6.19	Inviscid Transonic Pitching Airfoil. Iterative convergence of space-time residual.	114
6.20	Inviscid Transonic Pitching Airfoil. Time-Spectral versus time-accurate pitching moment coefficients for $N \in \{3, 9, 33\}$	116
6.21	Inviscid Transonic Pitching Airfoil. Time-Spectral versus time-accurate drag coefficients for $N \in \{3, 9, 33\}$	117
6.22	Inviscid Transonic Pitching Airfoil. Spectrum of Time-Spectral and time-accurate pitching moment coefficients.	118
6.23	Turbulent RANS Transonic Pitching Airfoil. Visualization of undamped turbulent eddy viscosity at $N = 5$ collocation points.	119
6.24	Turbulent RANS Transonic Pitching Airfoil. Continuous and Fourier representation of $\tilde{\nu}$ and its temporal derivative.	121
6.25	Turbulent RANS Transonic Pitching Airfoil. Time-Spectral versus time-accurate pitching moment coefficients for $N \in \{3, 9, 33\}$	122
6.26	Turbulent RANS Transonic Pitching Airfoil. Time-Spectral versus time-accurate drag coefficients for $N \in \{3, 9, 33\}$	123
6.27	V-22 Osprey Tiltrotor. (a) Full configuration and (b) definition of azimuthal angle, Ψ	125
6.28	Isolated V-22 Osprey Tiltrotor in Hover. Geometry and grid system.	127
6.29	Isolated V-22 Osprey Tiltrotor in Hover. Contours of vorticity magnitude.	129
6.30	Isolated V-22 Osprey Tiltrotor in Hover. Contours of undamped eddy viscosity.	129
6.31	Isolated V-22 Osprey Tiltrotor in Forward Flight. Top view of Isocontours of vorticity.	133
6.32	Isolated V-22 Osprey Tiltrotor in Forward Flight. Side view of Isocontours of vorticity.	134
6.33	Isolated V-22 Osprey Tiltrotor in Forward Flight. Reconstruction of Time-Spectral thrust coefficient.	135

7.1	Laminar Plunging NACA 0012 Airfoil. Thrust-producing case. Spurious oscillations on overset boundaries for relative motion.	141
7.2	Modal distribution for frequency-adaptive solution for the thrust-producing laminar plunging NACA 0012 airfoil.	144
A.1	Generalized curvilinear coordinate transformations.	149
B.1	Amplification spectra for the three-factor Time-Spectral scheme for fundamental frequencies 2π and 20π	157
B.2	Stability regions for the three-factor Time-Spectral scheme.	157
B.3	Amplification spectra for the four-factor Time-Spectral scheme for fundamental frequencies 2π and 20π	158
B.4	Stability regions for the four-factor Time-Spectral scheme.	159

Chapter 1

Introduction

The ongoing advancement of computational fluid dynamics (CFD) has greatly enhanced the ability to expedite the design of aerodynamic vehicles and associated components. This has been accomplished by iterating on simulations of the configuration prior to the physical construction of prototypes. Despite tremendous achievements in the field of CFD, the ever-increasing complexity of the simulations requires more efficient solution techniques in order to meaningfully integrate computation within the design process. In addition, costs grow significantly for simulations of unsteady phenomena, whose performance cannot be wholly described by a single state of the solution. A critical sub-class of unsteady flows are temporally periodic. Such flows arise in a broad range of fundamental aerodynamic applications including turbomachinery and rotorcraft. The application of valid assumptions in numerical methods can greatly increase the computational efficiency for appropriate classes of problems, and this is certainly true for the tremendously important case of time-periodic flow.

The standard procedure for simulating such flows involves integrating the unsteady governing equations forward in time until a statistically stationary *periodic steady state* is achieved. It is often necessary to simulate several periods to expel the initial transient and achieve a periodic response, making unsteady design optimization prohibitively expensive for many realistic problems. The motivation to improve the efficiency of these calculations led to the development of a variety of frequency-domain methods, which exploit the fact that temporally-periodic flow, while varying in the

time domain, is invariant in the frequency domain. Expanding the temporal variation at each spatial node as a truncated Fourier series transforms the unsteady governing equations into a coupled set of steady equations in integer harmonics. Upon convergence, the frequency-domain representation can be used to reconstruct the complete space-time solution.

The initial approaches were linearized frequency-domain methods, where each harmonic of the solution was solved independently. This strategy provided significant performance gains, however, the inability of such schemes to generally resolve nonlinear dynamics limited their application to a restrictive sub-class of periodic phenomena. Development of the Harmonic Balance (HB) method [1, 2] followed thereafter. This novel approach not only capitalized on the periodic nature of the solution, but did so while maintaining the ability to fully resolve nonlinearities inherent in the underlying physics. The spectral convergence rate of Fourier-based methods is vastly superior to the algebraic convergence rates of the finite-difference approximations typically used in traditional time-marching schemes for unsteady calculations. This disparity implies that a given level of accuracy can be achieved with significantly fewer degrees of freedom by employing a spectral method. Spectral methods have principally been applied in space, but can be motivated and developed in the same manner for the temporal dimension.

Time-domain Fourier pseudospectral methods were developed shortly thereafter with the objective of leveraging existing flow solvers and reducing the algorithmic complexity of the frequency-domain methods. In Fourier pseudospectral methods, the solutions at N equispaced time instances are coupled through the temporal derivative term, which is derived from the complex exponential basis functions of the Fourier series. The trigonometric representation of periodic phenomenon provides spectral convergence as the number of time-samples, N , and correspondingly, the number of resolvable harmonics, $K = (N - 1)/2$, increases.¹ Some approaches iterate the equations in the frequency domain; others operate solely in the time domain to further simplify the process of augmenting this capability to existing flow solvers. However, each approach harnesses the underlying steady solution in the frequency domain.

¹ N is taken to be odd unless otherwise indicated.

Fourier pseudospectral methods have provided an efficient means to predict periodic flows [2, 3, 4], but have yet to be generally and consistently applied to overset or Cartesian solvers due to the complications introduced by regions of the spatial domain that are periodically located within the impermeable boundaries of solid bodies. Overset and Cartesian grid methodologies are versatile techniques capable of handling complex geometry configurations with relative motion between components, and are commonly used for practical engineering applications.

In an arbitrary moving-body scenario for these approaches, a Lagrangian body moves through a fixed Eulerian mesh (or relative to another Lagrangian mesh). Mesh points interior to solid bodies are removed (*cut* or *blanked*), leaving a *hole* in the background mesh. Blanked mesh points are excluded from the computational domain within which the governing equations are solved; therefore, the solution at such nodes is undefined. In general, such grid points undergo dynamic blanking, where they are blanked for a portion of time and active otherwise. Such *dynamically-blanked* nodes lack a complete set of time samples, preventing the direct application of a standard Fourier pseudospectral approach due to the infinite support of the complex exponential basis functions of the Fourier series, upon which the method is based.

The objective of this thesis is to develop a robust and consistent Fourier pseudospectral method, capable of handling relative motion within an overset or Cartesian mesh framework, while still approaching the spectral convergence rate, and the associated computational efficiency of the standard approach. The integration of a Fourier pseudospectral approach with this general capability may potentially provide an enabling design and analysis tool.

1.1 Previous Research

With the goal of improving turbomachinery performance predictions, Hall et al. [1, 2] developed an approach that bridged the gap between time-linearized or time-averaged [5, 6, 7, 8, 9] frequency-domain methods that were unable to generally resolve the nonlinearities inherent in these types of flows, and traditional nonlinear time-marching schemes that were often intractably expensive for complex configurations due to their

protracted transients. The Harmonic Balance approach was more costly than its linearized analogue, as it coupled all of the solution modes together. However, this technique maintained a competitive advantage over the traditional approach of marching the nonlinear governing equations forward in physical time long enough to achieve a periodic steady state. This approach capitalized on the periodic nature of turbomachinery flows like the linearized frequency-domain solvers without abandoning the ability to fully resolve nonlinearities. Representing the temporal variation of the solution as a truncated Fourier-series afforded a highly-accurate description of the periodic flow with a limited number of temporal degrees of freedom.

Originally, the Harmonic Balance method was developed in such a way that the spatially-varying continuous Fourier-series coefficients were expressed analytically. This projection approach was consistent with a spectral method solving for the unknown frequency coefficients of the solution [10], with the desirable quality of inhibiting aliasing errors. However, the analytic approach was computationally expensive, required significant algorithmic development, and was not applicable to the more complex equations required to model turbulent flows [2]. Choosing to represent the solution with trigonometric interpolation, based on the *discrete* Fourier transform, enabled the calculation of equations of arbitrary complexity; the spatial residual evaluation routines of previously developed and well-validated time-domain flow solvers could be leveraged directly because the unknown Fourier coefficients are determined solely from time-domain data. The resulting method is a temporal pseudospectral (*collocation*) scheme that employs a uniform distribution of N temporal collocation points. The solution at each time sample is used to evaluate the nonlinear spatial residual and the temporal derivative is evaluated by the infinitely-supported, spectrally-accurate Fourier interpolation differentiation operator. This modified Harmonic Balance method has since been referred to as the High Dimensional Harmonic Balance (HDHB) method because varying N is a trivial endeavor permitting calculations using a large number of temporal degrees of freedom, contingent on the availability of the requisite computational resources. The discrete (time-domain) Fourier approach has demonstrated marked success in reducing the computational costs associated with simulating periodic flows [2, 3, 11, 4], and a significant body of

research has been undertaken since its initial introduction.

Other Fourier pseudospectral approaches, such as the Nonlinear Frequency Domain (NLFD) [3, 12, 11, 4, 13, 14, 15, 16] and Reduced Frequency [17] methods, elect to iterate the equations in the frequency domain, capitalizing on the efficient fast Fourier transform (FFT) and inverse transform (IFFT) operators. While the solution is updated in the frequency domain, it is transformed back into the time domain where the residual operator is evaluated. The residual is then transformed into the frequency domain to update the solution and the process iterates. Despite the algorithmic differences, this approach is mathematically consistent with the time-domain, HDHB scheme. Additionally, McMullen et al. [12] introduced a gradient-based approach for determining the fundamental frequency if it is unknown a priori, consequently widening the pool of candidate applications. The current work considers only forced periodic flows where the fundamental frequency of the response is assumed to be equivalent to the forcing frequency. The Time-Spectral method [18, 19], which iterates the discretized equations in the time domain, was developed in parallel and has been applied to turbomachinery [20, 21, 22], rotorcraft [23, 24] and aeroelasticity [25, 26, 27].

A major advantage of Fourier methods is the resulting transformation of the unsteady governing equations into a coupled set of steady equations; this led to the application of *steady* adjoint-based implementations capable of performing efficient unsteady design without resorting to the more costly unsteady-adjoint approach [24, 15, 28, 29]. Frequency-adaptive methods [30, 31, 32, 33, 34, 35, 36] more efficiently allocate temporal degrees of freedom by refining the number of temporal modes resolved at each spatial node to the frequency content in its solution, greatly reducing the computational effort and storage requirements. This adaptive temporal refinement technique may have significant application to the current work (See Chapter 7). Mavriplis and Yang [37] recently extended the Time-Spectral method to handle the class of *quasi-periodic* flows using the method of polynomial subtraction to resolve a slowly-varying aperiodicity in combination with strong periodic content. This hybrid Backward Difference Formula/Time-Spectral (BDF/TS) method has since been applied to rotorcraft in maneuver [38, 39] and aeroelastic flutter [26, 27]. Significant

effort has also been invested in developing more efficient implicit solution procedures for the Time-Spectral method by Mundis and Mavriplis [40, 27].

A number of previous efforts have focused on integrating Fourier pseudospectral temporal discretizations within overset or Cartesian flow solvers. Murman [17] demonstrated the Reduced-Frequency approach for a Cartesian solver with rigid domain motion, wherein the hole cutting remained fixed. Similarly, Custer [41] and Thomas et al. [42, 43] applied the HDHB method to the NASA overset OVERFLOW solver, but restricted the relative motion to maintain constant hole-cutting throughout the domain. Blanc et al. [25] implemented an iterative procedure to eliminate all dynamically-blanked nodes, if possible, to facilitate the use of the standard Time-Spectral method on dynamic overset grids. Soucy and Nadarajah [44] avoided dynamically-blanked nodes by extending the near-body grids such that they envelop the solid body at all time instances. The two previously mentioned approaches prove worthy for simple configurations, but are not a general solution for the case of arbitrary relative motion. These previous efforts focused on applying this method to Cartesian and overset meshes with constant blanking enabling the application of the standard Time-Spectral method, but Custer [41] and Blanc et al. [25] discuss the requirement for a novel approach to evaluate the time-derivative for the case of overset relative motion when complete time-histories are not universally available. More recently, Mavriplis et al. [39] demonstrated an approach to treat dynamically-blanked nodes by applying a Laplacian smoothing operator to populate the solution at blanked nodes with data from surrounding grid points.

1.2 Research Contributions

In the Time-Spectral method, the solution at every spatial node is represented as a truncated Fourier series. As described, overset calculations involving moving configurations generally introduce dynamically-blanked nodes which poses a fundamental obstacle in the pursuit of a general overset Time-Spectral method capable of handling arbitrary relative motion. The novel hybrid Time-Spectral approach developed in this thesis applies the standard Time-Spectral treatment at universally unblanked nodes,

but represents the solution at dynamically-blanked nodes in an alternative manner.

Numerical experiments demonstrate that the hybrid scheme mirrors the performance of the conventional Time-Spectral method, and monotonically converges to the comparable time-accurate simulations with increasing Time-Spectral modes. However, increased temporal resolution is required for the hybrid scheme for simulations involving large-amplitude relative-body motion due to the high-frequency response transferred from the moving body-conforming grids to the background grids. This section highlights the primary contributions of the research.

Development of the Hybrid Time-Spectral Scheme

Upon deciding to pursue the aforementioned hybrid approach, the primary developmental task was to determine an appropriate representation of the solution at dynamically-blanked nodes. The desire to leverage an existing codebase constrained the temporal collocation points at dynamically-blanked nodes to share the same uniform distribution as the standard Time-Spectral method. Three fundamental approaches were investigated:

- A global expansion of the solution using least-squares Fourier projection or least-norm Fourier interpolation from the available data at all unblanked time samples at a particular dynamically-blanked node. Preliminary results revealed that least-norm Fourier interpolation proved too unstable to merit further consideration, whereas least-squares Fourier projection demonstrated moderate success in the primary test case.
- A local expansion of the solution via periodization (Fourier continuation) or within the bounded interval of consecutively-unblanked time samples. Barycentric rational interpolants and splines proved most successful from the candidate pool of bounded interval approaches.
- A mixed expansion of the solution using a composite description of the solution.

The barycentric rational interpolant was selected here as the default treatment, but the hybrid Time-Spectral scheme has been developed as a general approach that only

requires a linear differentiation operator. Future investigation may uncover superior candidates that can be incorporated with minimal effort.

Time-Spectral Implementation within the OVERFLOW Solver

The Time-Spectral temporal discretization has been incorporated into NASA's implicit Reynolds-averaged Navier-Stokes (RANS) solver OVERFLOW, requiring a number of modifications to the existing source code:

- Both standard and hybrid Time-Spectral discretizations have been implemented within OVERFLOW, leveraging the well-established codebase and the ability to simulate complex, three-dimensional flows.
- The Time-Spectral discretization was incorporated into the implicit approximate-factorization (AF) scheme. The diagonalizability of the dense implicit temporal operator was exploited to reduce the $\mathcal{O}(N^3)$ complexity associated with a direct linear solve to the $\mathcal{O}(N \log N)$ complexity of the FFT and IFFT.
- Time-Spectral turbulence capabilities were incorporated within OVERFLOW's loosely-coupled scheme by applying a retroactive semi-implicit operator. A temporal gradient limiter was employed to maintain the positivity constraint of the one-equation Spalart-Allmaras model.
- OVERFLOW's spatial multigrid algorithm was extended to include the implicit Time-Spectral discretization on all grid levels, facilitating significantly-accelerated space-time convergence.
- Temporal vanishing viscosity was incorporated to mitigate the destabilizing effect of aliasing errors. The well-established modal dealiasing technique of spectral vanishing viscosity (SVV) was implemented for statically-blanked (Fourier-based) nodes. A nodal temporal dissipation operator derived from barycentric rational interpolants was applied to dynamically-blanked nodes at the partitioned interval boundaries.

Validation of the Time-Spectral OVERFLOW Solver

Validation was facilitated by direct comparisons between solutions computed with the augmented Time-Spectral OVERFLOW solver and the existing OVERFLOW time-accurate solver. High resolution time-accurate simulations were therefore taken as the exact solution and used to gauge the accuracy of the Time-Spectral calculations. The two-dimensional test cases were purposefully chosen as single-component configurations to enable both rigid and relative motion in order to compare the performance of the hybrid Time-Spectral scheme to the standard Time-Spectral method.

- Two-dimensional validation of the augmented OVERFLOW Time-Spectral solver was executed on a series of two-dimensional model problems for a number of relevant flow regimes; inviscid, laminar, transonic and turbulent test cases were included.
- The OVERFLOW Time-Spectral solver performed successful three-dimensional simulations of the isolated quarter-scale V-22 Tilt Rotor Aeroacoustic Model (TRAM) in hover and forward (edgewise) flight.

1.3 Dissertation Outline

The thesis continues in Chapter 2 with a brief derivation of the standard Time-Spectral method. A discussion of the difficulties associated with its application to nonlinear equations is provided and the SVV dealiasing technique is formally introduced. Chapter 3 reviews the overset grid methodology providing context for Chapter 4, which details the difficulty of applying the Time-Spectral method within such a framework for arbitrary relative motion. Chapter 4 then introduces a variety of potential strategies to resolve overset relative motion within the Time-Spectral framework. The governing equations of fluid dynamics are provided in Chapter 5 followed by a description of NASA's existing OVERFLOW solver. The augmented Time-Spectral OVERFLOW solver is then outlined, including the additional memory and computational overhead with respect to the steady-state solver. Various numerical experiments are presented in Chapter 6, beginning with two-dimensional

oscillating airfoils; inviscid, laminar, transonic and turbulent cases are investigated. Finally, computational results from the V-22 TRAM are presented for both hover and forward (edgewise) flight. A brief summary of the dissertation is provided in Chapter 7. A discussion of potential areas of future work concludes the dissertation.

Chapter 2

The Time-Spectral Method

Unsteady calculations have traditionally been performed by marching the governing equations forward in physical time from an initial condition. For a time-periodic problem, the solution, u , must be advanced long enough to eliminate the initial transient and establish a statistically stationary response such that $u(t + T) = u(t)$ for period T ; this may demand simulating dozens or more periods of the flow. Time-marching schemes typically employ an algebraically-accurate finite-difference-based temporal discretization. Due to the limited level of temporal accuracy, calculations typically employ hundreds to thousands of time steps per period, N , to adequately resolve the flow. Therefore, the overall cost of the calculation scales as a product of N with the number of periods required to achieve a periodic steady state. This costly endeavor may prove intractable for realistic high-fidelity three-dimensional unsteady design problems, including rotorcraft and turbomachinery, whose performance is primarily evaluated on the converged periodic steady-state characteristics of the flow.

Recognizing that the converged temporal response of forced periodic flows can be assumed periodic, it is reasonable to apply a periodic constraint a priori. This is accomplished by expressing the temporal variation of the solution at every point in space as a truncated Fourier series, with a fundamental frequency, $\omega = 2\pi/T$, that is equivalent to the forcing frequency. Due to the spectral convergence rate of the Fourier series [45], which is superior to any algebraic order, spectral methods provide significantly greater accuracy for a given number of degrees of freedom over

traditional algebraically-accurate time advancement schemes. Furthermore, the bulk of energy for forced flows is often highly concentrated in the lower modes justifying a truncation of the Fourier series to a relatively small number of terms. The tradeoff for this approach involves storing N states of the solution corresponding to the N evenly-spaced time samples or *collocation* points over the period, in contrast to just one or two states when using a backward difference formula (BDF) in the traditional time-marching approach. A calculation employing N time samples retains $K = (N - 1) / 2$ harmonics of the fundamental frequency.

Expressing the solution as a truncated Fourier series with spatially varying coefficients transforms the unsteady governing equations into a coupled set of steady equations where the combined *space-time* solution converges to its periodic steady state directly. This avoids the unsteady transient in physical time (a pseudotime transient does exist) and the need to solve time-accurately through multiple periods of the flow. Development of these Fourier pseudospectral schemes and their application to the overset framework is outlined in §1.1. Aliasing and a technique to mitigate its destabilizing effects, are introduced in §2.2. There exists a number of ways to derive the Time-Spectral method, but perhaps the most concise approach is its development as a Fourier collocation scheme in time [24, 37].

2.1 Fourier Collocation in Time

Spectral or pseudospectral methods have generally been applied to time-dependent problems where the spatial variation is expanded in a global basis set and the temporal discretization is treated by finite-differences or another algebraically-accurate approximation. Applying the routine combination of finite-differences and spectrally-accurate global differentiation operators leads directly to the development of the Time-Spectral method by exchanging the independent variables of space and time. The Time-Spectral method combines a spectrally-accurate global temporal differentiation operator with algebraically-accurate spatial discretizations such as finite differences, finite-volumes or finite-elements, that can be applied to spatial domains of arbitrary complexity.

To develop the Time-Spectral method, consider a general time-dependent scalar partial differential equation (PDE) where the solution, u , is assumed periodic in time with period T .

$$\frac{\partial}{\partial t} u(\mathbf{x}, t) + \mathcal{R}(u(\mathbf{x}, t)) = 0, \quad \mathbf{x} \in \Omega \quad (2.1)$$

The general residual operator, $\mathcal{R}(u)$, is free to be of linear or nonlinear form. Since the solution is assumed periodic in time, such that $u(\mathbf{x}, t + T) = u(\mathbf{x}, t)$, the solution at every point in space, \mathbf{x} , can be expressed as a Fourier series where the basis functions, $\phi_k(t)$, are the complex exponentials in integer harmonics of the fundamental frequency, $\omega = 2\pi/T$.

$$u(\mathbf{x}, t) = \sum_{k=-\infty}^{\infty} \hat{u}_k(\mathbf{x}) \phi_k(t), \quad \phi_k(t) = e^{i\omega kt} \quad (2.2)$$

The spatially-varying Fourier coefficients, $\hat{u}_k(\mathbf{x})$, are determined by the *continuous* Fourier transform (CFT).

$$\hat{u}_k(\mathbf{x}) = \frac{1}{T} \int_0^T u(\mathbf{x}, t) e^{-ik\omega t} dt \quad (2.3)$$

To solve Eq. 2.1 discretely, only a finite number of temporal modes, K , can be resolved, therefore requiring a truncation of Eq. 2.2 to $N = 2K + 1$ terms. Because of the spectral convergence properties of the Fourier series and the fact that the bulk of energy is often contained in the lower-frequency modes, truncation to a relatively small number of harmonics often resolves the meaningful content of the underlying solution.

The original Harmonic Balance method [2], and other spectral methods, solve for the Fourier coefficients, \hat{u}_k , in the transformed space where the solution is projected via P_N onto a subspace spanned by only the lowest K modes.

$$u_N(\mathbf{x}, t) \approx P_N u(\mathbf{x}, t) = \sum_{k=-K}^K \hat{u}_k(\mathbf{x}) \phi_k(t) \quad (2.4)$$

The sole error in the projection results from the truncation, $(I - P_N)u$. While this approach avoids the introduction of aliasing error, analytic evaluation of the CFT can become intractable with increasing modes or complexity of the PDE. Additionally, the residual routines for existing PDE solvers cannot be leveraged directly, therefore requiring significant upfront investment in software development.

The Time-Spectral method, in contrast to the original Harmonic Balance approach, is a time-domain method that takes as its unknowns the discrete values of the solution, $u_N(\mathbf{x}, t)$, at specific time instances throughout the period. This greatly simplifies its integration within existing solvers because the spatial operator, $\mathcal{R}(u_N)$, can be leveraged with minimal, if any, modifications. The time-domain formulation is a *collocation* or *pseudospectral* approach that seeks to eliminate the discrete residual at a set of predetermined collocation points, t_j , corresponding to the selected basis functions. For periodic phenomena, employing the complex exponential basis functions is formally referred to as *Fourier collocation*. A distribution of N evenly-spaced collocation points is defined over the period.

$$t_j = j\frac{T}{N} = j\Delta t, \quad j \in J_N, \quad J_N = \{0, \dots, N-1\} \quad (2.5)$$

The method of weighted residuals (MWR) is used to discretely approximate solutions to PDEs by minimizing a discrete residual, R_N , in some manner. The partial derivative with respect to time is now expressed as a full differential with respect to the discretized form of Eq. 2.1.

$$R_N(\mathbf{x}, t) = \frac{d}{dt}u_N(\mathbf{x}, t) + \mathcal{R}(u_N(\mathbf{x}, t)), \quad \mathbf{x} \in \Omega \quad (2.6)$$

The residual is eliminated by setting to zero the following inner product taken over the time period [46].

$$(R_N, \psi_j)_w = \int_0^T R_N \psi_j w dt = 0, \quad j \in J_N \quad (2.7)$$

The choice of weight and test functions, $w = 1$ and $\psi_j = \delta(t - t_j)$, respectively, defines a collocation scheme. Integrating Eq. 2.7 for each $j \in J_N$ results in a system

of N equations in N unknowns which simply sets to zero the discrete residual at each of the temporal collocation points, t_j .

$$R_N(\mathbf{x}, t_j) = 0, \quad \mathbf{x} \in \Omega, \quad j \in J_N \quad (2.8)$$

This can be expressed in terms of the original PDE in the following manner.

$$\frac{d}{dt}u_N(\mathbf{x}, t_j) + \mathcal{R}(u_N(\mathbf{x}, t_j)) = 0, \quad \mathbf{x} \in \Omega, \quad j \in J_N. \quad (2.9)$$

The time-derivative in Eq. 2.9 is evaluated by analytically differentiating the expansion of u_N in the complex exponential basis. However, the unknown expansion coefficients, \tilde{u}_k , in the pseudospectral method are defined by trigonometric interpolation, I_N , as opposed to trigonometric projection, P_N , because the unknowns are the time-domain variables as opposed to the continuous Fourier coefficients.

$$u_N(\mathbf{x}, t) \approx I_N u(\mathbf{x}, t) = \sum_{k=-K}^K \tilde{u}_k(\mathbf{x}) \phi_k(t) \quad (2.10)$$

The discrete Fourier coefficients $\tilde{u}_k(\mathbf{x})$ are determined from the solution values at the N collocation points by the *discrete* Fourier transform (DFT).

$$\tilde{u}_k(\mathbf{x}) = \frac{1}{N} \sum_{j=0}^{N-1} u_N(\mathbf{x}, t_j) e^{-i\omega k t_j} \quad (2.11)$$

Constructing a collocation approximation that sets to zero the residual $u - u_N$ at the collocation points uniquely determines the coefficients \tilde{u}_k by requiring that the approximation $I_N u$ interpolates u_N at the collocation points.

$$I_N u(\mathbf{x}, t_j) = u_N(\mathbf{x}, t_j) = \sum_{k=-K}^K \tilde{u}_k(\mathbf{x}) \phi_k(t_j), \quad j \in J_N \quad (2.12)$$

This results in an invertible system of N equations in N unknowns that can be

expressed in matrix form.

$$\mathbf{u}_N(\mathbf{x}) = \Phi \tilde{\mathbf{u}}(\mathbf{x}) \quad (2.13)$$

The inverse discrete Fourier transform operator, $\Phi_{jk} = \phi_k(t_j)$, maps the frequency domain into the time domain at the specified collocation points as a matrix-vector product. The vectors, $\mathbf{u}_N(\mathbf{x})$ and $\tilde{\mathbf{u}}(\mathbf{x})$, contain the time- and frequency-domain components of the solution, respectively, at a particular point in space, \mathbf{x} .

$$\begin{aligned} \mathbf{u}_N(\mathbf{x}) &= \{u_N(\mathbf{x}, t_0), \dots, u_N(\mathbf{x}, t_{N-1})\}^T \\ \tilde{\mathbf{u}}(\mathbf{x}) &= \{\tilde{u}_{-K}(\mathbf{x}), \dots, \tilde{u}_K(\mathbf{x})\}^T \end{aligned}$$

Because Eq. 2.13 is uniquely determined, and therefore invertible, the DFT is defined by Φ^{-1} .

$$\tilde{\mathbf{u}}(\mathbf{x}) = \Phi^{-1} \mathbf{u}_N(\mathbf{x}) \quad (2.14)$$

With the discrete Fourier coefficients uniquely determined, the time-derivative in Eq. 2.9 is defined by analytically differentiating Eq. 2.10.

$$\frac{d}{dt} u_N(\mathbf{x}, t) = \sum_{k=-K}^K \tilde{u}_k(\mathbf{x}) \frac{d}{dt} \phi_k(t) = \sum_{k=-K}^K i\omega k \tilde{u}_k(\mathbf{x}) e^{i\omega k t} \quad (2.15)$$

Substituting the definition of \tilde{u}_k from Eq. 2.11 into Eq. 2.15 defines the temporal derivative at a particular time sample, t_j , as a weighted sum of the solution at every

time sample.

$$\frac{d}{dt}u_N(\mathbf{x}, t_j) = \sum_{k=-K}^K i\omega k \tilde{u}_k(\mathbf{x}) e^{i\omega k t_j} \quad (2.16)$$

$$= \sum_{n=0}^{N-1} \left(\frac{1}{N} \sum_{k=-K}^K i\omega k e^{i\omega k(t_j - t_n)} \right) u_N(\mathbf{x}, t_n) \quad (2.17)$$

$$= \sum_{n=0}^{N-1} \partial_{jn} u_N(\mathbf{x}, t_n) \quad (2.18)$$

Defining the time derivative at each collocation point and collecting the expressions into matrix form arrives at the definition of the *Fourier interpolation* derivative operator [47].

$$\frac{d}{dt}\mathbf{u}_N(\mathbf{x}) = \mathcal{D}_N \mathbf{u}_N(\mathbf{x}) \quad (2.19)$$

The elements of \mathcal{D}_N correspond to the ∂_{jn} in Eq. 2.18. Direct assembly of this operator is provided in [22] and included below for N odd.

$$\partial_{jn} = \begin{cases} \frac{2\pi}{T} \frac{1}{2} (-1)^{j-n} \csc\left(\frac{\pi(j-n)}{N}\right), & j \neq n \\ 0, & j = n \end{cases} \quad (2.20)$$

This operator embeds three operations; the solution at the collocation points is transformed into the frequency domain, the frequency-domain representation is then differentiated by scalar multiplication and the result is transformed back into the time domain. This is perhaps more evident by writing equation 2.15 in matrix form.

$$\frac{d}{dt}\mathbf{u}_N(\mathbf{x}) = \Phi D_N \tilde{\mathbf{u}}(\mathbf{x}) \quad (2.21)$$

The diagonal differentiation matrix, D_N , with elements $d_{kk} = i\omega k$, operates in the transformed space on the vector of discrete Fourier coefficients. Substitution of Eq.

2.14 uncovers the underlying structure of $\mathcal{D}_N = \Phi D_N \Phi^{-1}$.

$$\frac{d}{dt} \mathbf{u}_N(\mathbf{x}) = \Phi D_N \Phi^{-1} \mathbf{u}_N(\mathbf{x}) = \mathcal{D}_N \mathbf{u}_N(\mathbf{x}) \quad (2.22)$$

The j -th component the temporal derivative approximation is inserted into Eq. 2.9 resulting in the Time-Spectral system of equations.

$$\sum_{n=0}^{N-1} \partial_{jn} u_N(\mathbf{x}, t_n) + \mathcal{R}(u_N(\mathbf{x}, t_j)) = 0, \quad \mathbf{x} \in \Omega, \quad j \in J_N \quad (2.23)$$

Equation 2.23 demonstrates how the Time-Spectral discretization results in a steady system of equations at the N collocation points coupled through the time-derivative term. The matrix form is used herein.

$$\mathcal{D}_N \mathbf{u}_N(\mathbf{x}) + \mathcal{R}(\mathbf{u}_N(\mathbf{x})) = \mathbf{0}, \quad \mathbf{x} \in \Omega \quad (2.24)$$

Equation 2.24 can be solved iteratively by integration in pseudotime until satisfactory convergence is achieved.

$$\frac{d}{d\tau} \mathbf{u}_N(\mathbf{x}) + \mathcal{D}_N \mathbf{u}_N(\mathbf{x}) + \mathcal{R}(\mathbf{u}_N(\mathbf{x})) = \mathbf{0}, \quad \mathbf{x} \in \Omega \quad (2.25)$$

Discretization of the spatial domain into N_x nodes (e.g. for a finite-difference approximation) results in an $(N \times N_x) \times (N \times N_x)$ system of equations that describe the complete space-time solution at all of the spatial nodes at each temporal collocation point. The Time-Spectral discretization results in a steady space-time system, suggesting that convergence can be aided by the powerful acceleration techniques afforded to steady solvers. Additionally, the steady nature of Eq. 2.25 has serious implications for design as a steady adjoint-based design approach [24, 15, 28, 29] can be used to optimize applications involving periodic flows. Finally, a continuous representation of the solution can be reconstructed by evaluating the DFT (Eq. 2.11) of the converged discrete solution and interpolating in time with Eq. 2.10.

In this thesis, the Time-Spectral discretization is incorporated within OVERFLOW, an implicit Navier-Stokes flow solver developed over the last quarter century

at NASA. OVERFLOW employs an approximately-factored scheme that integrates the implicit system of equations using a sequence of direct linear solves numbering the spatial dimensions of the simulation. The Time-Spectral discretization is incorporated into the approximate factorization by treating time like an additional spatial dimension, obviating the need to construct and solve the global system of equations. This results in a procedure that performs the computational tasks associated with a steady-state flow solve at the N temporal collocation points with an additional matrix-vector product and linear solve in the temporal dimension at every spatial node. As is the case with the spatial operator, $\mathcal{R}(u_N)$, the spatial boundary conditions, both physical and farfield, require no modifications (other than accounting for grid velocities in both cases) because they are applied to the solution at distinct instances in physical time. Implementation details of the approximately-factored Time-Spectral scheme are presented in Chapter 5 for the Navier-Stokes equations employed within the augmented OVERFLOW solver.

2.2 Aliasing & Spectral Vanishing Viscosity

In the previous section, a subtle distinction was drawn between the *continuous* and *discrete* Fourier coefficients, \hat{u}_k and \tilde{u}_k , respectively. The distinction is significant with respect to the error induced by Fourier projection, P_N , versus Fourier interpolation, I_N . Because the projection operator is computed in a continuous sense, by evaluating the integral in Eq. 2.3, only the frequencies included in the truncation contribute to the value of each coefficient, \hat{u}_k . However, in the case of the interpolation, the coefficients are evaluated by applying a discrete Fourier transform to time-domain data located at $N = 2K + 1$ equispaced time samples. Discrete sampling of time-domain data can admit information from frequencies higher than can be resolved on the discrete mesh, that are in turn aliased into the discrete Fourier coefficients, \tilde{u}_k .

Canuto et al. [48] offer a concise and intuitive description of the aliasing error, $A_N u$, induced by Fourier interpolation. This paragraph summarizes some of their key assertions. They show that substitution of Eq. 2.2 into Eq. 2.11 results in the following expression for the discrete Fourier coefficients (assuming convergence of Eq.

2.2 at the collocation points).

$$\tilde{u}_k = \hat{u}_k + \sum_{\substack{m=-\infty \\ m \neq 0}}^{\infty} \hat{u}_{k+Nm}, \quad k = -K, \dots, K \quad (2.26)$$

Equation 2.26 demonstrates the contribution of higher-frequency coefficients into the lower-frequency discrete coefficients; the interpolation error of a discrete signal is equivalent to the sum of its projection and aliasing error.

$$I_N u = P_N u + A_N u$$

Because the aliasing error,

$$A_N u = \sum_{k=-K}^K \left(\sum_{\substack{m=-\infty \\ m \neq 0}}^{\infty} \hat{u}_{k+Nm} \right) \phi_k,$$

is orthogonal to the truncation error (projection), the interpolation error is always larger than the projection error.

$$\|u - I_N u\|^2 = \|u - P_N u\|^2 + \|A_N u\|^2$$

Aliasing occurs because the complex exponentials of aliased frequencies evaluate to the same value at the N collocation points, $\phi_{k+Nm}(t_j) = \phi_k(t_j)$. This is further demonstrated in Fig. 2.1 for sine waves with frequencies $k = -10, -2$ and 6 with $N = 8$ nodes, which is taken from their discussion. Please refer to Section 2.1 in [48] for a more complete discussion and mathematical treatment of the matter.

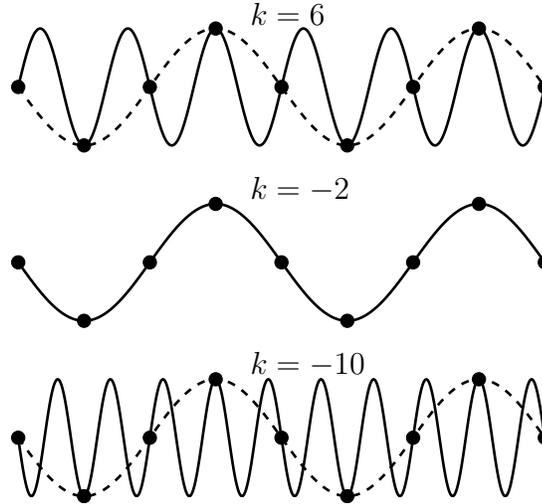


Figure 2.1: Aliasing of sine waves. The trigonometric interpolant of the $k = 6$ and $k = -10$ harmonic sine waves alias down to the $k = -2$ sine wave on the set of $N = 8$ samples, demonstrating the ability for high-frequency waves to corrupt a signal on a grid too coarse to resolve them. The $(k + Nm)$ -th harmonic is *aliased* to the k -th harmonic on a discrete grid of N nodes. Figure adapted from [47] and [48].

Despite the additional aliasing errors introduced by trigonometric *interpolation*, Snider [45] demonstrates that the discrete Fourier series maintains the spectral convergence rate for smooth functions. Nevertheless, aliasing can in fact impede solution convergence or destabilize calculations. Accordingly, several methods have been developed to mitigate aliasing for pseudospectral methods.

Several anti-aliasing techniques have emerged in the literature. Orszag [49] examined aliasing error for quadratic functions, and demonstrated that an oversampling approach wholly removes aliasing errors provided that $M = 3K + 1$ frequencies are included. The highest third of the spectrum is filtered out at each iteration, enabling the calculation to proceed absent aliasing. Of course, extra storage and computational effort is required, but this approach affords the user the ability to employ an interpolation scheme without investing the required development and computational costs associated with a projection scheme. However, this approach is only theoretically alias free for equations that are quadratic in the solution variables.

Another approach proposed by Tadmor [50], entitled *spectral vanishing viscosity* (SVV), applies an artificial viscosity term in the form of a second-degree global differentiation operator that is only applied above a certain cutoff frequency. This approach damps the high-frequency modes, attenuating their corrupting impact on the resolved spectrum. Maday and Tadmor [51] applied SVV to the periodic case where the basis functions are naturally taken to be the complex exponentials. Applying this concept to the original PDE results in the dealiased form of the equations.

$$\frac{\partial}{\partial t} u_N(\mathbf{x}, t) + \mathcal{R}(u_N(\mathbf{x}, t)) = \epsilon \frac{\partial}{\partial t} \left(\mathcal{Q}_N \frac{\partial}{\partial t} u_N(\mathbf{x}, t) \right), \quad \mathbf{x} \in \Omega \quad (2.27)$$

The *symmetric viscosity kernel*, \mathcal{Q}_N , is applied only to the higher frequencies by defining a smooth function in the frequency domain that is null below the cutoff frequency, M , and smoothly transitions to unity at the highest frequencies. The elements of the diagonal frequency-domain kernel, $\tilde{\mathcal{Q}}_N^M$, provided by Maday et al. [52] are defined below as

$$\tilde{\mathcal{Q}}_N^M(k) = e^{-\frac{(k-N)^2}{(k-M)^2}}, \quad k > M. \quad (2.28)$$

Selection of the coefficient, ϵ , cutoff frequency and viscosity kernel must be determined; here, we used $\epsilon = \Delta t/K$. The kernel is essentially a modal scaling of the second derivative operator that can be incorporated into the global differentiation operator. This approach avoids the extra computation required by the frequency-domain filtering process of oversampling and truncation. The default approach is to employ the *inviscid*, spectrally-accurate temporal differentiation operator, \mathcal{D}_N , and only apply SVV upon the realization of convergence deficiencies. If dealiasing is required, the SVV operator is embedded within the standard differentiation operator in the following manner.

$$\mathcal{D}_N \leftarrow \mathcal{D}_N - \frac{\Delta t}{K} \mathcal{D}_N \mathcal{Q}_N^M \mathcal{D}_N \quad (2.29)$$

The cutoff frequency, M , is a free parameter but taken most often as roughly two-thirds of the largest harmonic; other choices have been tested in the present research

for both the cutoff frequency and the kernel definition.

McMullen [11] employed this concept on the coarse grids in the multigrid framework to more rapidly damp out the high frequency modes without sacrificing accuracy on the finest-level mesh. This approach accelerated the convergence rate of a K -mode calculation to the optimal convergence rate of a steady-state, zero-mode, calculation. A detailed comparison between the classical HB and HDHB methods by Liu et al. [53] exposed the corrupting characteristics of Fourier interpolation. Labryer and Attar [54] investigated frequency- and finite-difference-based filtering to mitigate aliasing issues in HDHB calculations. Huang and Ekici [55] applied spectral vanishing viscosity to dealias HDHB calculations using a binary viscosity kernel. SVV is found to stabilize otherwise unstable or non-convergent calculations in Chapter 6 for the case of a high-frequency laminar plunging NACA 0012 airfoil (§6.1.2) and for the V-22 tiltrotor in forward flight (§6.2.2).

Chapter 3

Overset Grid Methodology

In the early stages of computational fluid dynamics (CFD), the availability of modest computational resources limited calculations to relatively simplistic geometries embedded within structured meshes [56, 57]. Structured-mesh-based flow solvers were quite efficient, as information at neighboring grid points was accessible directly at a known offset. Development of more powerful CFD tools, as well as the machines running them, resulted in both increased problem sizes and complexity. Structured grids were no longer easily applied to the increasingly-detailed geometries. Unstructured grids emerged as a viable candidate to successfully mesh geometries of arbitrary complexity, but forfeited the efficient data layout of the structured approach. Therefore, unstructured algorithms do not generally match the efficiency of their structured counterparts.

In the early 1980s, Benek, Steger and Dougherty [58] formally introduced the chimera grid approach – a system of overlapping structured grids. This novel concept facilitated the use of structured algorithms on complex geometries. This technique was extended by others [59, 60, 61, 62]. For example, a curvilinear structured grid could be used to discretize the span of the wing, and another grid could be used to discretize the tip region, contingent upon sufficient boundary communication between the two. Cartesian grids offer an additional level of algorithmic efficiency, and standard practice evolved towards meshing the body with one or more structured curvilinear grids embedded within a system of one or more Cartesian background

grids spanning the farfield domain. This chapter is by no means a comprehensive survey of the development and implementation of overset grid technology. Rather, a brief introduction of the fundamental concepts is provided to contextualize the extension of the standard Time-Spectral method for its accommodation within the overset grid framework in a consistent and general manner.

3.1 Overview

A representative two-dimensional overset grid system for a NACA 0012 airfoil calculation is presented in Fig. 3.1. In Fig. 3.1a, a curvilinear body-fitted grid is embedded within a background Cartesian grid. In Fig. 3.1b, mesh points interior to the airfoil are trimmed (*cut*) because the equations governing fluid dynamics do not apply. Additional points that lie outside the solid body are also trimmed to better match grid spacing in the overlap region, thereby increasing the size of the *hole cut*. Overset grids, while separate logical entities, must be linked to enable communication across their boundary interfaces. This requires overlap regions whose size depends on the stencil of the numerical scheme being employed. The overlap region is specified either by the user or by an automated domain connectivity tool, with the goal of providing commensurate spacing of the grids in the overlap region. Fringe points (the number of which are again determined by the stencil width) are identified on the boundary of the overlap region and tagged as receiver nodes, as seen in Fig. 3.1c. The domain connectivity software interpolates the solution from the donor grid in the vicinity of each fringe point in order to define the solution on boundary interfaces. This is repeated on all grids ensuring that every boundary node is either populated with interpolated data from a neighboring zone or is defined through a physical or farfield boundary condition. Mesh points that are ignored, because of their location within an impermeable boundary, or because they are trimmed in favor of a more accurate, finer grid, are considered *blanked*.

In a moving body scenario, grids can either move rigidly or relative to one another. For the case of rigid motion, where all grids move in unison, the blanking status for all grid points remains constant in time. However, for the case of relative motion,

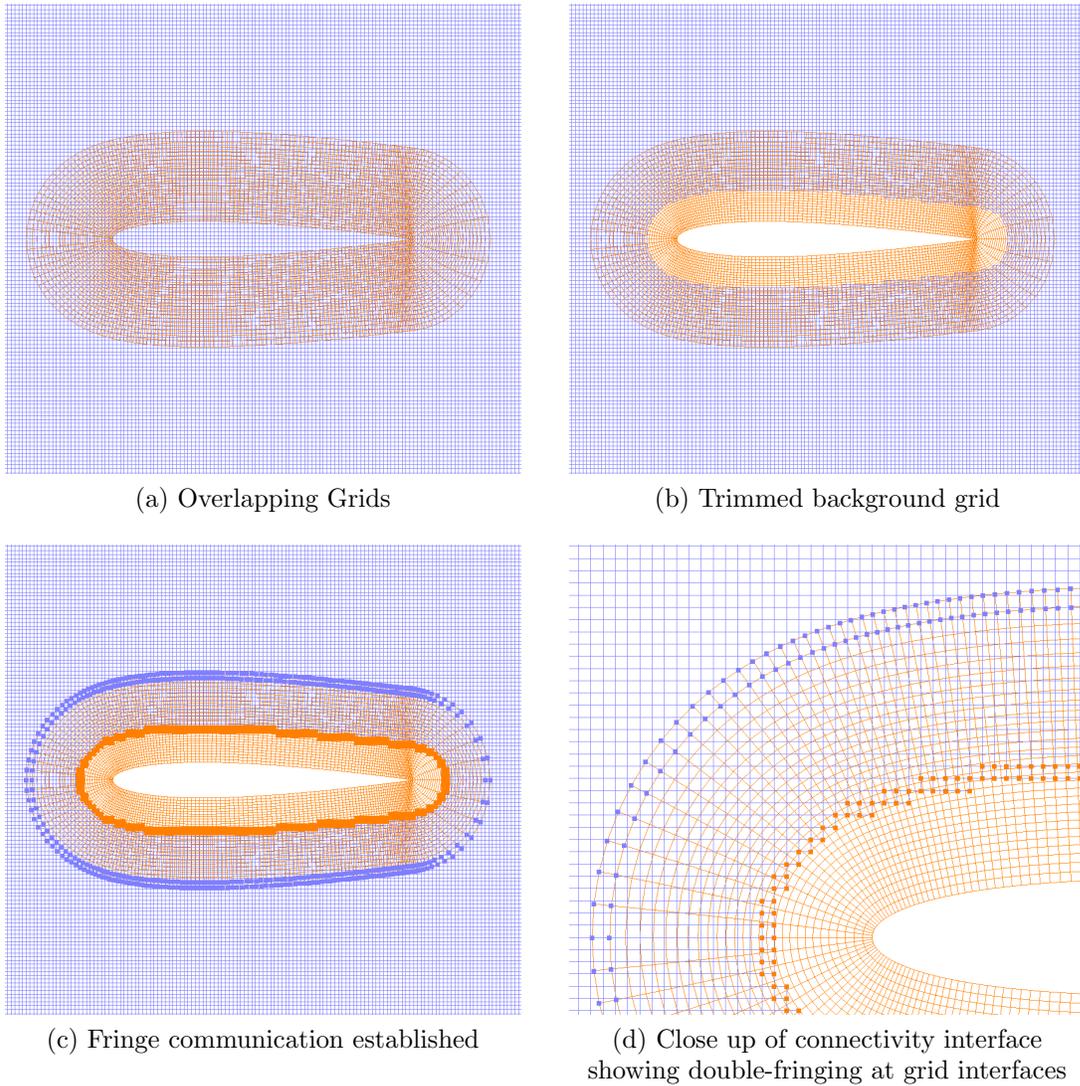


Figure 3.1: Two-dimensional overlapping grid-system of a NACA 0012 airfoil. (a) A curvilinear body-fitted *near-body* grid wraps the airfoil surface geometry and is embedded within a background Cartesian *off-body* grid. (b) The Cartesian grid is trimmed back from the surface of the geometry to remove points interior to the airfoil and to ensure commensurate grid-spacing in the overlap region. (c) Fringe points are established on grid interfaces to facilitate communication. (d) Magnified view of the grid interfaces. Double fringing is used enable a stencil width of 5 mesh points.

where grids move independently, the blanking status may change over time. Consider the case of a plunging airfoil. A curvilinear near-body grid is used to resolve the geometry of the airfoil and is embedded within a background Cartesian grid spanning the farfield domain. Figure 3.2 is used to examine the blanking status for the background Cartesian grid for both rigid and relative motion. The background grid is depicted at three time instances, corresponding to the three equispaced time samples associated with a Time-Spectral calculation using $N = 3$. The blue-colored nodes are unblanked for every time instance and the white region corresponds to the hole cut at the particular time instance. For the case of rigid-body motion, depicted in Fig. 3.2a, the hole cut remains constant. Every node that is unblanked at one time instance is unblanked for all three. This is not true for the case of relative motion depicted in Fig. 3.2b. Because the airfoil grid moves relative to the stationary background grid, the hole cut trims a different set of nodes at each time instance; some nodes are blanked at one time instance but unblanked at another time instance. These *dynamically-blanked* nodes are colored red in the figure and only contain physically meaningful solution data at a subset of the time samples. Herein lies the fundamental difficulty of applying the Time-Spectral approach to an overset solver; complete time histories are not available for every grid point and therefore the infinitely-supported, spectrally-accurate Fourier differentiation operator cannot be employed. An alternative treatment is sought for such nodes, which is the focus of the next chapter.

3.2 Representative One-Dimensional Grid System

The simplified one-dimensional, two-grid overset grid-system depicted in Fig. 3.3 is used to demonstrate a representative overset discretization. The integer `IBLANK` array assigns a value of either 1, 0 or -1 to every spatial node at a given time. Standard compute nodes that are not located on the physical or overset boundaries are defined as `IBLANK = 1`. Trimmed nodes are prescribed as `IBLANK = 0`, which implies that their information will not be accessed nor their solution updated. Nodes located on overset interfaces that require interpolated data from surrounding grids are assigned `IBLANK = -1`. This value is also assigned to nodes on physical or farfield boundaries.

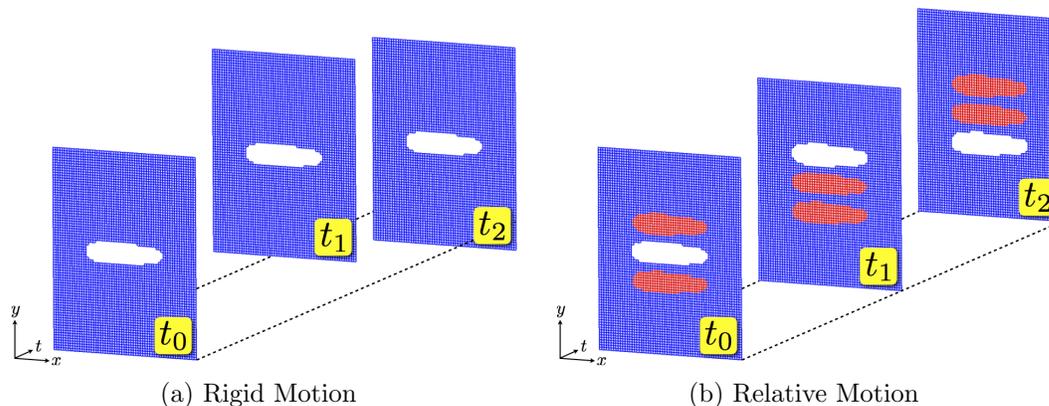


Figure 3.2: Hole cut at three time instances for a plunging airfoil for both rigid and relative motion. The time instances are defined by $t_j = jT/3$. For (a) rigid motion, the hole cut remains constant on the background grid. However, for (b) relative motion, the hole cut moves dynamically resulting in spatial nodes (in red) that are only defined for a subset of the global number of time samples and therefore lack complete time histories.

These points will not be updated, but can be accessed as part of neighboring stencils. The algorithms are written in a general way by incorporating the IBLANK array into the semi-discretization such that trimmed nodes are ignored and overset boundaries are handled appropriately.

Grid A in red, with grid-spacing Δx_A , has 12 nodes. The more coarsely-spaced grid B in blue, with grid-spacing $\Delta x_B > \Delta x_A$, also has 12 nodes. Presumably, grid A is embedded within grid B to resolve some feature of the solution in its vicinity. However, extending the finely-spaced grid A out to the farfield boundaries proves inefficient because its higher resolution is unnecessary away from the region containing the flow feature. Grid B is employed to efficiently extend the domain from the boundaries of grid A to the farfield boundaries and spans the entire width of the domain. Due to the superior resolution of grid A in the vicinity of its interior, nearby mesh points on grid B are blanked to avoid wasteful computation on the under-resolved mesh (these points would also be blanked if a solid body was placed in the location of grid A). The discretized governing equations are solved independently for each grid. After the solution is updated on all grids for a given iteration, the

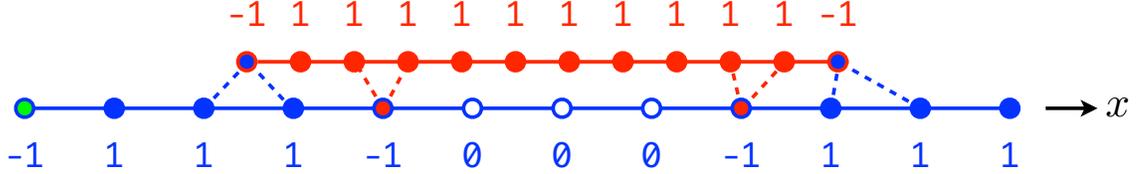


Figure 3.3: Representative two-grid, one-dimensional overlapping grid-system. Finely-spaced grid A is embedded within the coarse grid B . Integer IBLANK values are specified for each node. Node outlines indicate the associated grid (A in red and B in blue). Node interiors indicate blanking and boundary condition information; green indicates a physical boundary condition, white indicates a blanked node, and node interiors with a color corresponding to the other grid indicates a receiver node (with IBLANK = -1). The dotted lines indicate the donor nodes that provide interpolation information.

solution is interpolated to the receiver nodes by the connectivity routines within the flow solver.

To illustrate how the IBLANK array is incorporated into the discretization, consider the continuous linear advection equation in strong conservation-law form with advective flux, $f = au$.

$$\frac{\partial u}{\partial t} + \frac{\partial f}{\partial x} = 0 \quad (3.1)$$

The semi-discretization at node j employs a generic spatial difference operator, δ_x ,

$$\frac{du_j}{dt} + \delta_x f_j = 0 \quad (3.2)$$

For example, δ_x^{c2} refers to a second-order central-difference and δ_x^{b1} refers to a first-order backward-difference. A complete discretization employing a second-order central-difference operator acting on the advective flux and an implicit Euler time-advancement scheme results in the following system of equations, where $\Delta u_j = u_j^{n+1} - u_j^n$.

$$\frac{\Delta u_j}{\Delta t} + \delta_x^{c2} f_j^{n+1} + \mathcal{O}(\Delta x^2) = 0 \quad (3.3)$$

Dropping the error term of order Δx^2 results in the following expression

$$\frac{\Delta u_j}{\Delta t} + \frac{1}{2\Delta x} (f_{j+1}^{n+1} - f_{j-1}^{n+1}) = 0, \quad (3.4)$$

which becomes

$$\frac{\Delta u_j}{\Delta t} + \frac{a}{2\Delta x} (u_{j+1}^{n+1} - u_{j-1}^{n+1}) = 0. \quad (3.5)$$

Subtracting the spatial difference at iteration n , $\frac{a}{2\Delta x} (u_{j+1}^n - u_{j-1}^n)$, from both sides results in the *delta* form of the equations.

$$\Delta u_j + \frac{a\Delta t}{2\Delta x} (\Delta u_{j+1} - \Delta u_{j-1}) = -\frac{a\Delta t}{2\Delta x} (u_{j+1}^n - u_{j-1}^n) \quad (3.6)$$

Employing a second-order central-difference operator, with a stencil width of three points, requires the use of a single fringe-point on the overset boundaries of each grid. Therefore, nodes 2 through 11 are designated as compute nodes on grid *A* and prescribed with `IBLANK = 1`. Nodes 1 and 12 are designated as receiver nodes that require interpolated data from grid *B* and are prescribed accordingly with `IBLANK = -1`. Nodes 6 through 8 are blanked on grid *B* because grid *A* can more accurately resolve the solution in the region spanned by these nodes. Fringe points that receive data from grid *A* are positioned on either side of the blanked region so that blanked nodes are not included in the stencil of the spatial difference operator on the compute nodes of grid *B*. Nodes 2 through 4 and 10 through 12 are designated as compute nodes on grid *B* and therefore prescribed with `IBLANK = 1`. Nodes 5 and 9 are designated as receiver nodes requiring interpolated data from grid *A* and are therefore prescribed with `IBLANK = -1`. Node 1 on grid *B* requires a physical boundary condition specifying the inflow value of the solution and is therefore prescribed with `IBLANK -1` to indicate that it is not a compute node but contains physically meaningful data that can be accessed by its neighbor's stencil.

If node j is either blanked or a fringe point, then it is not considered a compute node and therefore lies outside the computational domain (`IBLANK =0`) or on the grid boundary (`IBLANK = -1`), respectively. The variable, ib_j is defined as the `IBLANK` value

of node j . Thus, the right-hand-side term, $\mathcal{R}(u_j)$, corresponding to node j is premultiplied by the term $\max(0, ib_j)$ to nullify its contribution to the system of equations. The left-hand-side term, $\mathcal{L}(u_j)$, must also be evaluated such that the $\Delta u_j = 0$ is prescribed for all blanked or fringe nodes; this ensures that the solution is not updated nor does it influence the other equations in the implicit system. Thus, the left-hand side and right-hand-side, respectively, of the equation corresponding to node j can be expressed as follows:

$$\mathcal{L}(u_j) = \Delta u_j + \frac{a\Delta t}{2\Delta x} (\Delta u_{j+1} - \Delta u_{j-1}) \times \max(0, ib_j) \quad (3.7)$$

$$\mathcal{R}(u_j) = -\frac{a\Delta t}{2\Delta x} (u_{j+1}^n - u_{j-1}^n) \times \max(0, ib_j) \quad (3.8)$$

Thus, the equation corresponding to compute node j , where $ib_j = 1$, collapses to

$$\frac{a\Delta t}{2\Delta x} \Delta u_{j+1} + \Delta u_j - \frac{a\Delta t}{2\Delta x} \Delta u_{j-1} = -\frac{a\Delta t}{2\Delta x} (u_{j+1}^n - u_{j-1}^n) \quad (3.9)$$

and the equation corresponding to a trimmed or boundary node j , where $\max(ib_j, 0) = 0$, is not updated.

$$\Delta u_j = 0 \quad (3.10)$$

The second-order central-difference operator can be applied to all but one node. A first-order backward-difference operator is applied at the outflow boundary (node 12 on grid B).

$$\delta_x^{b_1} f_{12}^{n+1} = \frac{a}{\Delta x_B} (u_{12}^{n+1} - u_{11}^{n+1}) + \mathcal{O}(\Delta x) \quad (3.11)$$

The system of equations for grid A resembles the typical case with boundary conditions applied at both the inflow and outflow; the boundary data are provided to the fringe points via the interpolated solution from grid B . Specifically, the boundary solution for node 1 is prescribed by a linear interpolation between the solutions at nodes 3 and 4 on grid B and the boundary solution at node 12 is prescribed by a linear interpolation between the solutions at nodes 10 and 11 on grid B . For grid B ,

the solution at node 5 is prescribed by a linear interpolation between the solutions at nodes 3 and 4 on grid A and the solution at node 9 is prescribed by a linear interpolation between the solutions at nodes 10 and 11 on grid A .

The 12×12 systems of equations for grid A and B are provided in Eq. 3.13 and Eq. 3.14, respectively. The following definitions are provided for compactness.

$$\alpha = \frac{a\Delta t}{2\Delta x_A}, \quad \beta = \frac{a\Delta t}{2\Delta x_B} \quad (3.12)$$

The matrix system for grid B is segmented into two disparate compute regions that are separated by a section of diagonal values of unity. Thus, the large system of equations can be decoupled into two independent linear systems that do not communicate directly and can therefore be solved independently (they are coupled only through their communication with grid A in the connectivity phase).

More sophisticated routines handle higher-order and higher-degree derivatives to ensure the correct operators are used near endpoints to avoid using blanked data, but those details are beyond the current scope. It should be clear that overlap regions must be large enough to accommodate the necessary fringing. Fringe points that cannot find a suitable set of donors are referred to as *orphan* points. There are several connectivity approaches ranging from fully manual to completely automatic. All computations in this thesis employ object X-rays [63] to determine the hole cut; overset communication is handled internally within OVERFLOW.

There exists a subtle distinction between the choice to blank nodes to reduce wasteful computation (as in the representative example) and the requirement to blank nodes that lie within the impermeable boundaries of solid bodies. Nodes in either category are assigned an `IBLANK = 0` value. However, physically meaningful information can be garnered at nodes in the former case (the solution on a fine grid can be restricted to the coarse grid it blanks), whereas the solution at nodes attributed to the latter case have no physical meaning; the solution is undefined. The implications are important to the Time-Spectral method but are largely tangential to the focus of the dissertation. Readers are directed to Chan [64] for a more comprehensive survey of overset technology.

$$\begin{bmatrix}
 1 & 0 & 0 & \cdots & \cdots & \cdots & \cdots & 0 \\
 -\alpha & 1 & \alpha & \ddots & & & & \vdots \\
 0 & -\alpha & 1 & \alpha & \ddots & & & \vdots \\
 \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & & \vdots \\
 \vdots & & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\
 \vdots & & & \ddots & -\alpha & 1 & \alpha & 0 \\
 \vdots & & & & \ddots & -\alpha & 1 & \alpha \\
 0 & \cdots & \cdots & \cdots & \cdots & 0 & 0 & 1
 \end{bmatrix}
 \begin{bmatrix}
 \Delta u_1 \\
 \Delta u_2 \\
 \Delta u_3 \\
 \vdots \\
 \vdots \\
 \Delta u_{10} \\
 \Delta u_{11} \\
 \Delta u_{12}
 \end{bmatrix}
 = -\alpha
 \begin{bmatrix}
 0 \\
 u_3^n - u_1^n \\
 u_4^n - u_2^n \\
 \vdots \\
 \vdots \\
 u_9^n - u_{11}^n \\
 u_{10}^n - u_{12}^n \\
 0
 \end{bmatrix}
 \quad (3.13)$$

$$\begin{bmatrix}
 1 & 0 & 0 & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & 0 \\
 -\beta & 1 & \beta & \ddots & & & & & & \vdots \\
 0 & -\beta & 1 & \beta & \ddots & & & & & \vdots \\
 \vdots & \ddots & -\beta & 1 & \beta & \ddots & & & & \vdots \\
 \vdots & & \ddots & 0 & 1 & 0 & \ddots & & & \vdots \\
 \vdots & & & \ddots & \ddots & \ddots & \ddots & \ddots & & \vdots \\
 \vdots & & & & \ddots & 0 & 1 & 0 & \ddots & \vdots \\
 \vdots & & & & & \ddots & -\beta & 1 & \beta & 0 \\
 \vdots & & & & & & \ddots & -\beta & 1 & \beta \\
 0 & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots & 0 & -2\beta & 1+2\beta
 \end{bmatrix}
 \begin{bmatrix}
 \Delta u_1 \\
 \Delta u_2 \\
 \Delta u_3 \\
 \Delta u_4 \\
 \Delta u_5 \\
 \vdots \\
 \vdots \\
 \Delta u_9 \\
 \Delta u_{10} \\
 \Delta u_{11} \\
 \Delta u_{12}
 \end{bmatrix}
 = -\beta
 \begin{bmatrix}
 0 \\
 u_3^n - u_1^n \\
 u_4^n - u_2^n \\
 u_5^n - u_3^n \\
 0 \\
 \vdots \\
 \vdots \\
 0 \\
 u_9^n - u_{11}^n \\
 u_{10}^n - u_{12}^n \\
 2(u_{12}^n - u_{11}^n)
 \end{bmatrix}
 \quad (3.14)$$

Chapter 4

Dynamic Hole Cutting

The fundamental obstacle of extending the Time-Spectral method to overset solvers concerns the nodes that dynamically move in and out of the physical domain due to the relative motion between the surface geometry and the background grid(s). Such nodes lack a complete time history of N physical states at the N equispaced temporal collocation points, rendering the infinitely-supported Fourier differentiation operator, and the standard Time-Spectral method upon which it is based, unsuitable. Thus, an alternative approach is required for representing the solution, and correspondingly its temporal derivative, at the unblanked collocation points at nodes undergoing dynamic hole cutting. Unfortunately, the incompatible Fourier interpolation-based temporal differentiation operator is the optimal choice for periodic problems, so the objective is to uncover and apply the best option among the remaining sub-optimal alternatives.

Figure 4.1 illustrates the issue at dynamically-blanked nodes for a representative one-dimensional oscillating piston. Fictional solutions in Fig. 4.1b corresponding to nodes a , b and c in Fig. 4.1a demonstrate dynamic hole cutting by the motion of the piston relative to the fixed background Eulerian grid. Node a never lies within the piston, and therefore its solution is defined over the entire time history of the period. The piston blanks node b briefly and node c twice. Thus, node b has one associated sub-periodic time interval while node c has two – one each represented by the solid and hashed lines in Fig. 4.1b. The shaded regions serve to highlight the time over which each node lies outside the physical domain (inside the impermeable boundary

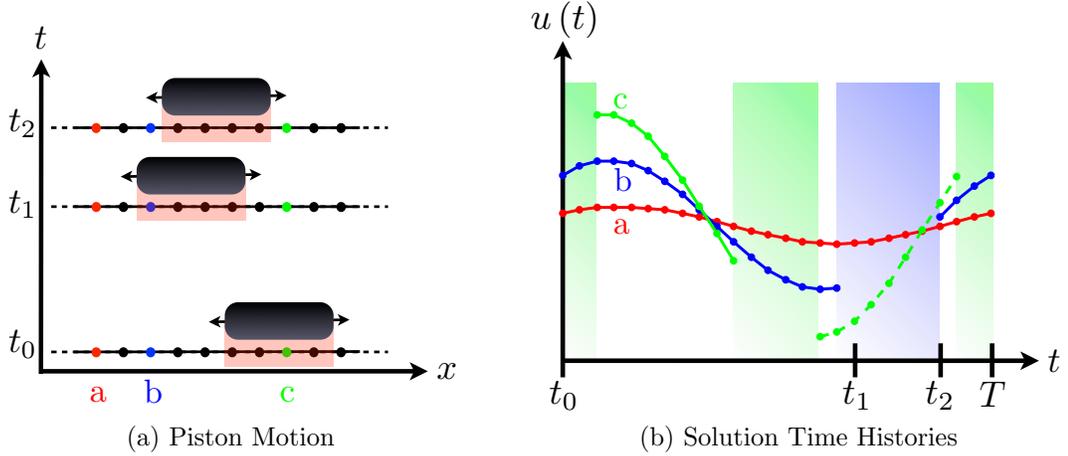


Figure 4.1: Figurative piston trajectory and solution histories. (a) Position of the piston in space at three time instances. The piston *cuts* nodes *b* and *c* as it oscillates over the background grid. (b) Intervals of defined solutions for nodes *a*, *b* and *c*. Shaded regions represent blanked regions through which the solution is undefined for nodes of the corresponding color.

of the piston) with an undefined solution. The standard Time-Spectral approach can be applied to node *a*, while special treatment is required for nodes *b* and *c*.

The objective of the proposed approach is to successfully approximate the solution, and therefore its temporal derivative, at each dynamically-blanked spatial node. Furthermore, it is critical to maintain equidistant collocation points at dynamically-blanked nodes. One of the more attractive features of the Time-Spectral method is the ability to apply it to existing flow solvers in a straightforward manner by leveraging existing spatial operators. Selecting a different distribution of temporal collocation points at dynamically-blanked nodes would require a complete overhaul of the spatial residual operators; spatial neighbors would no longer be guaranteed to be available at the same time instances, therefore requiring costly and potentially ill-conditioned temporal interpolation to evaluate spatial derivatives. Therefore, the distribution of evenly-spaced collocation points is constrained at all grid points, independent of their blanking status.

Three fundamentally different approaches were considered for treating dynamically-blanked nodes, and there exist several potential candidate treatments for each of the

three approaches. The various options are briefly introduced in each section, but only the most promising approaches are described in detail.

The first approach is global in nature and attempts to fit a basis of complex exponentials to the available data from all physically-defined nodes in some manner. The two candidate treatments for this approach result in an overdetermined system (least-squares) and an underdetermined system (least-norm); the least-squares option is detailed in §4.2.

The second fundamental approach is local in nature and partitions dynamically-blanked nodes into intervals of consecutively-unblanked nodes. The solution within each interval is independently expanded in a local basis that avoids extending the solution through undefined regions. A selection of candidate treatments are provided in §4.3 including both a periodization strategy (§4.3.1) and a bounded interval approach (§4.3.2).

Finally, a combination of the two approaches that uses a global description to recover the bulk characteristics of the solution in addition to a local correction that recovers an interpolation scheme consistent with the Time-Spectral method is examined in §4.4.

However, the chapter begins with an overview of the only generally applicable method published to date, where the solution in the blanked region is populated by a spatial smoothing operator. This process enables the standard Time-Spectral approach at dynamically-blanked nodes, but proves to be an inconsistent treatment.

The survey of potential treatments presented in this chapter is by no means an exhaustive list, but a sampling of some of the most natural options. The proposed framework is extensible in that the proposed hybrid Time-Spectral method requires only a linear differentiation operator. The hybrid Time-Spectral scheme was intended to be general, and it is possible that future investigation may uncover better choices than those described here which could then be incorporated with minimal effort.

4.1 Spatial Smoothing

The sole previous general approach attempting to tackle dynamic hole cutting for an overset Time-Spectral scheme is offered by Mavriplis et al. [39], where blanked nodes are populated with solution values by solving Laplace’s equation in the blanked regions at every iteration. The standard Time-Spectral method is then applied since all spatial nodes have complete time histories. While attractive for its straightforward implementation, the spatial-smoothing approach proves inconsistent; non-physical information provided by an alternative governing equation (Laplace’s equation) is propagated into the physical domain via the infinite support of the complex exponential basis functions. To demonstrate this inconsistency, see Fig. 4.2 which shows the solution of the one-dimensional linear advection equation, with $\Omega = [0, 1]$ and $\partial\Omega = 0$.

$$\frac{\partial u}{\partial t} + a \frac{\partial u}{\partial x} = 0, \quad x \in \Omega \quad (4.1)$$

$$u(0, t) = -\sin(\omega t), \quad x \in \partial\Omega \quad (4.2)$$

The analytic solution is $u(x, t) = \sin(\omega(x - at))$. An analytic expression for the spatial residual, $\frac{\partial u}{\partial x} = \omega \cos(\omega(x - at))$, is used to isolate the error introduced by the temporal discretization. A gap is defined in x to simulate a solid body, and the analytic solution is linearly interpolated (the solution to Laplace’s equation in one-dimension is the line connecting the Dirichlet endpoint data) through the blanked region to a blanked node at $x = x_b = 0.7$. The interpolation error across the spatial gap (Figure 4.2a) at one time sample corrupts the solution throughout the period (Figure 4.2b) via the infinite support of the Fourier basis. While the resulting solution may be smooth, it is inaccurate and dependent upon the arbitrary averaging of the solution through the spatial gap. In practice, the gaps can be any size and there is no guarantee that a spatially-smoothed value should result in an overall solution whose Fourier expansion successfully approximates the physically-defined portion of the solution. Consequently, spatial interpolation inhibits the desired temporal convergence (Figure 4.2c). If instead, a node located at $x = 0.5$ were blanked, the error

in the solution would be much lower because the solution at this location can be more accurately approximated by linear interpolation, reinforcing the assertion of the spatial-smoothing approach as arbitrary. Therefore, alternative approaches that do not depend on arbitrarily filling data at the blanked nodes were pursued.

4.2 Global Expansion of the Solution

The primary issue with the aforementioned spatial smoothing approach is its reliance on non-physical data provided by the Laplacian smoothing operator in the blanked regions. The infinite-support of the complex exponential basis functions couples this artificiality into the physical part of the domain, therefore corrupting the space-time solution. In the proposed global approach, the same complex exponential basis functions are used that implicitly satisfy the temporally periodic boundary conditions, but all unphysical data are ignored; sampling occurs only within the physically-defined regions of the temporal domain. Without a complete set of evenly-spaced collocation points distributed over the period, the standard differentiation operator cannot be applied because the discrete Fourier coefficients cannot be uniquely determined from the available data.

Two candidates were initially evaluated: a least-squares approach that approximates the known data in a reduced-frequency subspace, and a least-norm approach that uses the full spectrum of the global number of harmonics by minimizing the L_2 -norm of the N Fourier coefficients. While the complex exponential basis functions span the entire period, their expansion coefficients are based solely on data obtained at unblanked nodes that are populated with physically meaningful solutions. This is fundamentally different to the Laplacian-smoothing approach described in §4.1, where arbitrarily filled data are used to compute the temporal derivative at dynamically-blanked nodes.

The least-norm approach has been applied successfully in the field of compressed sampling [65], where a fundamental assumption is that information is contained within a small subset of the available modes. Compressed sampling often employs L_1 minimization to determine the coefficients. Only L_2 minimization is considered in the

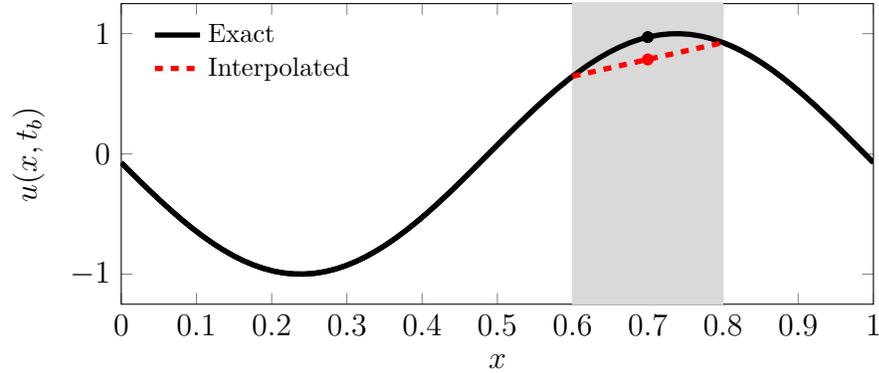
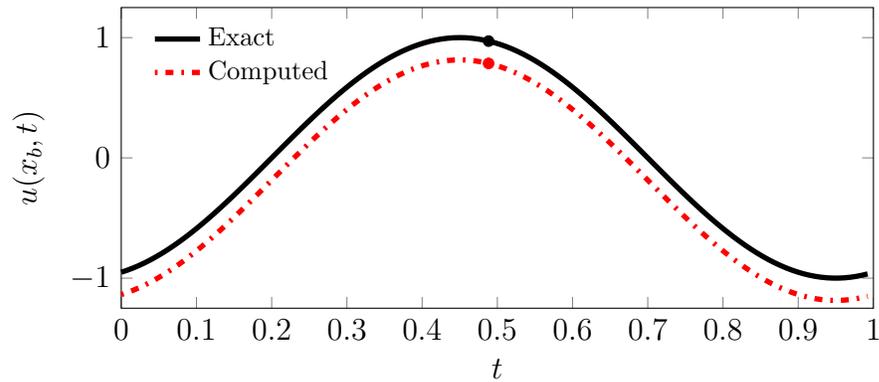
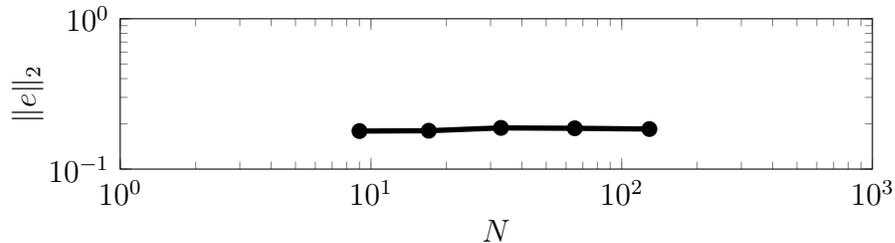
(a) Interpolation of the solution through the shaded *geometry* gap at time $t = t_b$ (b) Exact and computed solutions at $x = x_b = 0.7$ (c) Convergence of analytic error, $e = u - u_{ex}$, with N

Figure 4.2: Spatial smoothing Time-Spectral approach with a blanked node at $x = x_b = 0.7$. (a) The solution is interpolated through the gap in the shaded region to specify the solution at the blanked node. (b) The Time-Spectral solution resulting from spatial averaging may be smooth but is inconsistent, as the arbitrarily-filled value at the blanked node corrupts its complete time history. Severity of the corruption may be masked by fortuitous interpolation but it is highly dependent on the underlying solution and the gap through which it is filled. (c) The use of spatial averaging to populate the solution at the blanked point introduces an error that inhibits convergence with increased temporal resolution.

current approach, as it can be computed in a straightforward matrix operation. The least-norm approach appears quite desirable because it provides both interpolation and leverages the Fourier basis functions with their spectral accuracy and implicit periodic boundary conditions. However, the inclusion of so many frequencies, without enough data for them to be accurately determined, potentially results in extreme oscillations between the sample points, which in turn destroy the accuracy of its derivative approximation and proves unstable. Therefore, only the least-squares approach will be described in detail.

We seek to express the solution at the $N_P < N$ unblanked time-samples as a weighted sum of complex exponential basis functions. However, nonuniform sampling (the N_P active samples are no longer distributed evenly over the period) leads to ill-conditioning which generally prevents an invertible $N_P \times N_P$ transformation operator, Φ . In the least-squares approach, the solution at the N_P active nodes is orthogonally projected into a low-frequency subspace spanned by $M_P(N_P) < N_P$ harmonic basis functions. The frequency components of the projection can in turn be used to approximate the temporal derivative at the unblanked nodes.

The discrete projection, $\bar{u}_P(t)$, is expressed as a weighted sum of complex exponential basis functions, in a manner similar to the standard Time-Spectral approach.

$$\bar{u}_P(t) = \sum_{k \in \mathcal{K}(M_P)} \tilde{u}_k \phi_k(t), \quad \phi_k(t) = e^{ik\omega t} \quad (4.3)$$

Expressing Eq. 4.3 at the N_P active collocation points, t_j , with index set J_P

$$\bar{u}_P(t_j) = \sum_{k \in \mathcal{K}(M_P)} \tilde{u}_k \phi_k(t_j), \quad j \in J_P \quad (4.4)$$

results in the expression of Eq. 4.4 as a system of equations.

$$\bar{\mathbf{u}}_P = \Phi \tilde{\mathbf{u}} \quad (4.5)$$

The transformation operator, Φ , is composed of the $M_P < N_P$ basis functions.

$$\Phi_{jk} = \phi_k(t_j), \quad k \in \mathcal{K}(M_P), \quad j \in J_P, \quad (4.6)$$

The frequency set, $\mathcal{K} = \mathcal{K}(M_P)$, is determined by the number of basis functions.

$$\mathcal{K}(M_P) = \begin{cases} \{-M_P/2, \dots, M_P/2 - 1\}, & M_P \text{ even} \\ \{-(M_P - 1)/2, \dots, (M_P - 1)/2\}, & M_P \text{ odd} \end{cases} \quad (4.7)$$

The choice of M_P is a trade-off between accuracy and stability; typically the approximation improves with the inclusion of a broader spectrum, but the conditioning of the system degrades. Herein, the number of basis functions is assigned a value of roughly half the number of time samples.¹

$$M_P = \lfloor N_P/2 \rfloor \quad (4.8)$$

A solution to the overdetermined system, $\mathbf{u}_N = \Phi \tilde{\mathbf{u}}_P$, does not exist in general, resulting in a nonzero discrete residual, \mathbf{e}_P^N , between the sampled data, \mathbf{u}_N , and its discrete projection, $\bar{\mathbf{u}}_P$, at the sample points.

$$\mathbf{u}_N = \Phi \tilde{\mathbf{u}}_P + \mathbf{e}_P^N \quad (4.9)$$

The coefficients are uniquely determined in the least-squares sense, where the discrete residual is taken to be orthogonal to the subspace spanned by Φ . The normal

¹An odd number of basis functions, M_P , are retained for the complex exponential least-squares approximation in order to maintain a real-valued projection into the subspace spanned by the lowest K_P harmonics; $M_P = 2K_P + 1$. The complex exponential notation has been used to describe the Time-Spectral and least-squares Fourier approximation. However, trigonometric basis functions (sines and cosines) are used in practice for the Fourier least-squares approach. Both the input data and the trigonometric transformation operator are guaranteed to be real-valued, allowing the number of basis function, M_P , to be even or odd without risk of introducing imaginary data into the projection.

equations emerge by premultiplying Eq. 4.9 by the conjugate transpose of the transformation operator.

$$\Phi^* \mathbf{u}_N = \Phi^* \Phi \tilde{\mathbf{u}}_P + \Phi^* \mathbf{e}_P^N \quad (4.10)$$

Because the error is chosen to be orthogonal to the subspace spanned by Φ , the second term in Eq. 4.10 is identically zero.

$$\Phi^* \mathbf{u}_N = \Phi^* \Phi \tilde{\mathbf{u}}_P \quad (4.11)$$

The coefficients, $\tilde{\mathbf{u}}_P$, are uniquely determined by solving the invertible $M_P \times M_P$ system of equations.

$$\tilde{\mathbf{u}}_P = (\Phi^* \Phi)^{-1} \Phi^* \mathbf{u}_N \quad (4.12)$$

Substitution of Eq. 4.12 into Eq. 4.5

$$\bar{\mathbf{u}}_P = \Phi (\Phi^* \Phi)^{-1} \Phi^* \mathbf{u}_N \quad (4.13)$$

defines the operator, $P_P^N = \Phi (\Phi^* \Phi)^{-1} \Phi^*$, that projects time-domain data sampled at the set of active time-samples into the subspace spanned by $\mathcal{K}(M_P)$.

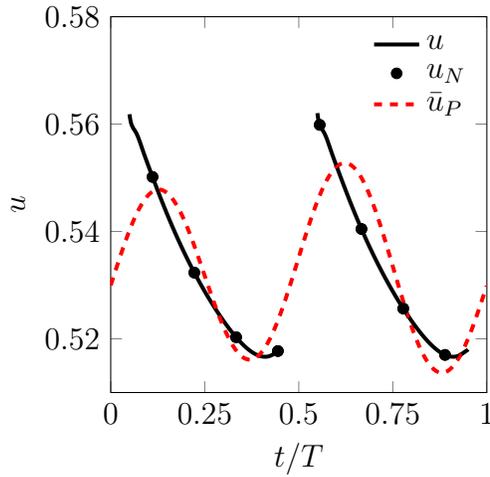
$$\bar{\mathbf{u}}_P = P_P^N \mathbf{u}_N \quad (4.14)$$

Note that the projection will not interpolate the data upon which it operates unless the data lie within the low-frequency subspace.

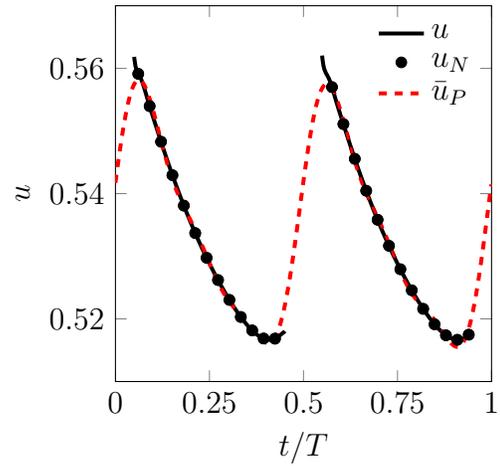
$$\bar{u}_P(t_j) \neq u_N(t_j), \quad j \in J_p \quad (4.15)$$

The projected solution can be interpolated via Eq. 4.3.

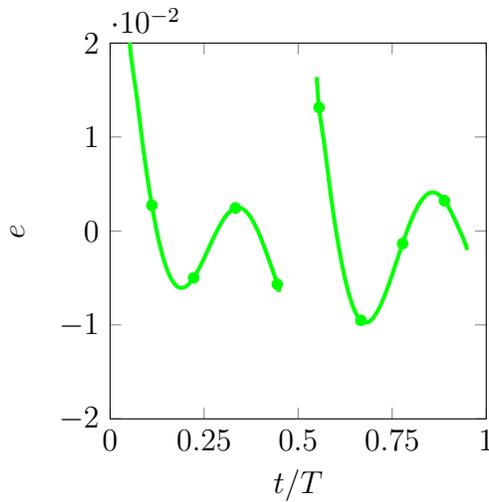
Figure 4.3 demonstrates the least-squares approximation to a representative signal (taken from the inviscid NACA 0012 plunging case in §6.1.1). For the first case, shown in Fig. 4.3a, only nine global collocation points are used; eight collocation points



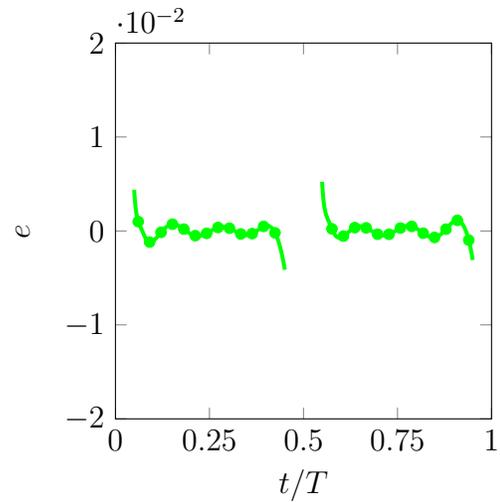
(a) Projected solution for $N_P = 8$ unblanked nodes



(b) Projected solution for $N_P = 26$ unblanked nodes



(c) Residual for $N_P = 8$ unblanked nodes



(d) Residual for $N_P = 26$ unblanked nodes

Figure 4.3: Global least-squares Fourier projection and residual, $e = u - \bar{u}_P$, for $N = 9$ and $N = 33$ global collocation points.

remain unblanked. The least-squares representation demonstrates poor agreement with the solution both at the collocation points and over its continuous expansion. However, for the case of thirty-three global collocation points shown in Fig. 4.3b, twenty-six nodes remain unblanked resulting in a least-squares approximation that matches the solution quite successfully, despite not interpolating the discrete data at the nodes. The residuals, $e = u - \bar{u}_P$, for both cases are plotted in Figs. 4.3c and 4.3d, respectively.

The differentiation operator in the frequency domain is the $M_P \times M_P$ matrix similar to standard diagonal operator, D_N , in §2.1, where $d_{kk} = i\omega k$. The real-valued trigonometric differentiation operator that is used in practice, for M_P even or odd, has a different form, but can be derived in a similar fashion.

A time-domain least-squares Fourier differentiation operator can be constructed in a similar manner to the standard Fourier differentiation operator by discretely transforming the scaled coefficients back into the time domain.

$$\frac{d}{dt} \mathbf{u}_N = \Phi D_P \tilde{\mathbf{u}}_P \quad (4.16)$$

The differentiation operator is constructed by substituting $\tilde{\mathbf{u}}_P$ with Eq. 4.12, which reveals that is simply the projection operator an embedded factor, D_P .

$$\mathcal{D}_N = \Phi D_P (\Phi^* \Phi)^{-1} \Phi^* \quad (4.17)$$

such that

$$\frac{d}{dt} \mathbf{u}_N = \mathcal{D}_N \mathbf{u}_N. \quad (4.18)$$

This form is similar to Eq. 2.22, where scalar differentiation is applied to the discrete frequency-domain expansion coefficients followed by a transformation back into the time domain.

As previously mentioned, the overdetermined system does not interpolate the data. Therefore, its projection results in an orthogonal residual. Thus, direct application within the Time-Spectral collocation (trigonometric *interpolation*) scheme is

inconsistent. The discrete residual orthogonal to the subspace spanned by Φ , \mathbf{e}_P^N , is expressed as

$$\mathbf{e}_P^N = \mathbf{u}_N - \bar{\mathbf{u}}_P \quad (4.19)$$

$$= (I - P_P^N) \mathbf{u}_N \quad (4.20)$$

We will return to the concept of a global least-squares approach in §4.4, with the development of an aggregate basis set that represents the solution as a combination of the least-squares projection and its associated residual, which is itself expanded in a localized basis. The addition of the residual term lifts the combined representation such that it interpolates the data resulting in an approach consistent with the Time-Spectral interpolation scheme.

4.3 Local Expansion of the Solution

The previous section outlined global representations of the solution at dynamically-blanked nodes. Another natural approach to avoid corruption from unphysical data is to relinquish the global description of the solution by partitioning the temporal domain into independent sub-periodic intervals. For a local description of the solution, time histories for dynamically-blanked spatial nodes are partitioned into intervals of consecutively-unblanked temporal collocation points. Independent temporal bases are used to represent the solution within each partition for such nodes. This does introduce a potential hazard for rapidly blanking nodes, i.e. if a node switches between blanked and unblanked at every point, the solution will be considered constant in time at each of the nodes and therefore the time derivative will be approximated as zero universally. This may prove problematic, but perhaps such situations would require higher temporal resolution to resolve the high-frequency motion, suggesting the use of temporal adaptation (See §7.3) to avoid costly global refinement. This section surveys several candidate treatments with the goal of finding the best available basis to resolve these local descriptions of the solution, with the objective of developing a consistent hybrid Time-Spectral scheme capable of resolving relative motion.

To achieve this goal, two primary approaches are evaluated. The first option is the periodization of the solution at unblanked intervals using complex exponentials to fit data from the physical portion of the time signal. Examples of this approach are Fourier continuation (FC), which extrapolates a non-periodic function into a periodic function on a larger domain [66, 67, 68], and compressed sampling, which applies L_1 or L_2 minimization to solve an underdetermined system [65]. These approaches are simply the localized versions of those found in §4.2 that fit a least-squares or least-norm approximation to the defined data *through* undefined regions, whereas the current proposal is that decoupled least-squares or least-norm approximations are used to fit each independent interval separately. A second approach is to expand the solution within bounded intervals of consecutively-defined nodes with a localized differentiable basis. Both of the aforementioned approaches are described and analyzed in the upcoming sections. Global temporal support is thus abandoned; however, this approach is consistent with the physics of disjoint domains separated by an impermeable boundary.

4.3.1 Periodization

By definition, a periodic phenomenon repeats itself every T , even if it is undefined over a portion of the period. Dynamic hole cutting results in valid time samples over a subset of the period. Despite being aperiodic on the subdomain, it is possible to assume a global periodic function exists that matches the solution on the sub-periodic domain. Furthermore, the domain can be subdivided into several intervals that are each treated with an independent basis. Again, complex exponential basis functions that span the global period are used, but their coefficients are determined using information only within each independent interval. No information is ever used or applied outside of the particular interval.

Two approaches are considered for the periodization approach: Fourier continuation which is simply the localized analogue of the global Fourier least-squares approach described in the previous section, and a local least-norm approach. Each case employs the complex exponential basis functions with their implicit boundary

conditions. However, the negative characteristics described in the previous section with respect to the least-norm approach would also apply to its local analogue, and therefore the periodized least-norm approach is omitted from further consideration.

Fourier continuation (commonly referred to as Fourier Extension) was developed to leverage the spectral accuracy properties of the Fourier series approximation in application to aperiodic signals. Standard Fourier series approximation to aperiodic signals results in a discontinuity at the boundaries resulting in spurious oscillations, referred to as Gibbs' phenomenon. To overcome this difficulty, Fourier continuation assumes a globally periodic surrogate function on an extended domain and attempts to match the signal within the original domain of interest (each unblanked interval) in some manner. The degree by which the *physical* boundary is extended is a parameter, b , that has serious implications for the accuracy of the resulting approximation.

The formal statement of the problem is to minimize Eq. 4.21 in a least-squares sense by using $M_P < N_P$ Fourier basis functions to best match the known sub-periodic data of N_P samples following Bruno et al. [68].

$$f(x_j) = \sum_{k \in \mathcal{K}(M_P)} \tilde{u}_k e^{\frac{2\pi}{b} i k x_j}, \quad j = 1, \dots, N_P \quad (4.21)$$

A smooth function, $f(x)$, is defined on the sub-periodic interval for $x \in [0, 1]$, and \mathcal{K} is the set of harmonics included in the subspace as defined in the previous section. Development of this approach is identical to the approach taken in Section 4.2 except that we now use N_P as the number of physically-defined nodes within a given interval of consecutively-defined collocation points and not the *total* number of defined collocation points at a particular spatial node (these are equivalent if there exists only a single interval of consecutively-unblanked nodes). The sample points, x_j , are not generally constrained to be evenly-spaced, but are so in the context of the proposed Time-Spectral approach. More sophisticated extension techniques have been developed [66] and in some cases applied to flow solvers [69, 70, 66, 71] in the spatial dimensions with a large number of degrees of freedom. However, the small number of degrees of freedom typically employed in the Time-Spectral method render those efforts unsuitable.

A common example of Fourier continuation, demonstrated by Boyd [67] and Bruno et al. [68], is the extension of the aperiodic function, $f(x) = x$, for $x \in [0, 1]$. The *period* is extended to a value $b > 1$ and the Fourier continuation technique is applied to best approximate f in the interval $x \in [0, 1]$. The approximation of the derivative can then be evaluated by applying the differentiation operator defined by Eq. 4.18. The choice of b in this example is arbitrary with significant implications on the quality of the approximation of f and its derivative. Figure 4.4 demonstrates the sensitivity of the continuation to the parameter b . In the case of $b = 1.25$, there is not enough space for the continued function to smoothly transition from one endpoint to the other, resulting in significant oscillations (Figure 4.4a), which destroy its derivative approximation (Figure 4.4b). On the other hand, choosing $b = 2.0$ provides a better projection that accurately represents the function (Figure 4.4c) and its derivative (Figure 4.4d) for most of the *physical* interval, $x \in [0, 1]$.

For the case of dynamically-blanked Time-Spectral nodes, the surrogate domain is imposed by the global period of the problem ($b = T$) and the aperiodic subinterval is provided by the number of consecutively-unblanked time samples. Noting that $\omega = 2\pi/b$ in 4.21 is the fundamental frequency of the solution, the truncated transformation matrix, Φ , is defined in the same manner as in Section 4.2; the $N_P \times M_P$ Fourier continuation basis matrix is composed of the M_P column vectors, $\phi_k(t_j) = e^{ik\omega t_j}$, whose length is the number of valid time samples within a particular interval, N_P . The projection can also be differentiated using the result from §4.2.

Figure 4.5 demonstrates the Fourier continuation treatment. Revisiting the same representative signal from the previous section, the case of $N = 9$ global collocation points results in only a single blanked node. The current implementation interprets only a single sub-periodic interval, whereas the continuous signal is partitioned into two distinct segments. Therefore, the localized continuation approach (Figure 4.5a), and its associated residual (Figure 4.5c) for this particular case, are identical to the global least-squares result (cf. Figures 4.3a and 4.3c). This is not the case for the example employing $N = 33$ global collocation points. Two distinct intervals, each composed of thirteen consecutive time samples, are detected. The solution at the nodes within each partition is projected into a subspace spanning $M_P = \lfloor \frac{13}{2} \rfloor =$

6 harmonic basis functions (Figure 4.5b) compared to the least-squares approach where the solution through both consecutively-defined intervals is spanned by a total of $M_P = \lfloor \frac{26}{2} \rfloor = 13$ harmonic basis functions (Figure 4.3b). The residual in the partitioned case (Figure 4.5d) is somewhat smaller than the least-squares approach for this particular example. Despite this fact, the projection does not interpolate the data provided, $u_N(t_P) \neq \bar{u}_P(t_P)$, which is inconsistent with the interpolation property of collocation methods.

4.3.2 Bounded Interval

Alternatively, we can abandon the search for a periodization of the sub-periodic data, and instead employ localized bases whose support is limited to the bounded interval of consecutively-unblanked time samples. This poses an issue due the requirement of maintaining equispaced sampling. Orthogonal polynomials such as Chebyshev polynomials, offer spectrally-accurate interpolation and differentiation on aperiodic domains, but are severely ill-conditioned on equispaced samples, often exhibiting violent oscillations on the boundaries referred to as Runge's phenomenon [72]. Endpoint clustering resolves this issue but violates the prescribed constraint of uniform time sampling in the current application. Thus the problem is now reduced to finding a robust and efficient method for evaluating the temporal derivatives of non-periodic functions on equidistant nodes.

In this work, several candidate approaches have been considered for the approximation of the temporal derivative within discrete intervals of consecutively-defined nodes. The pool of candidates include differentiation operators derived from finite-difference approximations, a least-squares projection of orthogonal polynomials, wavelets, barycentric rational interpolants and splines. Only the rational interpolant- and spline-based approaches have been implemented as they were deemed the most promising options, but all of the aforementioned candidates are briefly introduced.

Finite-difference formulas are commonly used to approximate derivatives in both time and/or space in many numerical methods. They are constructed by choosing a stencil width and using Taylor-series approximations to determine the coefficients, a_i ,

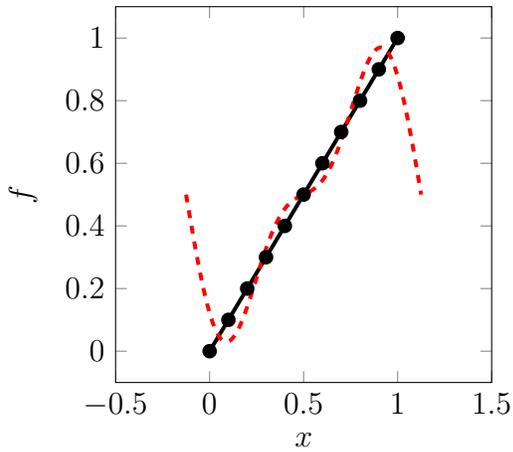
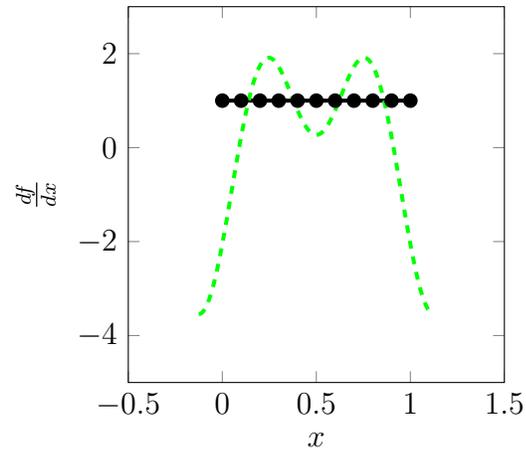
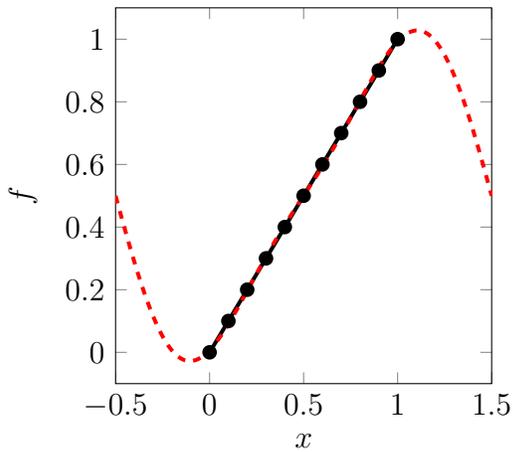
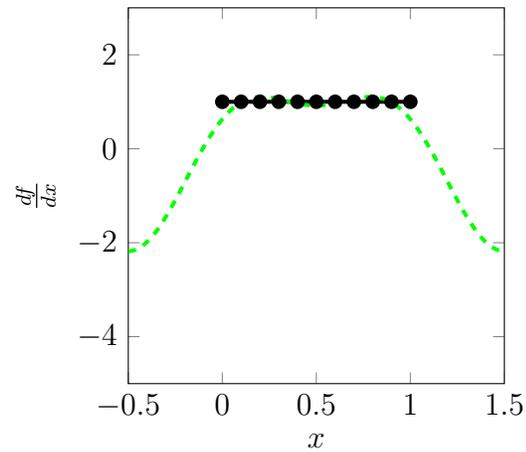
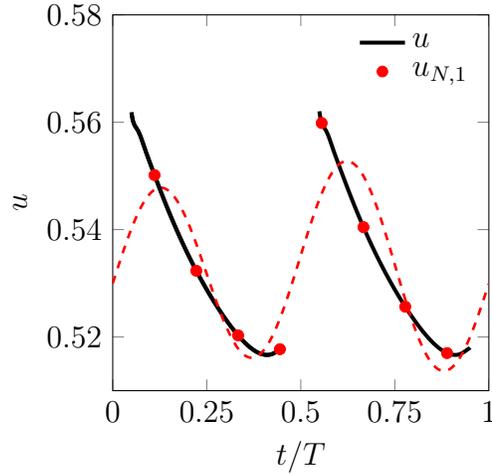
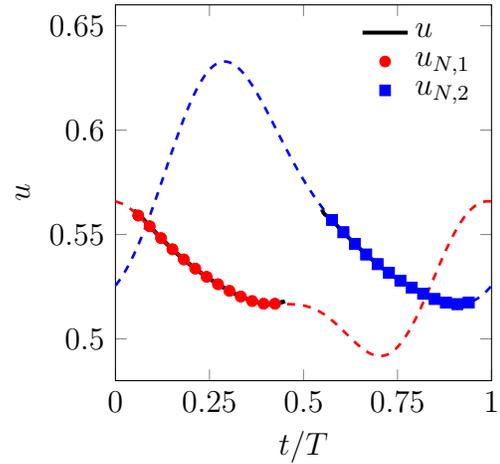
(a) Function $f(x) = x$ continued to $b = 1.25$ (b) Derivative of continuation, $b = 1.25$ (c) Function $f(x) = x$ continued to $b = 2.0$ (d) Derivative of continuation, $b = 2.0$

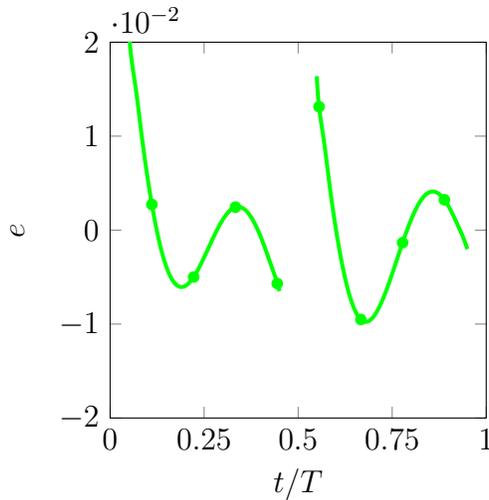
Figure 4.4: Fourier continuation of $f(x) = x$ for $x \in [0, 1]$ and its derivative approximation using $N_P = 11$ samples and $M_P = 5$ basis functions for $b = 1.25$ and $b = 2.0$. Exact values in black.



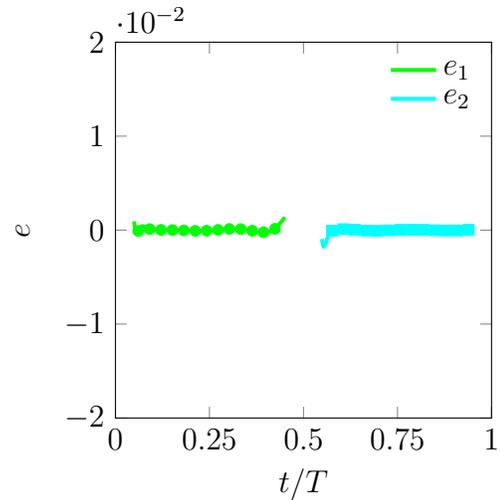
(a) Projected solution for $N_P = 8$ unblanked nodes within a single interval



(b) Projected solution for $N_P = 26$ unblanked nodes within two independent intervals



(c) Residual for $N_P = 8$ unblanked nodes within a single interval



(d) Residual for $N_P = 26$ unblanked nodes within two independent intervals

Figure 4.5: Local least-squares Fourier continuation (projection) and residual, $e = u - \bar{u}_P$, for $N = 9$ and $N = 33$ global collocation points.

that provide the best order of accuracy in approximating an m -th degree derivative at node j . The stencil width spans $p + q + 1$ nodes [73] resulting in a truncation error ϵ_t dependent on each of the parameters.

$$\left(\frac{\partial^m u}{\partial x^m}\right)_j - \sum_{i=-p}^q a_i u_{j+i} = \epsilon_t \quad (4.22)$$

To reiterate the assertion in §2.2, the algebraic convergence rate of any finite-difference approximation proves inferior to the spectral convergence afforded by a global scheme such as a Fourier collocation. Compact or Padé schemes [74] may also be used. However, the number of active nodes, N_P , for a given sub-periodic interval is often quite small. Therefore, we seek a more global approximation with improved approximation and differentiation properties over traditional finite-differences. Furthermore, the finite-differences are not a basis set. They are defined by differentiating local Lagrange interpolants but cannot be expanded as a single solution spanning the domain. However, the localized nature of the finite-differences leads to better handling of discontinuities. They also lead to fixed-bandwidth differentiation operators, that can be directionally factored to provide efficient inversions to solve implicit systems of equations. Nevertheless, a finite-difference approximation is not further considered.

Wavelets emerged in the late twentieth century as a novel orthogonal, compactly-supported basis to gauge frequency information about a signal without abandoning temporal localization, which is the case for the Fourier transform [75, 76]. Incorporation of wavelets was suggested by Blanc et al. [25] as a potential tool in developing a method to handle dynamically-blanked nodes for an overset Time-Spectral scheme. Wavelets have been employed in finite element methods [77]. Because the wavelets are active within a limited range of a particular node, it is possible to apply them as a basis upon which the solution within each sub-periodic partition can be expanded. Many wavelets are not smooth and are not defined in closed form. The global wavelet approximation is also assumed to be periodic, and therefore requires special treatment on interval boundaries [78]. They have yet to be implemented for the hybrid Time-Spectral scheme but may be utilized in the future for either the local or global approach (nonuniform sampling on a globally periodic domain).

Eisinberg and Fedele [79] demonstrated that a least-squares approximation of orthogonal polynomials mitigates the severe ill-conditioning exhibited with traditional polynomial interpolation on equidistant nodes. This property can be used to develop a transformed temporal differentiation operator based on orthogonal polynomials in a similar manner as for the least-squares Fourier-based differentiation operator. However, the least-squares approximation again results in a non-interpolating scheme and is therefore not further considered; it offers neither interpolation (offered by splines or the barycentric rational interpolants introduced hereafter) nor implicit periodic boundary conditions (provided by a least-squares Fourier approximation).

Splines offer yet another method of representing a solution within a sub-periodic interval by linking piecewise polynomials into a single approximation within the bounded domain. Splines are typically limited to a given approximation order, p , that is a function of its piecewise polynomial representation, suggesting that spectral convergence is not to be expected for dynamically-blanked nodes. Cubic B-Splines are considered but any type of splines can be applied if a linear differentiation operator can be constructed. In the case of cubic B-splines, the solution is comprised of piecewise cubic polynomials that are determined by enforcing continuity of the first and second derivatives at node interfaces and by supplying a particular boundary condition. Both *natural* splines, which assume the solution is linear at the endpoints ($u'' = 0$), and *not-a-knot* splines, which enforce continuity of the third derivative at the endpoints, are considered [74]. The spline-based differentiation operators are not explicitly derived, but both the natural and knot-a-knot approaches have been implemented and employed in §6.1.1

The approach finally adopted in this work is to use barycentric rational interpolants (BRI). These provide a viable alternative to orthogonal polynomials for equidistant data on an aperiodic domain [80]. Bos et al. [81] demonstrated their superior approximation and differentiation properties on equidistant nodes over conventional orthogonal polynomials. Efforts to explore rational interpolants, and their utility as a pseudospectral basis for spectral collocation methods, include but are not limited to [82, 83, 84, 85].

Baltensperger and Berrut [86] define the rational interpolant, $r(x)$, approximating

the function f in barycentric form

$$r(x) = \frac{\sum_{k=0}^N \frac{w_k}{x-x_k} f(x_k)}{\sum_{k=0}^N \frac{w_k}{x-x_k}} \quad (4.23)$$

and its corresponding differentiation operator, D_N , such that $\frac{d}{dx}r = D_N r$

$$d_{jk} = \begin{cases} \frac{w_k}{w_j} \frac{1}{(x_j-x_k)} & \text{if } j \neq k \\ -\sum_{i=0, i \neq k}^N d_{ji} & \text{if } j = k \end{cases} \quad (4.24)$$

where $r_k = r(x_k) = f(x_k)$. The interpolation can be reformulated as a weighted sum of nodal basis functions, $\phi_k(x)$, with coefficients equal to the function value at each node, $f_k = f(x_k)$.

$$I_N f(x) = \sum_{k=0}^N f_k \phi_k(x), \quad \text{with } \phi_k(x) = \frac{\frac{w_k}{x-x_k}}{\sum_{k=0}^N \frac{w_k}{x-x_k}} \quad (4.25)$$

The barycentric rational basis functions nodally interpolate the solution data, $\Phi_{jk} = \phi_k(x_j) = \delta_{jk}$, resulting in an identity transformation operator that implies discrete orthogonality at the sample points. It follows that the transformed differentiation operator is identical to D_N : $\mathcal{D}_N = \Phi D_N \Phi^{-1} = D_N$.

While similar in form to the Lagrange interpolant, a key distinction of the barycentric rational interpolant is in how the weights, w_k , are defined. The Lagrange interpolant is constrained to pass through $N + 1$ points as a polynomial of degree N . In contrast, while the rational interpolant passes through the data, it is not constrained to do so as a polynomial of degree N . This relaxation aids in preventing spurious oscillations at the endpoints of equispaced nodes associated with the traditional Lagrange or Chebyshev polynomial bases, while retaining powerful interpolation and differentiation properties.

Floater and Hormann [87] derived weights that provide an approximation order

$d + 1$ while avoiding poles. The weights, w_k , control the accuracy and stability characteristics of the rational interpolant which is a blend of polynomials of degree d . Weights guaranteeing an absence of poles

$$w_k = (-1)^{k-d} \sum_{i \in J_k} \prod_{j=i, j \neq k}^{i+d} \frac{1}{|x_k - x_j|} \quad (4.26)$$

are defined for $N + 1$ samples where

$$I := \{0, 1, \dots, N - d\} \text{ and } J_\alpha := \{i \in I : \alpha - d \leq i \leq \alpha\}$$

for the desired approximation order, d . For uniform nodes, the resulting weights can be scaled to integer values without changing the interpolant.

$$w_k = (-1)^{k-d} \sum_{i \in J_k} \binom{d}{k-i} \quad (4.27)$$

Readers are directed to the derivation in [87] for additional detail. Figure 4.6 illustrates three basis functions corresponding to the first, second and fourth nodes with approximation order $d = 3$ over 8 equispaced sample points with the values highlighted at the nodal points. Note the interpolation property of the basis, $\phi_k(x_j) = \delta_{jk}$.

The application of barycentric rational interpolation to the representative signal is demonstrated in Fig 4.7 (with approximation order limit, $d_{\max} = 1$). Note that using a resolution of only nine global collocation points, depicted in Fig. 4.7a, does not interpret an interruption to the signal and represents the solution spanning the eight unblanked nodes in a single basis (Figure 4.7a). Using $N = 33$ global collocation points does, in fact, result in two distinct intervals that are each represented by independent barycentric rational interpolants (Figure 4.7b). Residual plots in Figs. 4.7c and 4.7d demonstrate excellent agreement with the continuous segments of the solution away from interval boundaries. Extrapolation outside of the bounding collocation points is not required for differentiation and included in the figures for consistency with previously demonstrated approaches. It may be beneficial to detect

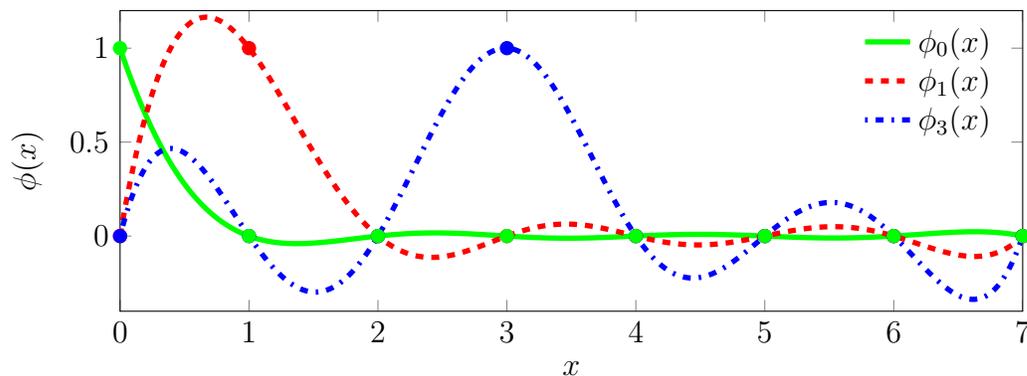
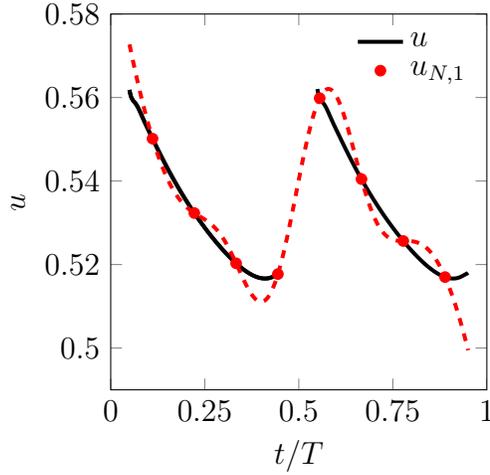


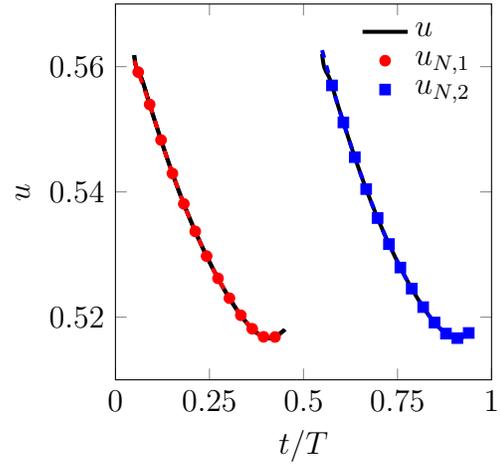
Figure 4.6: Barycentric rational interpolant basis functions $\phi_0(x)$, $\phi_1(x)$ and $\phi_3(x)$ corresponding to nodes 0, 1 and 3, respectively, for approximation order $d = 3$ over 8 equispaced sample points.

the blanked region for the case of $N = 9$. This would require a potentially expensive preprocessor that uses a predetermined or user-supplied granularity to cycle through the motion, recording the dynamically-blanked boundaries at a finer resolution. Using this information, the solution in Fig. 4.7a would also be partitioned and not interpolated through undefined regions.

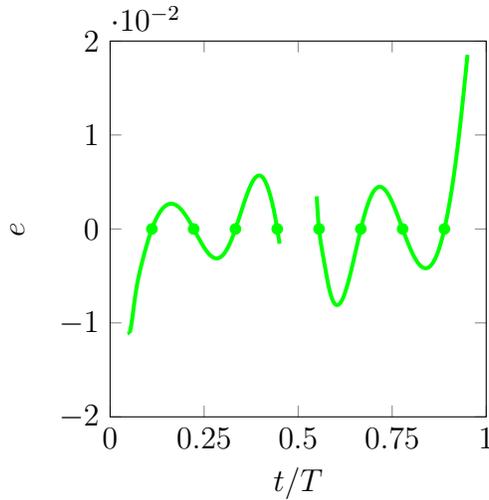
In the Time-Spectral method, differentiation of the expansion is required and therefore the accuracy properties for each candidate basis must be assessed. Here, the differentiation performance of the rational interpolant is compared to that of finite differences, the Fourier differential operator (generally optimal for periodic functions), the Chebyshev differentiation operator (generally optimal for non-periodic functions) on Chebyshev nodes and differentiation operators derived from cubic splines. An even-odd harmonic function (Figure 4.8a) and Runge's function (Figure 4.8b) are differentiated by a selection of the aforementioned methods. Convergence of differentiation error, $e = \|\frac{df}{dx} - \mathcal{D}_N f\|_2$, versus N is plotted in Fig. 4.9 for the harmonic function and Fig. 4.10 for Runge's function. The rational interpolant approximation order, $d = \min(\lfloor \frac{N-1}{2} \rfloor, d_{\max})$, is defined for different values of d_{\max} . Fourier differentiation exhibits roundoff error for the harmonic function with increased N and offers poor convergence for Runge's function. Chebyshev differentiation demonstrates



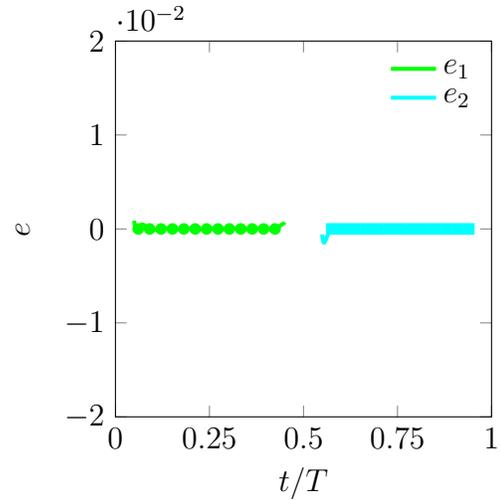
(a) Interpolated solution for $N_P = 8$ unblanked nodes within a single interval



(b) Interpolated solution for $N_P = 26$ unblanked nodes within two independent intervals



(c) Residual for $N_P = 8$ unblanked nodes within a single interval



(d) Residual for $N_P = 26$ unblanked nodes within two independent intervals

Figure 4.7: Barycentric rational interpolation and residual, $e = u - I_N u$, for $N = 9$ and $N = 33$ global collocation points with the default setting of $d_{\max} = 1$.

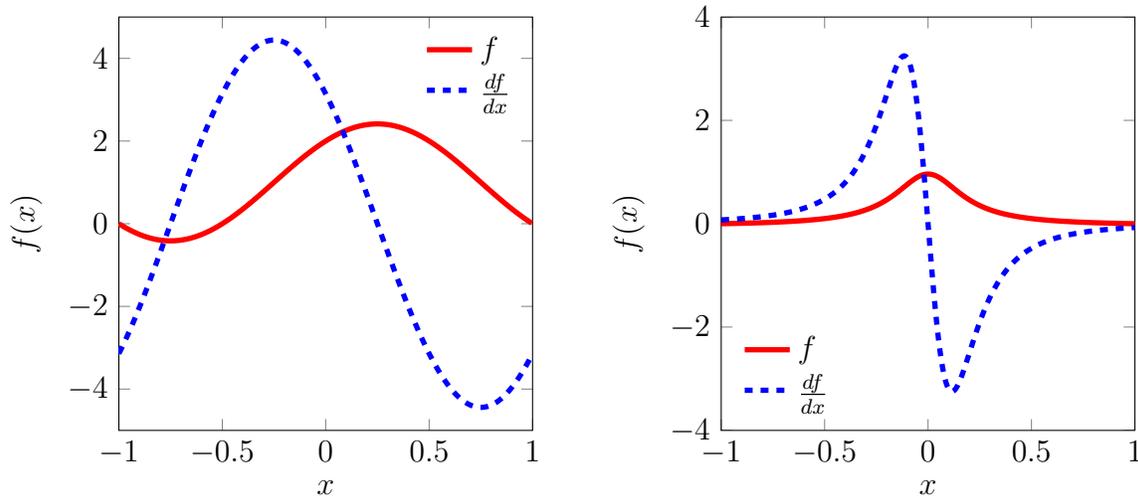
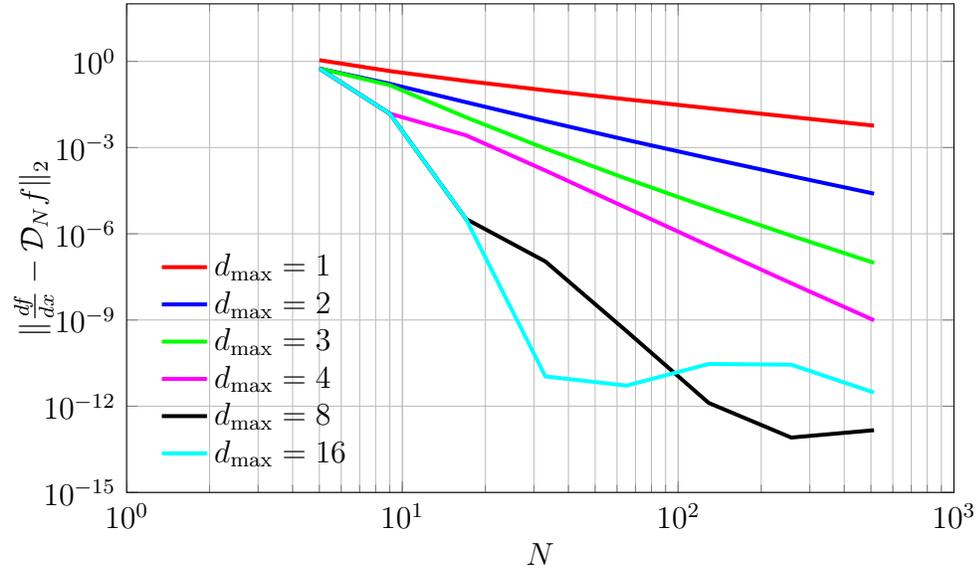
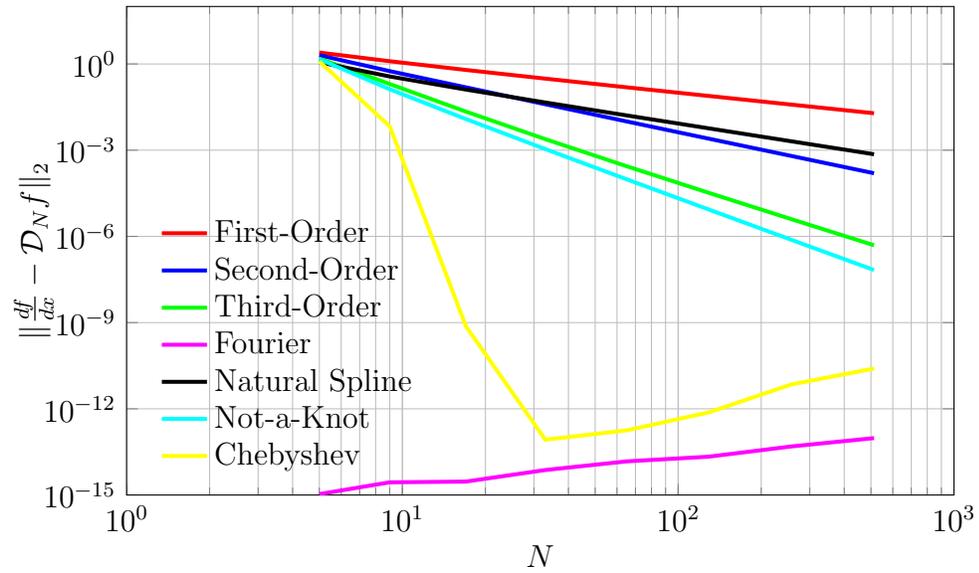
(a) Harmonic function, $f(x) = 1 + \cos(\pi x) + \sin(\pi x)$ (b) Runge's function, $f(x) = \frac{1}{1+25x^2}$

Figure 4.8: Analytic functions, f , and analytic derivatives, $\frac{df}{dx}$, for $x \in [-1, 1]$. (a) Even-odd harmonic function and (b) Runge's function.

spectral convergence for both functions (albeit delayed for Runge's function in contrast to the harmonic function) and boundary condition selection makes a significant difference for spline-based differentiation. While the rational interpolant-based differentiation operator demonstrates spectral-like convergence for $d_{\max} \in \{8, 16\}$ on the smooth harmonic function, it diverges for those values of d for small N in the case of Runge's function. Taking the best rational result for every N in the Runge's function example produces spectral-like convergence (outpacing the Chebyshev convergence) reinforcing the importance of selecting an appropriate d . Using a lower value of d is associated with higher compactness, and is therefore more successful for high-frequency functions (e.g. Runge's function) for small N . For smooth functions like the harmonic example, a large value of d poses no problem. Optimal selection of d and d_{\max} is problem dependent and an area of continuing research [84, 88]. It is expected that the choice of d will strongly influence the stability of the matrix operator; stability considerations are addressed in Appendix B.

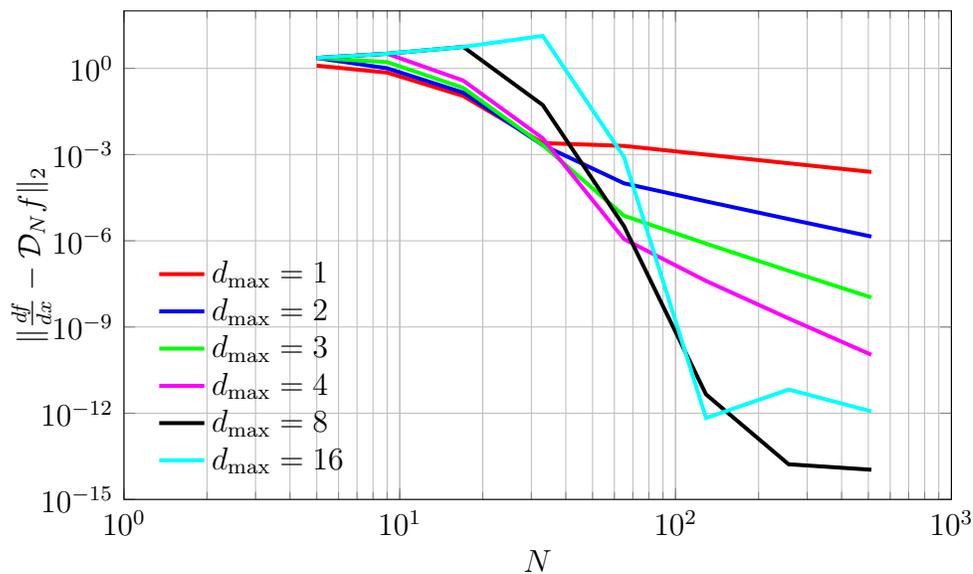


(a) Convergence of barycentric rational interpolant-based differentiation varying the values of d_{\max} .

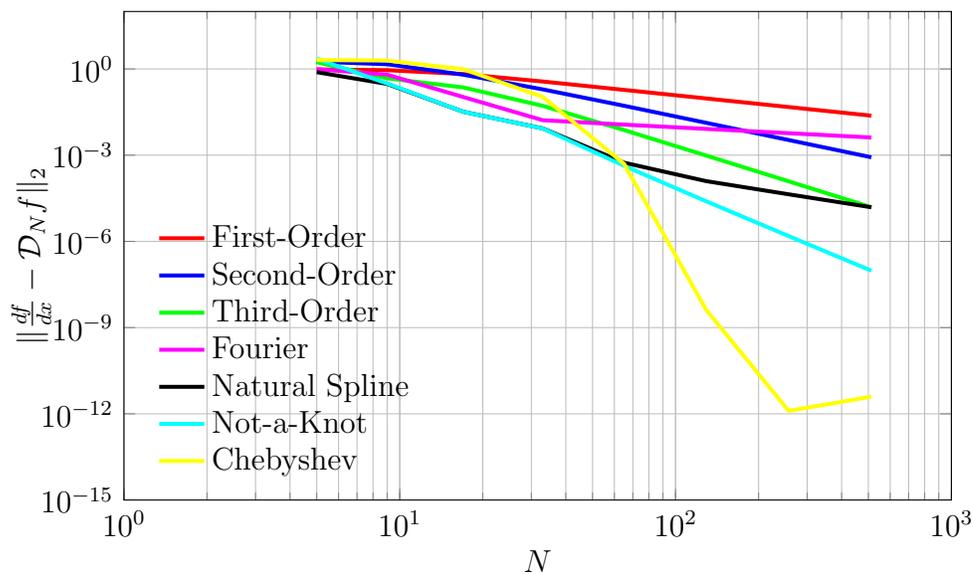


(b) Convergence of finite-difference-, Fourier-, spline- and Chebyshev-based differentiation.

Figure 4.9: Convergence of differentiated harmonic function, $f(x) = 1 + \cos(\pi x) + \sin(\pi x)$, $x \in [-1, 1]$, using the barycentric rational interpolant-based differentiation operator for different values of d_{\max} and other differentiation operators on equispaced nodes (Chebyshev on Chebyshev nodes).



(a) Convergence of barycentric rational interpolant-based differentiation varying the values of d_{\max} .



(b) Convergence of finite-difference-, Fourier-, spline- and Chebyshev-based differentiation.

Figure 4.10: Convergence of differentiated Runge's function, $f(x) = \frac{1}{1+25x^2}$, $x \in [-1, 1]$, using the barycentric rational interpolant-based differentiation operator for different values of d_{\max} and other differentiation operators on equispaced nodes (Chebyshev on Chebyshev nodes).

Well-Posedness & Boundary Conditions

In the previous sections, two fundamentally different approaches were presented that each introduce a potential difficulty. The least-squares Fourier projection, in either the global (§4.2) or localized (§4.3.1) formulation, is not an interpolation scheme but it does satisfy the implicit temporal periodic boundary conditions. The bounded interval treatments, such as barycentric rational interpolants or splines, naturally interpolate the data but lack implicit (or explicit) temporal periodic boundary conditions.

In the continuous sense, a node on a background grid transitions between an unblanked state (*a*) to a position where it lies coincident to the solid wall of the impermeable boundary (*b*) and finally, internal to the boundary (*c*). However, the precise location in time of state *b*, where the node is collocated with the solid wall, is in general not one of the discrete temporal collocation points. This fact is demonstrated in Figs. 4.11 and 4.12. The node at times $t = t_1$ and $t = t_5$ is unblanked and representative of state *a*. The node at $t = t_3$ lies within the impermeable boundary and is therefore blanked and an example of state *c*. At precisely $t = t_2$ and $t = t_4$, the node lies coincident with the solid wall of the impermeable boundary. Figure 4.12 demonstrates that while the continuous blanked interval spans I_c , the discrete blanked interval, as detected by the limited sampling of the Time-Spectral collocation scheme, is interpreted as I_d , with no information regarding the location of the solid wall crossings (hashed lines).

In the bounded interval approach, no boundary conditions are provided at the outermost nodes because they are not generally located at solid walls. Instead, the basis spans the interior of I_d only. It is conceivable to construct wall boundary conditions at the crossover points (state *b* corresponding to times $t = t_2$ and $t = t_4$) if their location was discernible from available information. This would require a fine-grained query as part of a preprocessor that would traverse the trajectory of all impermeable boundaries and store the temporal locations of all solid-wall boundaries for each node. Appropriate physical boundary conditions would then have to be derived and implemented. This would also remove the cases where continuously partitioned intervals are spanned by a single discrete partitioned interval. Approaches lacking physical boundary conditions on the bounded interval have demonstrated

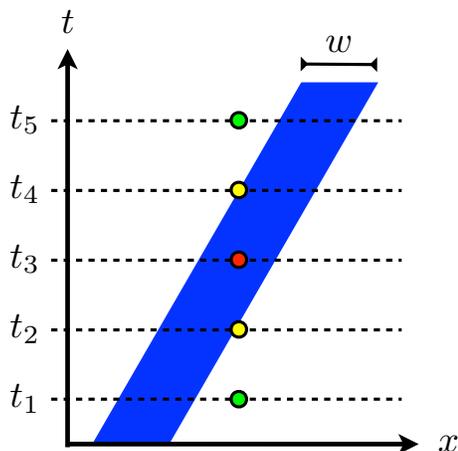


Figure 4.11: A continuous trajectory of a moving solid object with width w and its associated blanked region (blue) demonstrating active (green), solid-wall (yellow) and blanked (red) states for a representative node at different times. The solid boundary has width w and moves from left to right in x . At times $t = t_1$ and $t = t_5$ the node is active (unblanked). At times $t = t_2$ and $t = t_4$ the node coincides with the solid boundary of the moving object. At time $t = t_3$ the node lies within the impermeable boundary of the moving object and is therefore blanked.

success in practice, but numerical stability, and likely accuracy, would be enhanced with the imposition of physical boundary conditions. This is an area of potential future work but is not included in the current scope.

Another potential remedy is the use of a mixed description of the solution, where a least-squares Fourier projection is coupled with a secondary, localized basis such that the combined representation interpolates the data, and completely or partially satisfies the implicit periodic boundary conditions. This approach is detailed in the next section.

4.4 Mixed Expansion of the Solution

This section outlines a mixed approach where the discrete solution is projected by a Fourier least-squares operator, which does not interpolate the data, and *lifted* to match the discrete data with a secondary basis. In this construction, the lifting function has less of an effect as the least-squares projection improves with increased

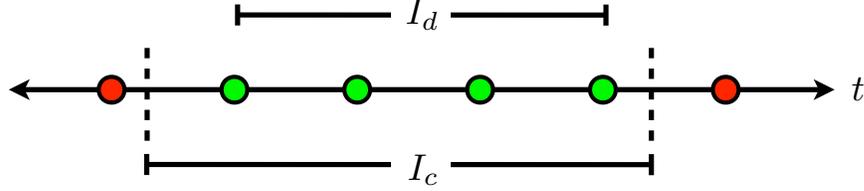


Figure 4.12: Bounded interval for a representative dynamically-blanked spatial node with four consecutively-unblanked collocation points (green). The continuous interval spans I_c whereas the discrete interval, which only detects blanking information at the temporal collocation points, spans I_d . Determining the instantaneous solid-wall crossover points (dashed lines) would require an expensive preprocessing stage.

resolution. With increased resolution, the discrete solution may exist wholly within the Fourier subspace spanned by the orthogonal projection operator; the resulting representation would be purely harmonic. It is not clear if the least-squares projection should include all of the available data points spanning all discrete intervals as described in §4.2 or if multiple least-squares projections should be employed following the description of the periodized approach in §4.3.1.

In Fig. 4.13, a single, global Fourier least-squares approximation projects the discrete solution to the discrete temporal collocation points shown in red. Each independent lifting function only operates within its associated interval. The solution can be expressed as a combination of the Fourier, modal projection, $P_P^N = \Phi (\Phi^* \Phi)^{-1} \Phi^*$, and the nodal lifting bases, Ψ^1 and Ψ^2 , corresponding to the first and second interval of unblanked nodes, respectively.

$$\bar{\mathbf{u}}_P = P_P^N \mathbf{u}_N \in J_P \quad (4.28)$$

$$\mathbf{e}_1 = (I - P_P^N) \mathbf{u}_N \in J_1 \quad (4.29)$$

$$\mathbf{e}_2 = (I - P_P^N) \mathbf{u}_N \in J_2 \quad (4.30)$$

The set J_P defines all of the active nodes independent of the particular interval. Sets J_1 and J_2 refer to the active nodes within intervals 1 and 2, respectively. Therefore, the nodal residual at each active node is simply the difference between the discrete

value and its projection. The lifting functions expand the residual in a secondary basis. Ideally one of the aforementioned boundary conditions would be applied but the example that follows is consistent with the description used in §4.3.2 for a nodal basis such as the barycentric rational interpolants.

$$e_1(t) = \sum_{k \in J_1} e_1(t_k) \psi_k^1(t) \quad (4.31)$$

$$e_2(t) = \sum_{k \in J_2} e_2(t_k) \psi_k^2(t) \quad (4.32)$$

Figure 4.13 demonstrates the concept of a mixed expansion using a global least-squares projection with localized corrections \mathbf{e}_1 and \mathbf{e}_2 within intervals 1 and 2, respectively, by interpolating the projection residual, $\mathbf{e} = \mathbf{u}_N - \bar{\mathbf{u}}_N$, at the nodes with barycentric rational interpolants. It is clear that the corrected representation successfully reconstructs the solution within each partition.

Just as the solution is expanded in a pair of bases, differentiation is supplied by a compound operator. Consider the corrected solution within the first interval.

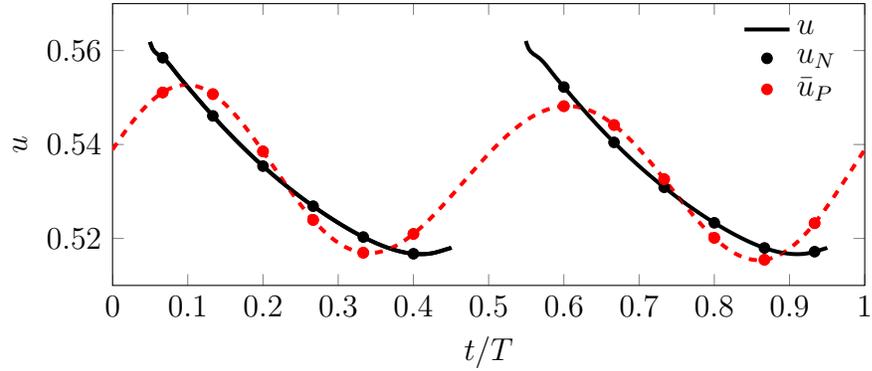
$$u_N(t_j) = \bar{u}_P(t_j) + e_1(t_j), \quad j \in J_1 \quad (4.33)$$

The expression can be written in matrix form

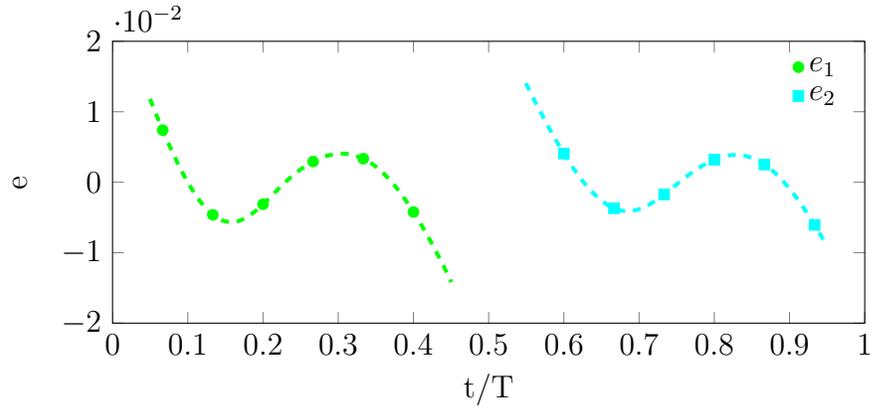
$$\mathbf{u}_N = \Phi \tilde{\mathbf{u}} + \Psi^1 \mathbf{e}_1 \in J_1 \quad (4.34)$$

where $\tilde{\mathbf{u}}_P = (\Phi^* \Phi)^{-1} \Phi^* \mathbf{u}_N$. Taking \mathcal{D}_Ψ^1 as the time-domain differentiation operator for the basis Ψ^1 , and recalling D_P from §4.2, we can differentiate Eq. 4.34 in the following manner.

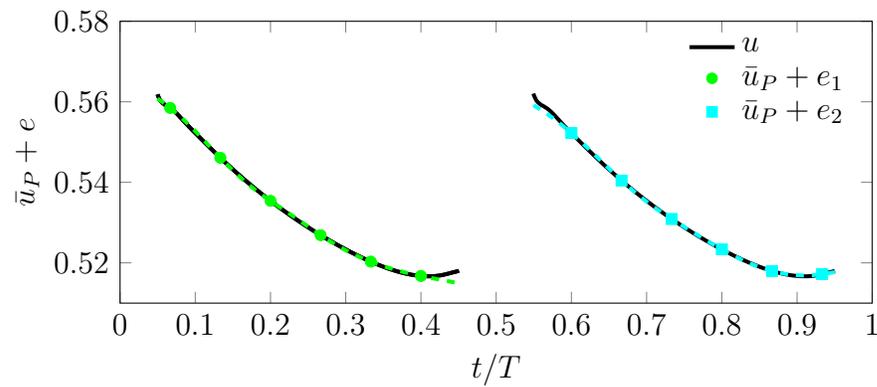
$$\frac{d}{dt} \mathbf{u}_N = \Phi D_P \tilde{\mathbf{u}} + \mathcal{D}_\Psi^1 \mathbf{e}_1 \in J_1 \quad (4.35)$$



(a) Unsteady data and its least-squares Fourier projection $\bar{u}_P = P_P^N u_N$



(b) Lifting functions for each interval determined by interpolating the residual between the solution and its projection, $(I - P_P^N)u_N$



(c) Mixed expansion of the solution, $\bar{u}_P + e_i$, for each interval, i

Figure 4.13: Mixed expansion of the solution with $N = 15$ global collocation points over two intervals; each composed of six unblanked nodes. A global least-squares projection is corrected by localized barycentric rational interpolant lifting functions. The mixed expansion satisfies the interpolation property of the Time-Spectral collocation method.

The compound differentiation acting on the solution vector is expressed by the substitution of $\mathbf{e}_1 = (I - P_P^N) \mathbf{u}_N \in J_1$.

$$\frac{d}{dt} \mathbf{u}_N = \Phi D_P \tilde{\mathbf{u}} + \mathcal{D}_\Psi^1 (I - P_P^N) \mathbf{u}_N \in J_1 \quad (4.36)$$

$$= \Phi D_P (\Phi^* \Phi)^{-1} \Phi^* \mathbf{u}_N + \mathcal{D}_\Psi^1 (I - P_P^N) \mathbf{u}_N \in J_1 \quad (4.37)$$

$$= [\Phi D_P (\Phi^* \Phi)^{-1} \Phi^* + \mathcal{D}_\Psi^1 (I - P_P^N)] \mathbf{u}_N \in J_1 \quad (4.38)$$

$$\frac{d}{dt} \mathbf{u}_N = \mathcal{D}_N \mathbf{u}_N \in J_1 \quad (4.39)$$

From Eq. 4.38, it is clear that the localized differentiation operator is only applied to the discrete projection error, $(I - P_P^N) \mathbf{u}_N$.

The lifting function can in general fall into either the periodized or bounded interval categories. For the examples discussed in the bounded interval case in §4.3.2, no implicit periodic or physical boundary conditions are available (as discussed in the previous section). However, in the case of a lifting function, it is reasonable to *enforce* that it is active in the discretely defined region, I_d , and zero outside. This serves to lift the projected solution such that the combination interpolates the discrete data and goes to zero outside of the bounded interval. However, it can be constrained to do so at only the two endpoints (one collocation point outside of the active region where by definition they are blanked) or it can be set to zero at every collocation point not in the interval. Either option is compatible with dynamically-blanked nodes having multiple partitions. We could also use a periodic lifting function or cluster *ghost* nodes outside of interval to transition the lifting function smoothly to zero. These extensions are considered part of the future work.

The combination of the global and local representations results in a collocation method satisfying the desired interpolation property of the Time-Spectral method, while using a harmonic approximation that reverts to a fully spectral description once the discrete projection lies within the least-squares subspace. This approach fits nicely within the Time-Spectral framework by inclusion of the harmonic basis functions. However, the low-frequency subspace introduces a high-frequency residual. Future investigation may explore strategies to retain some high frequencies in the

least-squares projection to facilitate a more accurate derivative approximation for the lifting functions on less-oscillatory data.

The numerical experiments described in §6.1.1 support the conclusion that the local approach using splines or barycentric rational interpolants is the most accurate.

Chapter 5

OVERFLOW

A primary objective of this research is the development of the computational machinery necessary to solve complex fluid simulations using overset grids with the Time-Spectral method. To validate the hybrid Time-Spectral approach developed in Chapter 4 and employ it to solve realistic flows, the Time-Spectral discretization was implemented within a mature flow solver rather than developing an entirely new codebase. This choice also demonstrates the limited development overhead associated with augmenting existing solvers with Time-Spectral capability. As mentioned in Section 1.1, researchers at Duke University [41, 42, 43] have already augmented the OVERFLOW code with the Harmonic Balance method. However, their applications were limited to simulations involving constant hole cutting. However, Custer [41] cited the need to develop a novel approach for evaluating the temporal derivative at dynamically-blanked nodes to facilitate a general Time-Spectral overset scheme.

This chapter first introduces the governing equations of fluid dynamics that are solved by OVERFLOW. The discussion proceeds with an overview of the standard solver; turbulence modeling and convergence acceleration techniques are addressed. Implementation details for the OVERFLOW Time-Spectral augmentation follow. A brief synopsis of the additional resource requirements (over the steady-state solver) concludes the chapter. All of the OVERFLOW discussion and results presented for both the time-accurate and Time-Spectral implementations are based on the version 2.2e release of the source code.

OVERFLOW is a highly-parallelized, overset, time-accurate Reynolds-averaged Navier-Stokes (RANS) flow solver that has been under development at NASA since the 1970s with roots in the ARC3D code [89]. With a wide user base spanning academia, industry and the government, OVERFLOW has been used tackle a wide variety of problems in fluid dynamics. Because of its accessible codebase, OVERFLOW serves as an apt testbed for the implementation of the Time-Spectral method for both the conventional approach and the proposed hybrid approach to treat overset relative motion.

5.1 Governing Equations

OVERFLOW is an implicit finite-difference-based code that solves the compressible Reynolds-averaged Navier-Stokes (RANS) equations in conservation-law form in time-dependent generalized curvilinear coordinates. The Navier-Stokes equations are statements of the conservation of mass, momentum and energy that can be expressed in strong conservation-law form.

$$\frac{\partial Q}{\partial t} + \nabla \cdot \mathbf{F}(Q) = 0 \quad (5.1)$$

The state vector, $Q = \{\rho, \rho \mathbf{u}, e\}^T$, consists of the N_Q conserved quantities in Cartesian coordinates. The flux vector, \mathbf{F} , is composed of both convective and viscous contributions [See §A.1]. For brevity, only two spatial dimensions are considered, but an extension to three spatial dimensions is straightforward.

In order to solve the discretized equations on arbitrarily-oriented moving meshes, the equations are expressed in generalized curvilinear coordinates. The strong conservation-law form of the governing equations in the physical domain

$$\frac{\partial Q}{\partial t} + \frac{\partial F}{\partial x} + \frac{\partial G}{\partial y} = 0 \quad (5.2)$$

is transformed into strong conservation-law form in the generalized curvilinear coordinate system [See §A.2].

$$\frac{\partial \mathcal{Q}}{\partial t} + \frac{\partial \mathcal{F}}{\partial \xi} + \frac{\partial \mathcal{G}}{\partial \eta} = 0 \quad (5.3)$$

Flux contributions arising from mesh motion are embedded within the transformed flux vectors, \mathcal{F} and \mathcal{G} .

5.2 The Standard OVERFLOW Solver

This section outlines the discretization and relevant solution procedures for the standard OVERFLOW solver. Source-code modifications required to add the Time-Spectral discretization to OVERFLOW are described in §5.3.

5.2.1 Implicit Factored Scheme

OVERFLOW is both a steady-state and time-accurate flow solver with a number of discretization options in both time and space. The time-accurate discretization using the second-order backward difference formula (BDF2) is outlined below for the approximate-factorization (AF) scheme developed by Beam and Warming [90]. The BDF2 scheme will be used in Chapter 6 for all time-accurate simulations. The steady-state discretization is also introduced.

Semi-Discretization

The following statement is a semi-discretization of Eq. 5.3, where δ_ξ and δ_η are general finite-difference operators in the ξ - and η -directions, respectively.

$$\frac{d}{dt} \mathcal{Q} + \delta_\xi \mathcal{F} + \delta_\eta \mathcal{G} = 0 \quad (5.4)$$

Time-Accurate Discretization

In the general case of time-dependent flow, a dual time-stepping scheme [91] is applied to sufficiently reduce the unsteady residual at each physical time step. The time-derivative of the state vector of conserved quantities, \mathcal{Q} , can be discretized by a first- or second-order approximation but the second-order BDF2 approximation is used for all of the calculations in Chapter 6.

$$\frac{d}{dt}\mathcal{Q} = \frac{3\mathcal{Q}^{n+1} - 4\mathcal{Q}^n + \mathcal{Q}^{n-1}}{2\Delta t} + O(\Delta t^2) \quad (5.5)$$

It is important to note that two values of the solution must be stored (Q^n and Q^{n-1}) for the second-order scheme. The fully discrete form of the unsteady governing equations is constructed by substituting Eq. 5.5 into Eq. 5.4.

$$\frac{3\mathcal{Q}^{n+1} - 4\mathcal{Q}^n + \mathcal{Q}^{n-1}}{2\Delta t} + \delta_\xi \mathcal{F}^{n+1} + \delta_\eta \mathcal{G}^{n+1} = 0 \quad (5.6)$$

The dual-time advancement scheme is achieved by adding a pseudotime derivative approximated by the first-order backward Euler scheme and substituting $s+1 \leftarrow n+1$, where s is the subiteration index.

$$\frac{\mathcal{Q}^{s+1} - \mathcal{Q}^s}{\Delta \tau} + \frac{3\mathcal{Q}^{s+1} - 4\mathcal{Q}^n + \mathcal{Q}^{n-1}}{2\Delta t} + \delta_\xi \mathcal{F}^{s+1} + \delta_\eta \mathcal{G}^{s+1} = 0 \quad (5.7)$$

The nonlinear form of the Navier-Stokes flux vectors requires a linearization of \mathcal{F}^{s+1} and \mathcal{G}^{s+1} about Q^s [90].

$$\mathcal{F}^{s+1} = \mathcal{F}^s + \Delta \mathcal{Q} \left. \frac{\partial \mathcal{F}}{\partial \mathcal{Q}} \right|_s + O(\Delta \tau^2) \quad (5.8)$$

$$\mathcal{G}^{s+1} = \mathcal{G}^s + \Delta \mathcal{Q} \left. \frac{\partial \mathcal{G}}{\partial \mathcal{Q}} \right|_s + O(\Delta \tau^2) \quad (5.9)$$

Substitution of the linearizations into Eq. 5.10 results in the following fully-discrete implicit system of equations

$$\frac{Q^{s+1} - Q^s}{\Delta\tau} + \frac{3Q^{s+1} - 4Q^n + Q^{n-1}}{2\Delta t} + \left[\delta_\xi^L \mathcal{A} + \delta_\eta^L \mathcal{B} \right] \Delta Q = - \left[\delta_\xi^R \mathcal{F}^s + \delta_\eta^R \mathcal{G}^s \right] \quad (5.10)$$

The flux Jacobians in the ξ - and η -directions are represented by $\mathcal{A} = \frac{\partial \mathcal{F}}{\partial Q} \Big|_s$ and $\mathcal{B} = \frac{\partial \mathcal{G}}{\partial Q} \Big|_s$, respectively. This general expression accommodates distinct left- and right-hand-side differencing operators (δ^L versus δ^R) to account for stability and efficiency considerations.

Taking $\Delta Q = Q^{s+1} - Q^s$ and subtracting all explicit terms from the left-hand-side (LHS) results in the *delta* form of the equations, where $\Delta\tilde{\tau} = \Delta\tau / (1 + \frac{3\Delta\tau}{2\Delta t})$.

$$\left[I + \Delta\tilde{\tau} \delta_\xi^L \mathcal{A} + \Delta\tilde{\tau} \delta_\eta^L \mathcal{B} \right] \Delta Q = -\Delta\tilde{\tau} \left[\frac{3Q^s - 4Q^n + Q^{n-1}}{2\Delta t} + \delta_\xi^R \mathcal{F}^s + \delta_\eta^R \mathcal{G}^s \right] \quad (5.11)$$

The implicit solution of the equation involves an approximate factorization whereby each direction is solved independently without loss of the global order of convergence. An approximate factorization of the LHS of Eq. 5.11 results in the final form of the system of equations, where $\mathcal{R}(Q^s)$ is the residual vector evaluated by the right-hand-side (RHS) of Eq. 5.11.

$$\left[I + \Delta\tilde{\tau} \delta_\xi^L \mathcal{A} \right] \left[I + \Delta\tilde{\tau} \delta_\eta^L \mathcal{B} \right] \Delta Q = \mathcal{R}(Q^s) \quad (5.12)$$

The system of equations can be written compactly, where \mathcal{L}_ξ and \mathcal{L}_η are the implicit operators in the ξ - and η -directions, respectively.

$$\mathcal{L}_\xi \mathcal{L}_\eta \Delta Q = \mathcal{R}(Q^s) \quad (5.13)$$

These equations are solved by applying a sequence of inversions onto $\mathcal{R}(Q^s)$.

$$\mathcal{L}_\xi \Delta \bar{Q} = \mathcal{R}(Q^s) \quad (5.14)$$

$$\mathcal{L}_\eta \Delta Q = \Delta \bar{Q} \quad (5.15)$$

While the global operators \mathcal{L}_ξ and \mathcal{L}_η are each of dimension $N_\xi \times N_\eta \times N_Q$, each system of equations can be solved in parallel due to the directionality of the operators. The Beam-Warming block tridiagonal solver [90] is used for most of the computations presented; it applies a sequence of direct linear solves of block tridiagonal systems in each spatial direction as part of the AF process in Eq. 5.13. This is accomplished by solving the $N_\xi \times N_Q$ block tridiagonal system at each node in η followed by solving the $N_\eta \times N_Q$ block tridiagonal system at each node in ξ . The diagonalized scheme is also commonly used, which instead applies a sequence of scalar pentadiagonal inversions (See Pulliam and Chaussee [92]). OVERFLOW stores the physical state vector so the solution is advanced by dividing the transformed update by the volume, V .

$$Q_j^{s+1} = Q_j^s + \frac{\Delta Q_j}{V_j} \quad (5.16)$$

After executing the specified number of subiterations or reaching a predetermined convergence threshold, the solution is advanced from t^n to t^{n+1} ,

$$Q^{n-1} \leftarrow Q^n \quad (5.17)$$

$$Q^n \leftarrow Q^{s+1} \quad (5.18)$$

and the subiteration process is initiated at the new physical time step.

The default spatial discretization employs second-order central differences with second- and fourth-difference dissipation operators [93]. The second-difference dissipation is activated by a pressure switch but the fourth-difference dissipation is universally applied. To retain the block tridiagonal structure of the linear system of equations, the LHS fourth-difference artificial dissipation operator, with a stencil width of five, is replaced by a second-difference operator with a stencil width of three. The LHS operator is scaled by twice the fourth-difference coefficient to retain stability. Further details concerning artificial dissipation operators are provided by Pulliam [94] and Pulliam and Zingg [95]. The complete operators are defined in the

following manner.

$$\delta_\xi^L \leftarrow \delta_\xi^C + 2\epsilon_4 \Delta\xi \delta_{\xi\xi}^C \quad (5.19)$$

$$\delta_\xi^R \leftarrow \delta_\xi^C + \epsilon_4 \Delta\xi^3 \delta_{\xi\xi\xi\xi}^C \quad (5.20)$$

For $\mathcal{Q}_{ij} = \mathcal{Q}(i\Delta\xi, j\Delta\eta)$, the particular difference operators are defined as follows.

$$\delta_\xi^C \mathcal{Q}_{i,j} = \frac{1}{2\Delta\xi} (\mathcal{Q}_{i+1,j} - \mathcal{Q}_{i-1,j}) \quad (5.21)$$

$$\delta_{\xi\xi}^C \mathcal{Q}_{i,j} = \frac{1}{\Delta\xi^2} (\mathcal{Q}_{i+1,j} - 2\mathcal{Q}_{i,j} + \mathcal{Q}_{i-1,j}) \quad (5.22)$$

$$\delta_{\xi\xi\xi\xi}^C \mathcal{Q}_{i,j} = \frac{1}{\Delta\xi^4} (\mathcal{Q}_{i+2,j} - 4\mathcal{Q}_{i+1,j} + 6\mathcal{Q}_{i,j} - 4\mathcal{Q}_{i-1,j} + \mathcal{Q}_{i-2,j}) \quad (5.23)$$

The same definitions hold for the η direction by varying the index j . This second-order accurate (in space and time) dual time-stepping procedure is used for the time-accurate calculations presented in Chapter 6. OVERFLOW offers higher-order central-difference discretizations and upwind schemes but the current work employs the default settings described above.

Steady-State Discretization

For steady-state calculations, the discrete physical time derivative in Eq. 5.11 is omitted, resulting in a similar expression that is iterated in pseudotime without the need for subiterations.

$$[I + \Delta\tau \delta_\xi^L \mathcal{A}] [I + \Delta\tau \delta_\eta^L \mathcal{B}] \Delta\mathcal{Q} = -\Delta\tau [\delta_\xi^R \mathcal{F}^s + \delta_\eta^R \mathcal{G}^s] \quad (5.24)$$

Algorithm 1 illustrates the iterative process for a steady-state computation that will be contrasted in §5.3.1 against the augmented Time-Spectral OVERFLOW process outlined in Algorithm 2.

Algorithm 1: OVERFLOW Steady-State Flow Solver

```

Input : Reference grid,  $\mathbf{X}_{\text{ref}}$ , flow and discretization parameters,  $\mathbf{P}$ 
Output: Solution,  $\mathbf{Q}$ 

// Iterate NSTEP time steps
for  $ISTEP \leftarrow 1$  to  $NSTEP$  do
    // Loop over NG grids
    for  $IG \leftarrow 1$  to  $NG$  do
         $\mathbf{Q} \leftarrow \text{BC}(\mathbf{Q}, \mathbf{X}, \mathbf{M}, \mathbf{IB}, \mathbf{S}, \mathbf{P})$            // Apply physical BC
         $\mathbf{S} \leftarrow \text{RHS}(\mathbf{Q}, \mathbf{X}, \mathbf{M}, \mathbf{IB}, \mathbf{S}, \mathbf{P})$        // Apply standard RHS
         $\mathbf{S} \leftarrow \text{LHS}(\mathbf{Q}, \mathbf{X}, \mathbf{M}, \mathbf{IB}, \mathbf{S}, \mathbf{P})$    // Apply standard LHS
         $\mathbf{Q} \leftarrow \mathbf{Q} + \mathbf{S}$                                    // Update solution
         $\mathbf{S} \leftarrow \mathbf{0}$                                        // Zero out S array
    end
     $\mathbf{Q} \leftarrow \text{CBCXCH}(\mathbf{Q})$                                // Overset and MPI data exchange
end
return

```

5.2.2 Multigrid Acceleration

Multigrid is a powerful acceleration technique designed to improve convergence rates by damping projections of low-frequency errors more rapidly on coarser grids where they are represented as higher-frequency signals. A comprehensive overview of multigrid methods is provided by Wesseling [96] and OVERFLOW-specific treatment is outlined in [97]. Typical calculations involve three-level multigrid in combination with coarse-grid sequencing, referred to herein as full multigrid (FMG). The solution is computed initially on a series of coarse grids to more rapidly damp the initial transient. The Navier-Stokes equations are nonlinear and therefore OVERFLOW employs the Full Approximation Storage (FAS) technique.

5.2.3 Turbulence Model

Turbulent phenomena play a crucial role in many of the applications of interest. Despite the ever-increasing computational resources available, the cost of resolving the entire spectrum is grossly prohibitive for most realistic flows. Turbulence modeling

is a rich field that has provided significant improvements in both theoretical understanding and computational implementation.

OVERFLOW employs a loosely-coupled turbulence model whereby the turbulent and fluid variables are updated sequentially. Within an iteration, the turbulent variables are integrated first and the turbulent update is used to advance the fluid equations. Several turbulence treatments are available, ranging from algebraic models to one- and two-equation models. All of the time-accurate and Time-Spectral calculations presented employ the well-established Spalart-Allmaras (SA) one-equation turbulence model. Readers are directed to [98] for a detailed derivation and explanation of the model.

5.2.4 Connectivity

Intergrid connectivity is required for overset communication and domain decomposition. Both cases are handled by the same process within OVERFLOW. Domain decomposition balances the workload among the processors resulting in grids that are split into separate logical entities. Grid units are distributed to the available processors and solved independently. Hole-cutting facilitated by object X-Rays [63].

5.3 The Time-Spectral OVERFLOW Solver

Discretization of the augmented Time-Spectral OVERFLOW solver is outlined starting from the development presented in §5.2 for the standard OVERFLOW solver. The Time-Spectral implementation shares the majority of the codebase with the steady OVERFLOW solver. The primary difference between the two procedures is that the Time-Spectral algorithm requires that that N states of the solution be stored, corresponding to the N equispaced temporal collocation points; in contrast, the BDF2 time-accurate scheme only stores two states of the solution. An additional array of the same dimension must be allocated to store the temporal residual, $\mathcal{D}_N Q^s$, at every node. Additionally, the spatial residuals and linear solves in the spatial dimensions must be executed at each time level and finally an additional implicit solve must be

executed in the temporal dimension for every grid point. Multigrid and turbulence modeling implications are described, and an estimate is provided of the additional computational resources overhead for the Time-Spectral implementation above its steady-state analogue.

5.3.1 Implicit Factored Scheme

The Time-Spectral discretization is a modification to the discretization of the time-derivative in Eq. 5.6 that replaces an algebraically-accurate finite-difference approximation with an infinitely-supported, spectrally-accurate differentiation operator, \mathcal{D}_N . No subiterations are required for the modified discretization because the solution converges to a steady state in the combined space-time domain.

$$\mathcal{D}_N \mathcal{Q}^{s+1} + \delta_\xi \mathcal{F}^{s+1} + \delta_\eta \mathcal{G}^{s+1} = 0 \quad (5.25)$$

In the Time-Spectral formulation, the solution and flux vectors \mathcal{Q} , \mathcal{F} and \mathcal{G} are now taken over both the spatial and temporal dimensions. Therefore, the spatial and temporal differentiation operators are constructed accordingly, as demonstrated explicitly by Thomas et al. [43] and Naik et al. [29].

As with the time-accurate case, the nonlinear flux vectors are linearized as provided by Eqs. 5.8-5.9 and distinct spatial discretization operators are specified for the right- and left-hand sides whose definitions are those specified in Eqs. 5.19-5.20. The temporal differentiation operator is linear; therefore, no linearization is required.

$$\mathcal{D}_N \mathcal{Q}^{s+1} + [\delta_\xi^L \mathcal{A} + \delta_\eta^L \mathcal{B}] \Delta \mathcal{Q} = - [\delta_\xi^R \mathcal{F}^s + \delta_\eta^R \mathcal{G}^s] \quad (5.26)$$

As with the time-accurate case, a backward Euler pseudotime advancement scheme is applied to Eq. 5.26.

$$\frac{\mathcal{Q}^{s+1} - \mathcal{Q}^s}{\Delta \tau} + \mathcal{D}_N \mathcal{Q}^{s+1} + [\delta_\xi^L \mathcal{A} + \delta_\eta^L \mathcal{B}] \Delta \mathcal{Q} = - [\delta_\xi^R \mathcal{F}^s + \delta_\eta^R \mathcal{G}^s] \quad (5.27)$$

Equation 5.27 is put into *delta* form, where $\Delta \mathcal{Q} = \mathcal{Q}^{s+1} - \mathcal{Q}^s$, by subtracting $\mathcal{D}_N \mathcal{Q}^s$

from both sides.

$$[I + \Delta\tau\delta_\xi^L\mathcal{A} + \Delta\tau\delta_\eta^L\mathcal{B} + \Delta\tau\mathcal{D}_N] \Delta\mathcal{Q} = -\Delta\tau [\delta_\xi^R\mathcal{F}^s + \delta_\eta^R\mathcal{G}^s + \mathcal{D}_N\mathcal{Q}^s] \quad (5.28)$$

Equation 5.28 can be approximately-factored into $N_{sd}+1$ directional operators, where $\mathcal{R}(\mathcal{Q}^s)$ is the residual vector evaluated by the RHS of Eq. 5.28.

$$[I + \Delta\tau\delta_\xi^L\mathcal{A}] [I + \Delta\tau\delta_\eta^L\mathcal{B}] [I + \Delta\tau\mathcal{D}_N] \Delta\mathcal{Q}_j = \mathcal{R}(\mathcal{Q}^s) \quad (5.29)$$

This system is expressed compactly,

$$\mathcal{L}_\xi\mathcal{L}_\eta\mathcal{L}_t\Delta\mathcal{Q} = \mathcal{R}(\mathcal{Q}^s), \quad (5.30)$$

where the innermost operator, \mathcal{L}_t , is the implicit temporal differentiation operator applied to the time-history of the solution at a given spatial location, and the spatial operators \mathcal{L}_ξ and \mathcal{L}_η are unchanged from the standard OVERFLOW solver. These equations are solved by applying a sequence of $N_{sd} + 1$ inversions onto $\mathcal{R}(\mathcal{Q}^s)$.

$$\mathcal{L}_\xi\Delta\bar{\bar{\mathcal{Q}}} = \mathcal{R}(\mathcal{Q}^s) \quad (5.31)$$

$$\mathcal{L}_\eta\Delta\bar{\bar{\mathcal{Q}}} = \Delta\bar{\bar{\mathcal{Q}}} \quad (5.32)$$

$$\mathcal{L}_t\Delta\mathcal{Q} = \Delta\bar{\bar{\mathcal{Q}}} \quad (5.33)$$

As with the time-accurate case, the sequence of three $N_\xi \times N_\eta \times N \times N_Q$ systems of equations are not solved directly. This update is accomplished by looping over each of the N time samples and solving the $N_\xi \times N_Q$ block tridiagonal system at each node in η , followed by solving the $N_\eta \times N_Q$ block tridiagonal system at each node in ξ . The final update involves looping over every node in ξ and η and performing N_Q direct inversions of dimension N . The physical solution is advanced by the dividing the transformed update by the volume, V , as in the standard solver.

$$Q_j^{s+1} = Q_j^s + \frac{\Delta\mathcal{Q}_j}{V_j} \quad (5.34)$$

The required modifications to the existing solver are limited to an additional linear solve of at most dimension N at every grid point for the Time-Spectral approximately-factored operator, and an evaluation of the temporal residual term ($\mathcal{D}_N \mathcal{Q}^s$) at every grid point. In other words, time is treated in a similar fashion as the spatial independent variables when solving for the steady-state solution in the combined space-time domain.

For spatial nodes with complete time histories, the standard Fourier-based differentiation operator is used. However, for dynamically-blanked nodes, a temporal differentiation operator of dimension $N_P < N$ is dynamically generated and used for the intervals of N_P consecutively-defined time samples (for the bounded interval approach). For statically-blanked (Fourier-based) nodes, the approximately-factored Time-Spectral implicit temporal operator, \mathcal{L}_t , can be diagonalized by application of the discrete Fourier transform if $\Delta\tau$ is constant for each time-sample ($\Delta\tau$ is still free to change in space to enable local pseudotime stepping). Thus, Eq. 5.33 is reduced from an implicit linear system of equations to a series of scalar equations removing the need for an inversion of the dense LHS operator. Replacing the direct linear solution with a fast Fourier transform (FFT) and inverse fast Fourier transform (IFFT) reduces the complexity from $\mathcal{O}(N^3)$ to $\mathcal{O}(N \log N)$.

$$\begin{aligned}\tilde{\mathcal{R}} &\leftarrow \text{FFT} \left(\Delta \bar{\mathcal{Q}} \right) \\ \Delta \tilde{\mathcal{Q}}_k &\leftarrow \frac{\tilde{\mathcal{R}}_k}{1 + \Delta\tau i\omega k}, \quad \forall k \\ \Delta \mathcal{Q} &\leftarrow \text{IFFT} \left(\Delta \tilde{\mathcal{Q}} \right)\end{aligned}$$

Similarly, explicit evaluation of $\mathcal{D}_N \mathcal{Q}^s$ on the RHS can be evaluated by applying the FFT to \mathcal{Q}^s , differentiating in the frequency domain by multiplication of $i\omega k$ and applying an IFFT to the differentiated data. This results in an implicit Time-Spectral scheme that matches the overhead of the explicit NLFD [11] scheme for the temporal terms. If the time steps vary for each time sample then the standard approach of performing a direct solve of $[I + \Delta\tau \mathcal{D}_N]$ can still be used. This is also required for dynamically-blanked nodes because there exists no similarity transform of the default rational interpolant-based differentiation operator to the author's knowledge.

The memory requirements for a Time-Spectral solution are greater than its time-accurate analogue because the solution and residual must be stored for each time sample. However, careful implementation limits the storage requirements and allows for a significant number of temporal modes, even for complex three-dimensional problems. Details of memory scaling for the OVERFLOW Time-Spectral solver are outlined in [99] and suggest that a large number of temporal modes can be applied without risk of exhausting the memory budget on a suitably parallelized calculation. Mavriplis and Yang [37] parallelized their Time-Spectral solver in both time and space; this is certainly a topic of interest and a focus of future work. It would be beneficial to add processors to account for a calculation involving more temporal modes while retaining the same spatial domain decomposition.

The general OVERFLOW Time-Spectral solution procedure is outlined in Algorithm 2 and its preprocessor is detailed in Algorithm 3. The Time-Spectral process collapses to OVERFLOW's steady-state solver outlined in Algorithm 1 when a single time sample is chosen. Naturally this corresponds to a steady, zero-mode calculation ($K = (N - 1) / 2 = 0$) with a null temporal differentiation operator. Therefore, neither the temporal residual evaluation nor the implicit solve in the temporal dimension are required in this case.

5.3.2 Multigrid Acceleration

Fully-implicit spatial multigrid acceleration has been incorporated into the Time-Spectral OVERFLOW implementation by leveraging the existing multigrid routines of the OVERFLOW solver. This was accomplished by including the Time-Spectral temporal residual evaluations (RHS) and linear solves (LHS) on all coarse-grid levels, facilitated by appropriate restriction and prolongation of relevant data at each of the N collocation points. The multigrid implementation has been validated on single and overset grids for both rigid and relative motion.

Figure 5.1 compares the convergence of the complete space-time residual (Figure 5.1a) and first harmonic of drag coefficient (Figures 5.1b & 5.1c) for a three-level FMG calculation with coarse-grid sequencing to that for a single grid computation. The

Algorithm 2: OVERFLOW Time-Spectral Flow Solver

```

Input : Time Samples,  $N$ , reference grid,  $\mathbf{X}_{\text{ref}}$ , flow, discretization and
          motion parameters,  $\mathbf{P}$ 
Output: Solution,  $\mathbf{Q}$ 

// Initialize all arrays with Time-Spectral preprocessor
preprocessTS( $N$ ,  $\mathbf{Q}$ ,  $\mathbf{X}$ ,  $\mathbf{X}_{\text{ref}}$ ,  $\mathbf{M}$ ,  $\mathbf{IB}$ ,  $\mathbf{S}$ ,  $\mathbf{P}$ ,  $\mathbf{IBTS}$ )

// Iterate NSTEP time steps
for  $ISTEP \leftarrow 1$  to  $NSTEP$  do
  // Loop over NG grids
  for  $IG \leftarrow 1$  to  $NG$  do
    // Loop over spatial nodes on grid IG
    for  $J \leftarrow 1$  to  $NPTS$  do
       $\mathbf{S} \leftarrow \text{RHSTS}(\mathbf{Q}, \mathbf{IBTS}_J)$  // Compute temporal residual and
      store in  $\mathbf{S}$  for each node  $j$ 
    end
    // Loop over N time-levels
    for  $I \leftarrow 1$  to  $N$  do
       $\mathbf{Q}_I \leftarrow \text{BC}(\mathbf{Q}_I, \mathbf{X}_I, \mathbf{M}_I, \mathbf{IB}_I, \mathbf{S}_I, \mathbf{P})$  // Apply physical BC
       $\mathbf{S}_I \leftarrow \text{RHS}(\mathbf{Q}_I, \mathbf{X}_I, \mathbf{M}_I, \mathbf{IB}_I, \mathbf{S}_I, \mathbf{P})$  // Apply standard RHS
       $\mathbf{S}_I \leftarrow \text{LHS}(\mathbf{Q}_I, \mathbf{X}_I, \mathbf{M}_I, \mathbf{IB}_I, \mathbf{S}_I, \mathbf{P})$  // Apply standard LHS
    end
    // Loop over spatial nodes on grid IG
    for  $J \leftarrow 1$  to  $NPTS$  do
       $\mathbf{S} \leftarrow \text{LHSTS}(\mathbf{S}, \mathbf{IBTS}_J, \mathbf{M})$  // Apply final implicit operator
    end
    // Loop over N time-levels
    for  $I \leftarrow 1$  to  $N$  do
       $\mathbf{Q}_I \leftarrow \mathbf{Q}_I + \mathbf{S}_I$  // Update solution
    end
  end
  // Loop over N time-levels
  for  $I \leftarrow 1$  to  $N$  do
     $\mathbf{Q}_I \leftarrow \text{CBCXCH}(\mathbf{Q}_I, \mathbf{IB}_I)$  // Overset and MPI data exchange
  end
end
return

```

Algorithm 3: Time-Spectral Preprocessor

```

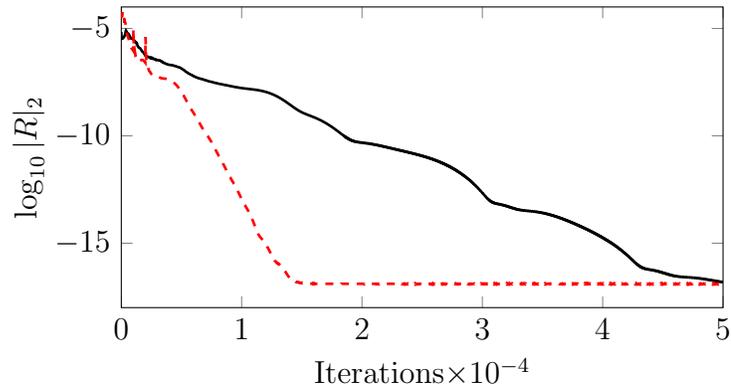
Subroutine preprocessTS( $N, \mathbf{Q}, \mathbf{X}, \mathbf{X}_{\text{ref}}, \mathbf{M}, \mathbf{IB}, \mathbf{S}, \mathbf{P}, \mathbf{IBTS}$ )
  // Loop over time samples to move grid, compute metrics and
  // evaluate iblanks
  for  $I \leftarrow 1$  to  $N$  do
    |  $t_I \leftarrow (I - 1)T/N$ 
    |  $\mathbf{X}_I \leftarrow \text{moveGrid}(\mathbf{X}_{\text{ref}}, \mathbf{P}, t_I)$  // Move grid
    |  $\mathbf{M}_I \leftarrow \text{computeMetrics}(\mathbf{X}_I, \mathbf{P})$  // Compute metrics
    |  $\mathbf{IB}_I \leftarrow \text{computeIblanks}(\mathbf{X}_I)$  // Evaluate iblanks
  end
  // Process time-histories of  $\mathbf{IB}$  for each node  $j$ 
  for  $J \leftarrow 1$  to  $NPTS$  do
    |  $\mathbf{IBTS}_J \leftarrow \text{processIblanks}(\mathbf{IB})$ 
  end
  return

```

space-time residual converged to machine zero in under one third of the iterations, and the first harmonic of c_d converged in about one third of the iterations. Convergence acceleration using Time-Spectral multigrid is problem dependent but these statistics serve as a representative example of observed results. The ability to employ the accelerated multigrid algorithm enables the investigation of more complex problems that would otherwise prove intractable. Only spatial multigrid has been implemented but adding multigrid in the temporal dimension is included as part of future work.

5.3.3 Temporal Vanishing Viscosity

Periodic approximate factorization is known to be unstable in three or more dimensions. A Time-Spectral calculation in two spatial dimensions results in a three-dimensional space-time approximate-factorization, which is therefore unstable. The Time-Spectral method can be thought of as using a spectrally-accurate, purely-imaginary (non-dissipative), central difference operator in time. Because of this lack of dissipation, one can anticipate instabilities that are not damped by dissipative eigenvalues in the other *spatial* dimensions. Linear stability analysis in Appendix B



(a) Space-time residual

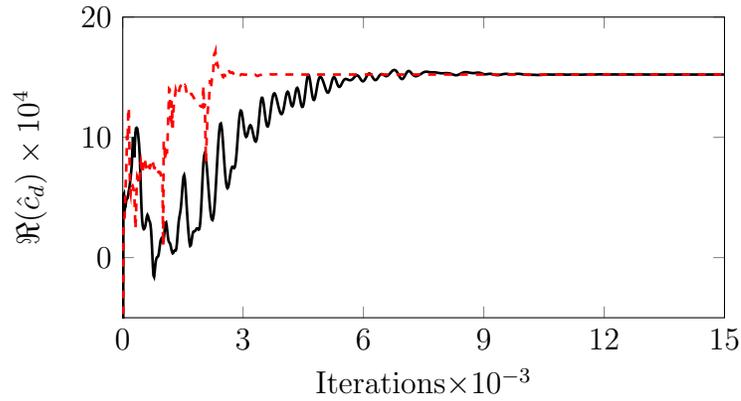
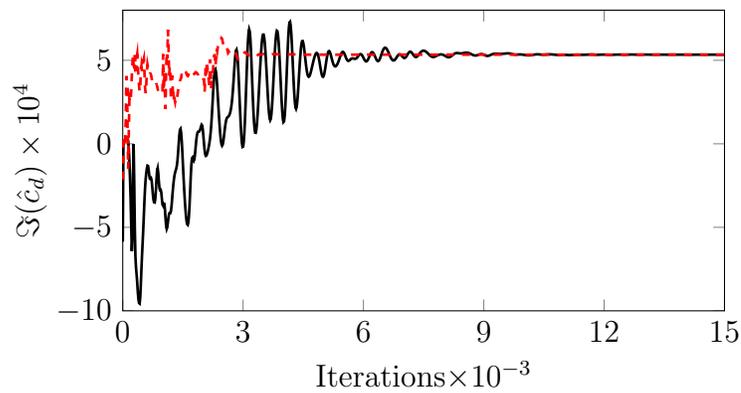
(b) Real component of c_d for first harmonic(c) Imaginary component of c_d for first harmonic

Figure 5.1: Multigrid acceleration for a representative calculation. Convergence of single grid (black line) and three-level FMG (hashed red line) calculations of a laminar plunging airfoil case with $St = 0.288$.

demonstrates the conditional stabilization of the three- and four-factor Time-Spectral AF schemes by applying artificial dissipation to only the spatial dimensions, therefore maintaining the spectral accuracy of the temporal differentiation operator.

Because the Time-Spectral method is a pseudospectral scheme that employs the discrete Fourier Transform, aliasing errors can be introduced that may destabilize calculations, as described in section §2.2. Convergence stalling has been observed for calculations involving significant high-frequency content. For such cases, spectral vanishing viscosity (SVV) is applied, as described in §2.2, on all grid levels.¹

As a modal dissipation operator, SVV is only applied at statically-blanked (Fourier-based) nodes. It is sometime necessary to apply added temporal dissipation at the dynamically-blanked nodes when using a local basis such as the barycentric rational interpolants. In contrast to the modal complex exponential basis set upon which SVV is based, the nodal basis of the barycentric rational interpolant can not be decomposed into modal components to damp a selection of the highest-frequencies. Instead, a dissipation operator derived from barycentric rational interpolants (scaled in the same manner as SVV) is applied at the two interval boundary nodes, which has shown to sufficiently stabilize otherwise unstable calculations. This approach avoids sacrificing accuracy within the interior of the partitioned temporal domain.

5.3.4 Turbulence Model

As previously mentioned, OVERFLOW employs a loosely coupled turbulence scheme whereby the turbulent variable is updated initially and held constant for the flow-equation iteration. In the current work, the one-equation Spalart-Allmaras turbulence model, with the turbulent working variable of undamped eddy viscosity, $\tilde{\nu}$, is used for all calculations presented. The code structure prevents a direct Time-Spectral implicit treatment for the turbulent variable without a significant overhaul. Instead, a semi-implicit process retroactively applies the implicit operator to the temporally-explicit turbulent update, $\Delta\tilde{\nu}^{s+\frac{1}{2}}$. Initially, the turbulence equation is updated implicitly in the spatial dimensions where the Time-Spectral temporal residual, $\mathcal{D}_N\tilde{\nu}^s$ is explicitly

¹McMullen [11] employed a binary SVV kernel on coarse grids to further accelerate NLFDM multi-grid calculations without sacrificing accuracy on the finest mesh.

added to the RHS.

$$[I + \Delta\tau\mathcal{A}][I + \Delta\tau\mathcal{B}]\Delta\tilde{\nu}^{s+\frac{1}{2}} = -\Delta\tau[\mathcal{R}(\tilde{\nu}^s) + \mathcal{D}_N\tilde{\nu}^s] \quad (5.35)$$

The intermediate solution is then held fixed while iterating the fluid variables from \mathcal{Q}^s to \mathcal{Q}^{s+1} .

$$\mathcal{Q}^{s+1} = \mathcal{Q}^s + \Delta\mathcal{Q}\left(\tilde{\nu}^{s+\frac{1}{2}}\right) \quad (5.36)$$

The implicit update $\Delta\tilde{\nu}^{s+1}$ is computed and retroactively applied to $\tilde{\nu}^s$ to advance the solution prior to the subsequent time step.

$$[I + \Delta\tau\mathcal{D}_N]\Delta\tilde{\nu}^{s+1} = \Delta\tilde{\nu}^{s+\frac{1}{2}} \quad (5.37)$$

This approach has proved successful in practice.

A second factor concerning the Time-Spectral treatment of the turbulence equation is the necessity to maintain the positivity of $\tilde{\nu}$. The time-derivative is added to the turbulent residual prior to the implicit spatial solves. A limiter is employed to ensure that the explicit temporal treatment does not violate positivity. This is addressed in more detail in §6.1.3.

5.3.5 Connectivity

The Time-spectral implementation leverages the connectivity software from the standard OVERFLOW solver. Blanking information is evaluated in the preprocessor stage (See Algorithm 3); time histories of the IBLANK array are processed to determine the intervals over which the solution is defined at every dynamically-blanked node. This information is then used to assign the appropriate differentiation operator depending on the basis and distribution of unblanked temporal collocation points at a given spatial node.

5.3.6 Additional Resource Requirements

Naturally, there are both storage and computational overhead costs associated with a Time-Spectral calculation above the steady-state requirements. The extra cost is offset by the capability of converging to a space-time solution more rapidly than a time-accurate calculation. The efficiency merits of Fourier pseudospectral methods have been demonstrated in great detail by McMullen [11, 13]. Therefore, only the costs specific to the current implementation will be described.

Memory

Complete space-time solution and residual arrays must be stored. If N_x is the number of spatial nodes and N_Q is the number of solution variables, then the total memory for the global solution and residual arrays are each $N \times N_x \times N_Q$. Metric terms are also stored for each of the N time samples to avoid recomputing them at every iteration. The metric array is $13 \times N_x$ (in three dimensions), so an additional N copies are stored. The **IBLANK** array, **IB**, is also stored at each time level. The **IB** array simply stores a single integer at every spatial node. However, if memory is limited then a low-storage option can be used where the **IBLANK** and metric arrays are reevaluated at every iteration. Figure 5.2 demonstrates that a computation using 125 million nodes distributed across fifty 20-core Ivy Bridge nodes can still employ upwards of one hundred time samples, which is beyond what is typically used in practice.

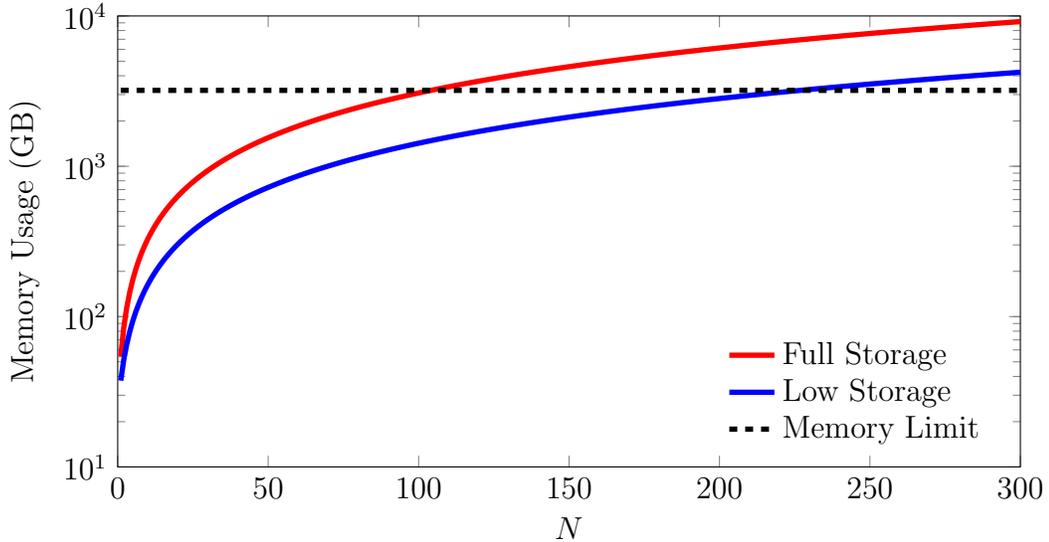


Figure 5.2: Estimated OVERFLOW Time-Spectral memory usage. Memory usage estimate in gigabytes (GB) versus N for fifty Ivy Bridge twenty-core nodes for a sample problem of 125 million mesh points. Each Ivy Bridge node on the NASA Pleiades supercomputer is equipped with 64GB of memory.

Computation

The primary computational overhead resulting from the implicit approximately-factored Time-Spectral solver is attributed to the evaluation of the temporal residual and the linear solve in the temporal dimension at every mesh point. There are other processes such as loading and unloading data as the process loops through each time sample, but their impact should be minimal. As mentioned in §5.3.1, the $\mathcal{O}(N \log N)$ fast Fourier transforms can be used to reduce the computational cost of both the residual computation and the linear solve if constant time-stepping is used at a particular spatial node (again, the time step is free to vary in space). Dynamically-blanked nodes cannot employ the FFT and therefore require a matrix-vector product and a linear solve, but the number of dynamically-blanked nodes is generally a tiny fraction of the overall number of spatial degrees of freedom. The matrix-vector product and linear solves can also be used for the Fourier-based nodes, but the savings afforded by the FFT become quite significant as the number of temporal modes increases.

Chapter 6

Numerical Results

This chapter presents the numerical results obtained with the standard and augmented Time-Spectral OVERFLOW solvers, applied to flows in two and three spatial dimensions. The first section investigates two-dimensional oscillating airfoils to validate the proposed approach. The analysis includes subsonic inviscid and laminar plunging airfoils in addition to transonic inviscid and turbulent RANS pitching airfoils. Accuracy of the various treatments for the dynamically-blanked nodes are addressed. Additionally, the topics of aliasing, resolution and turbulence modeling are discussed. Finally, the isolated V-22 Osprey tiltrotor test cases of hover and forward (edgewise) flight are included to demonstrate the ability of the Time-Spectral method to successfully predict realistic three-dimensional flows.

Computational efficiency considerations of the Time-Spectral method are not addressed. The literature provides strong support for Fourier collocation schemes due to their ability to compute periodic steady-state flows up to an order of magnitude more rapidly than traditional time-marching schemes. Instead, the primary focus of the upcoming numerical experiments is to demonstrate the ability of the standard and hybrid Time-Spectral schemes to produce results in agreement with the time-accurate solver. Well-resolved time-accurate calculations are considered as the exact solutions and are therefore used to assess the accuracy of the Time-Spectral results.

The chapter begins in §6.1.1 with a subsonic inviscid plunging NACA 0012 airfoil test case. This relatively innocuous case serves as the main testbed for analysis

because it is relatively inexpensive. Choosing a large plunging amplitude results in a large region of dynamically-blanked nodes that are swept out by the airfoil. Results using time-accurate, standard Time-Spectral (absent relative motion) and the hybrid Time-Spectral schemes will be presented and compared. A comparison of frequency content resulting from rigid and relative-body motion is provided, suggesting that more frequencies are fundamentally required to resolve the flow when grids move relative to each other. Next, analysis of a pair of low-speed laminar plunging cases is presented in §6.1.2 to investigate the dependence of Time-Spectral convergence on the fundamental frequency, ω , of the oscillation and the need to damp aliasing errors for large ω . Inviscid and turbulent RANS transonic pitching NACA 0012 airfoil test cases follow in §6.1.3. Finally, three-dimensional calculations of the quarter-scale V-22 Tilt Rotor Aeroacoustic Model (TRAM) are provided for both hover and forward flight configurations in §6.2.

6.1 Two-Dimensional Oscillating Airfoils

A series of oscillating airfoil simulations have been used to both validate the Time-Spectral OVERFLOW implementation and analyze the ramifications of solving such problems using rigid- or relative-body motion with respect to the different Time-Spectral approaches. These problems can be solved using rigid-body motion, and correspondingly the standard Time-Spectral method, because only a single component is included. However, this is not the case in general; these relatively simple test cases are selected to compare the relative-motion results computed with the hybrid Time-Spectral scheme directly against their rigid-motion analogues computed with the standard Time-Spectral method.

6.1.1 Inviscid Plunging NACA 0012 Airfoil

A large-amplitude inviscid plunging NACA 0012 airfoil test case provides a meaningful demonstration of the hybrid Time-Spectral scheme as it results in a large number of dynamically-blanked nodes. This section first examines the accuracy of the schemes

discussed in Chapter 4 by evaluating how closely the Time-Spectral solution approximates the time-accurate solution at a selection of representative nodes. Next, modal convergence in the force and moment coefficients is demonstrated for both the rigid- and relative-motion cases.

The instantaneous vertical displacement, $y(t) = h \sin(kt)$, is defined as a single sinusoid of the reduced frequency, $k = \omega c/V_\infty$. The plunging amplitude is selected as half the chord length, resulting in a nondimensional plunging amplitude, $h = a/c = 0.5$. A reduced frequency of $k = 0.1627$ radians per nondimensional time unit was selected. The freestream Mach number is chosen as $M_\infty = 0.5$ to maintain subsonic flow throughout the domain. A near-body O-mesh with 241×30 points in the chordwise and body-normal directions, respectively, is embedded within a 341×261 rectilinear off-body grid that stretches approximately $100c$ to the farfield boundaries in both the x - and y -directions. The hole cut is targeted at $0.1c$ off the solid wall of the airfoil.

Figure 6.1 delineates the portion of the domain that undergoes dynamic blanking for the plunging airfoil case with $N = 9$ time samples. The airfoil sweeps through the red-colored region resulting in blanked spatial nodes for some subset of the period of motion. For example, in the bounded interval approach, the temporal collocation points of these spatial nodes are partitioned into intervals spanning fewer than $N = 9$ time samples and differentiated with a barycentric rational interpolant- or spline-based operator. Nodes in the blue-colored region remain unblanked for the duration of the period and are therefore differentiated by the Fourier-based operator. The length of the time interval associated with each dynamically-blanked point varies and is plotted in Fig. 6.1b for the temporal collocation point at $t = 0$. The white region shows the nodes blanked by the airfoil grid in its neutral position. Nodes colored blue have complete time histories and correspond to the blue-colored nodes in Fig. 6.1a. The solution at these nodes can be expanded in a Fourier series and are differentiated by the Fourier operator. Nodes colored cyan, green, yellow and red are blanked for a portion of the period and only have defined solutions for a subset of the time samples. For example, nodes colored green are included in temporal intervals spanning seven temporal collocation points including the $t = 0$ time sample.

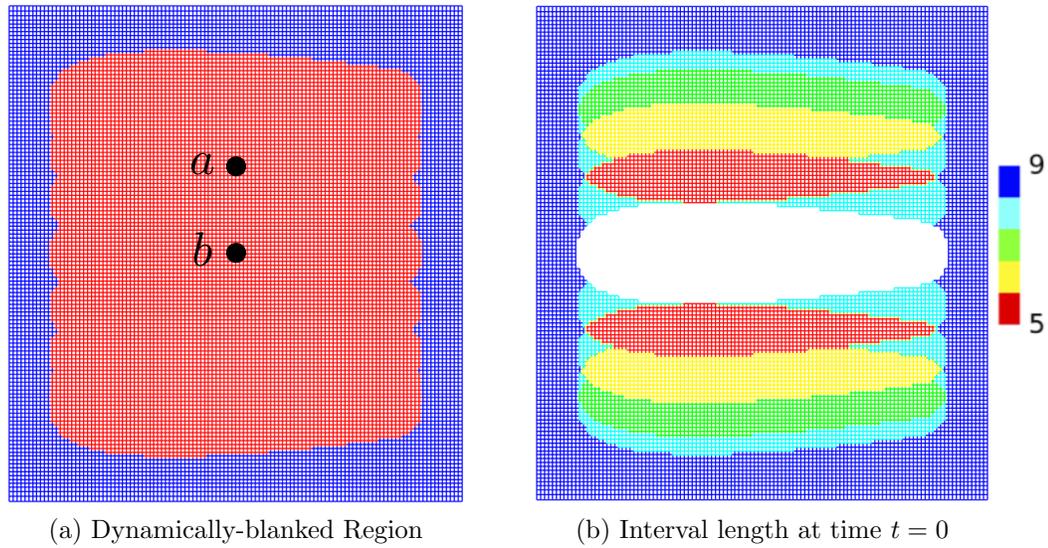
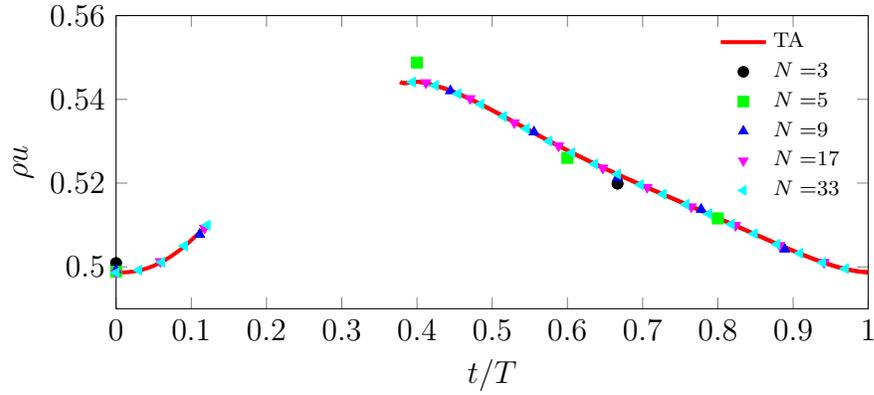
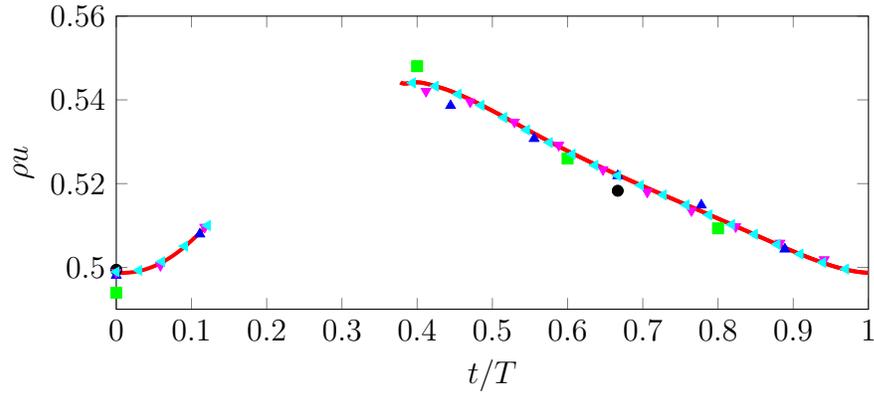


Figure 6.1: Inviscid Plunging Airfoil. Dynamically-blanked region on the stationary background mesh. (a) The Fourier differentiation operator is applied in the blue region where spatial nodes have complete time histories. Alternative differentiation operators are applied to the dynamically-blanked nodes in the red region. Labels a and b locate nodes whose solutions are plotted in Figs. 6.2 and 6.4, respectively. (b) Length of time interval associated with each spatial node at the $t = 0$ temporal collocation point for the $N = 9$ case. The white region corresponds to the hole cut.

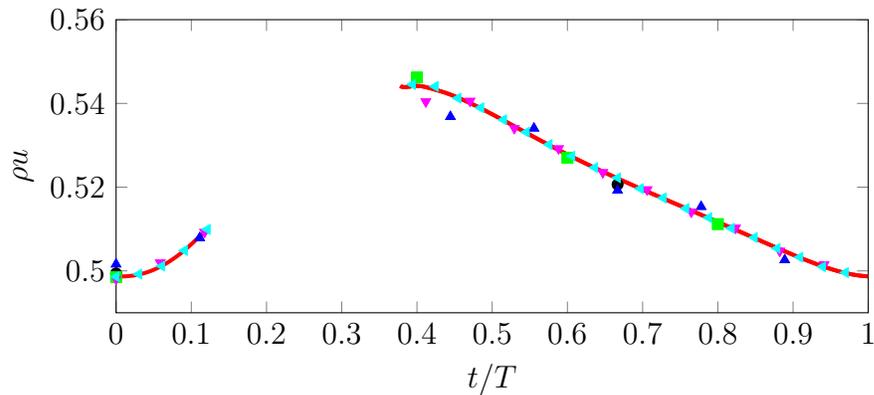
The distribution of interval length changes as the number of modes varies but this example is representative of the general case.

Figures 6.2 and 6.4 track the time history of streamwise momentum at dynamically-blanked nodes a and b (as indicated in Figure 6.1a), respectively, for several treatments. Streamwise momentum is the most sensitive of the conserved quantities at the nodes selected and will therefore be used to analyze each scheme. Node a is located above the mid-chord position of node b , and is only blanked briefly. The airfoil initially blanks node b as it moves upwards and blanks it again as it makes the downward stroke. The node is blanked again as the airfoil completes its cycle. Accuracy is measured by computing the root-mean square (RMS) error of the Time-Spectral versus time-accurate solution at the unblanked collocation points. Convergence for the conserved quantities of density, streamwise and normal momentum is plotted in Fig. 6.5 for node a and Fig. 6.3 for node b .

Figure 6.2 shows the Time-Spectral solution of streamwise momentum, ρu , located at node a located above the mid-chord position. Solutions using the bounded interval approach using barycentric rational interpolants are shown in Fig 6.2a and compared against the use of the globally expanded least-squares Fourier projection in Fig. 6.2b and the spatial smoothing approach in Fig. 6.2c. The three methods all perform similarly for the case of $N = 3$ and $N = 5$ global collocation points. The bounded interval approach demonstrates a better agreement for the case of $N = 9$ and $N = 17$. The superior convergence of the bounded interval approach is confirmed in Fig. 6.3. Three choices of d_{\max} are presented for the barycentric rational interpolant in addition to both the natural and knot-a-knot cubic splines. A constant, zero-derivative result is plotted for reference. Also included are results obtained using the Laplacian spatial smoothing approach, the global least-squares approach (LSF) and a mixed approach using local expansions of barycentric rational interpolants with $d_{\max} = 1$. The rational interpolant- and spline-based bounded-interval approaches converge more rapidly than the spatial-smoothing, least-squares Fourier or the mixed expansion for density (Figure 6.3a) and normal momentum (Figure 6.3c). The mixed approach performs as well as the natural spline and barycentric rational interpolant treatments for streamwise momentum using $N = 65$ (Figure 6.3b).

(a) Barycentric rational interpolant, $d_{\max} = 1$ 

(b) Global least-squares Fourier projection



(c) Spatial smoothing

Figure 6.2: Inviscid Plunging Airfoil. Time-Spectral versus time-accurate values of streamwise momentum, ρu , at node a in Fig. 6.1a using (a) barycentric rational interpolants with $d_{\max} = 1$, (b) a global least-squares Fourier projection and (c) spatial smoothing to treat dynamically-blanked nodes.

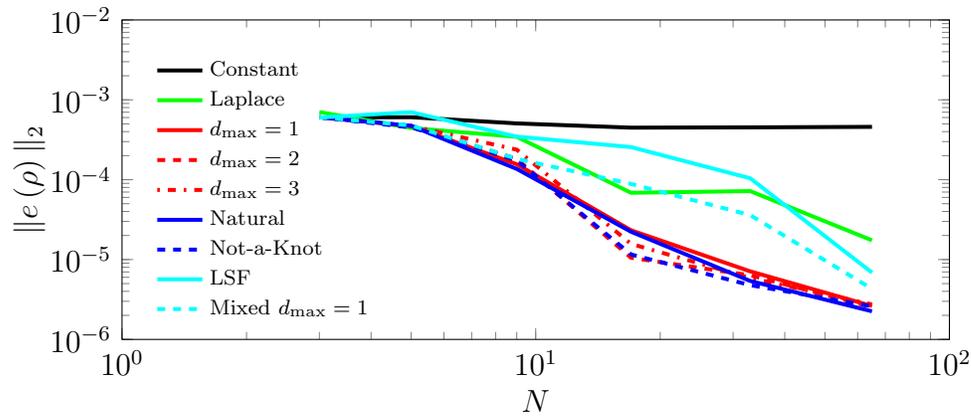
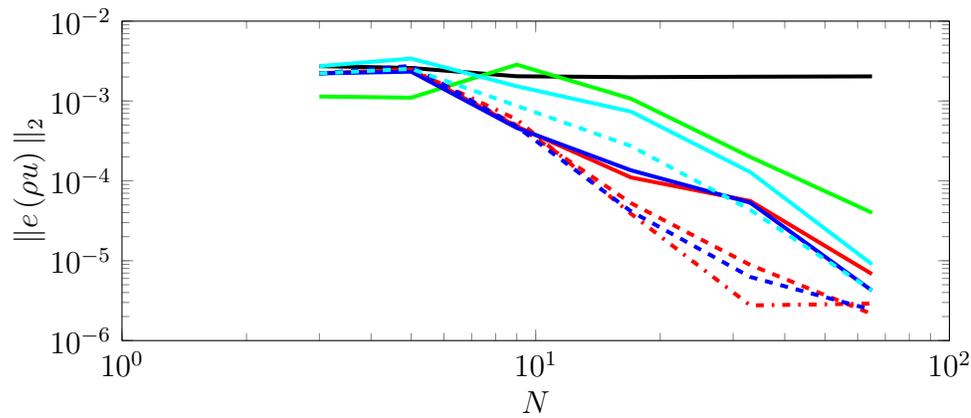
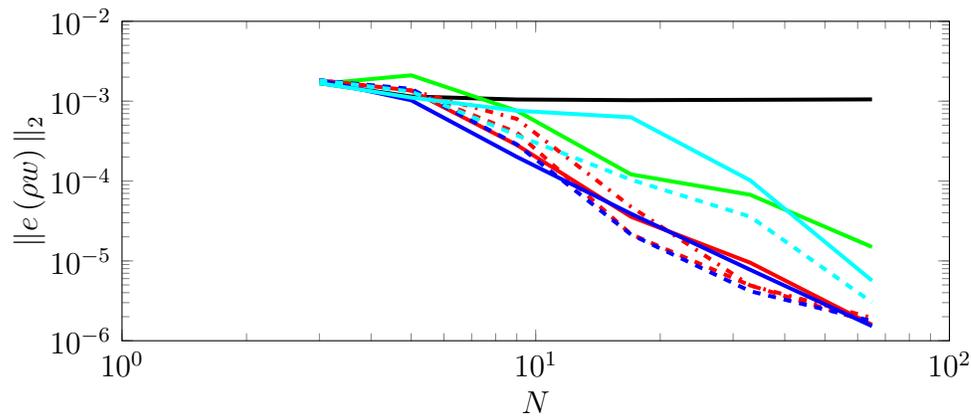
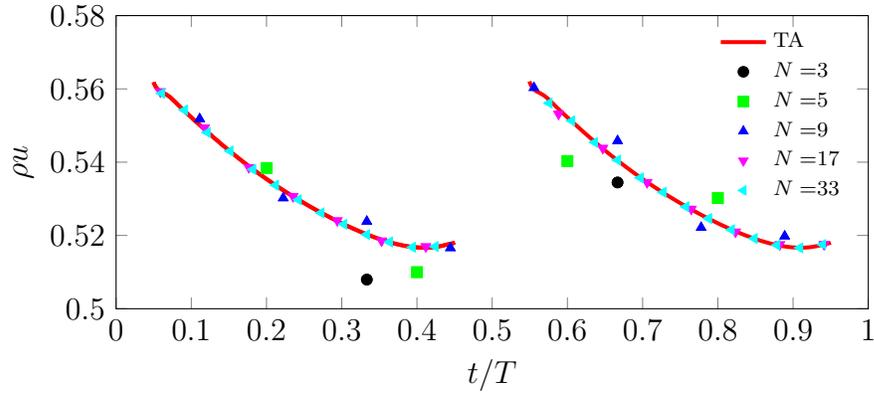
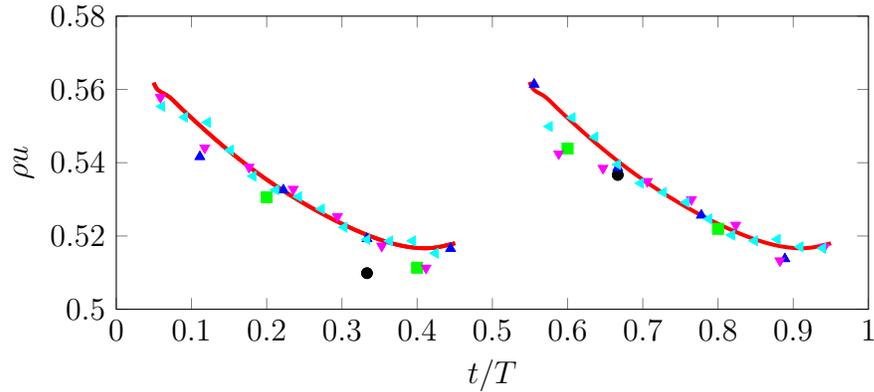
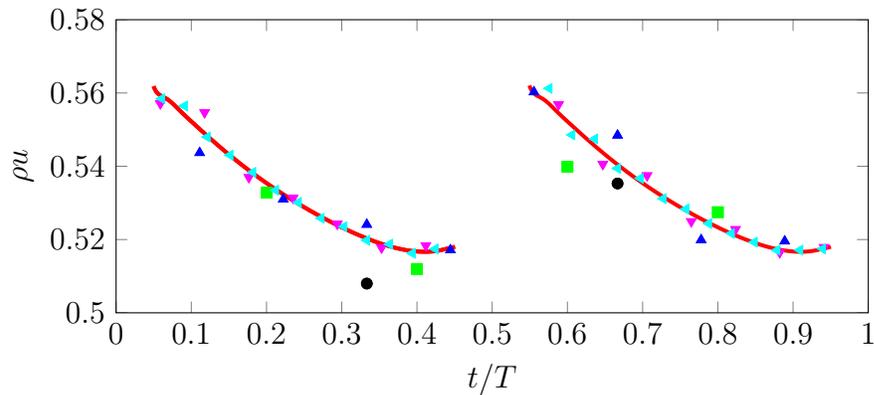
(a) Density, ρ (b) Streamwise momentum, ρu (c) Normal momentum, ρw

Figure 6.3: Inviscid Plunging Airfoil. Convergence of the RMS error of various hybrid Time-Spectral schemes versus the time-accurate solution at node a in Fig. 6.1a.

(a) Barycentric rational interpolant, $d_{\max} = 1$ 

(b) Global least-squares Fourier projection



(c) Spatial smoothing

Figure 6.4: Inviscid Plunging Airfoil. Time-Spectral versus time-accurate values of streamwise momentum, ρu , at node b in Fig. 6.1a using (a) barycentric rational interpolants with $d_{\max} = 1$, (b) a global least-squares Fourier projection and (c) spatial smoothing to treat dynamically-blanked nodes.

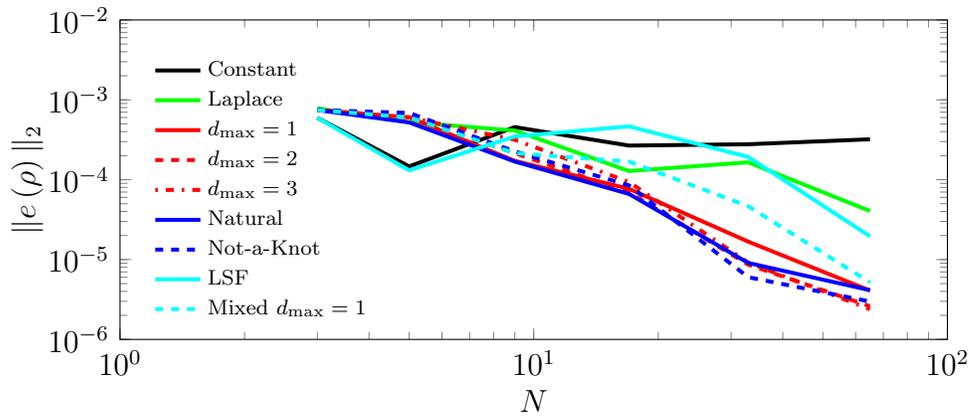
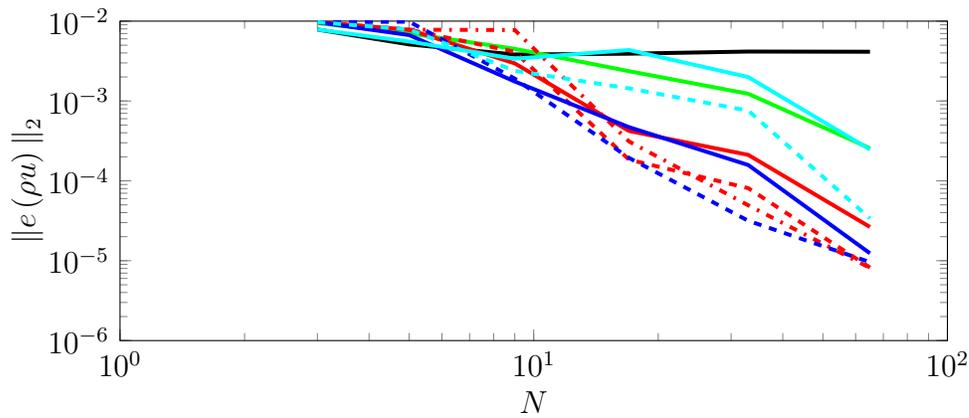
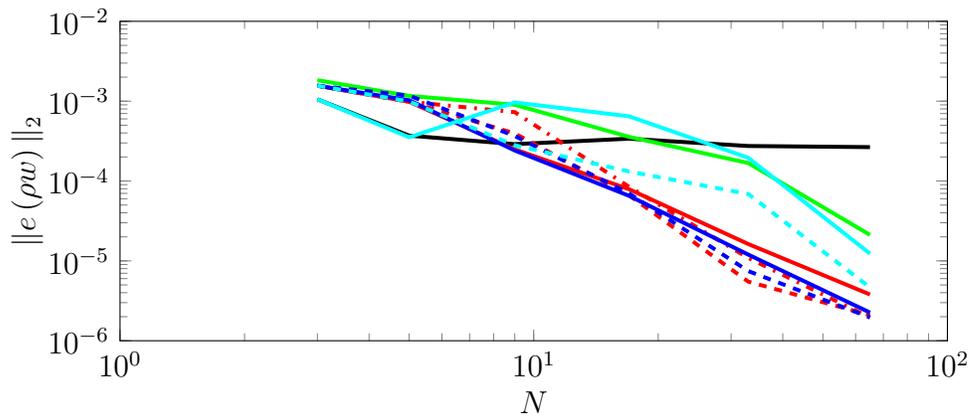
(a) Density, ρ (b) Streamwise momentum, ρu (c) Normal momentum, ρw

Figure 6.5: Inviscid Plunging Airfoil. Convergence of the RMS error of various hybrid Time-Spectral schemes versus the time-accurate solution at node b in Fig. 6.1a.

Figure 6.4 demonstrates more clearly the superior approximation properties of the bounded interval approach in contrast to the global least-squares or spatial smoothing approaches. Even using $N = 33$ global collocation points, both of the Fourier expansions exhibit high-frequency oscillations in the solution whereas the result computed using two independent partitioned segments spanned by barycentric rational interpolants demonstrates no such oscillations (in the scale of the figure). This performance is confirmed in Fig. 6.5 where the bounded interval approach again demonstrates superior convergence towards the time-accurate solution at the collocation points. While these are just two representative nodes, they do demonstrate the general behavior at all of the nodes surveyed to date. For the remainder of the dissertation, the default treatment for dynamically-blanked nodes is the barycentric rational interpolant using $d_{\max} = 1$.

The previous paragraphs examined the accuracy of the hybrid Time-Spectral approach at particular dynamically-blanked grid points. However, design is often driven by scalar performance measures. In the next paragraphs, drag coefficient polars are used to evaluate how well the Time-Spectral solutions match the time-accurate result. The Time-Spectral solutions are interpolated from their collocation points to reconstruct a continuous signal over the period. The interpolated solution can then be used to evaluate force and moment coefficients. The results from both the standard Time-Spectral and hybrid Time-Spectral calculations are then compared to the time-accurate solution.

Figure 6.6 plots the time-accurate drag coefficient, c_d , versus vertical displacement, h , over the period against the Time-Spectral results for $N \in \{3, 9, 33\}$. The time-accurate calculation is initialized with a steady state solution and simulated for ten periods of motion. The drag signal settles into a periodic steady state relatively quickly. For the case of $N = 3$, depicted in Fig. 6.6a, neither the standard Time-Spectral method (in green) nor the hybrid Time-Spectral method (in blue) accurately predicts the unsteady drag coefficient signal. However, the solution at the collocation points strongly agree with the time-accurate result, especially for the rigid-motion case using the standard Time-Spectral method. However, as the temporal resolution

is refined to $N = 9$, as depicted in Fig. 6.6b, both the rigid- and relative-motion interpolated signals demonstrate strong agreement with the time-accurate signal. However, there are slight discrepancies between the standard and hybrid Time-Spectral methods. Finally, for the case of $N = 33$, depicted in Fig. 6.6c, the two Time-Spectral methods demonstrate complete agreement with each other and the time-accurate signal.

Figure 6.7 demonstrates the monotonic convergence of both the standard (Figure 6.7a) and hybrid (Figure 6.7b) Time-Spectral calculations by plotting the continuous drag coefficient signals for $N \in \{3, 5, 9, 17, 33\}$ versus the periodic steady-state time-accurate signal. For the case of rigid motion using the standard Time-Spectral method, the signal computed using two harmonics ($N = 5$) nearly matches the time-accurate result. However, the hybrid scheme for the relative motion case required four harmonics ($N = 9$) to show similar agreement with the time-accurate signal. Thus, the relative motion case required more modes to achieve the same qualitative level of accuracy as the rigid case.

This case has demonstrated the ability of the proposed hybrid Time-Spectral scheme to successfully predict the performance for a plunging airfoil that moves relative to a background grid. However, the convergence of the relative-motion Time-Spectral scheme lags behind its rigid-motion counterpart. The sub-optimal basis functions used to describe the solution at dynamically-blanked nodes is likely partially responsible for this but so too is the fact that the stationary background grid is not moving with the body. The solution on the background grid remains periodic, however, its frequency content is driven by the relative motion between itself and the plunging near-body grid. Consider the solution at a point on the near-body grid moving with the airfoil. Over the course of the period of motion, it experiences smooth transitions between compressions and expansions. While a stationary point on the background grid in the vicinity of the airfoil experiences periodic compressions and expansions, the transitions become more rapid as the plunging amplitude increases; the airfoil moves past the stationary point with a larger velocity and therefore the transition occurs over a shorter duration. This increases the frequency of

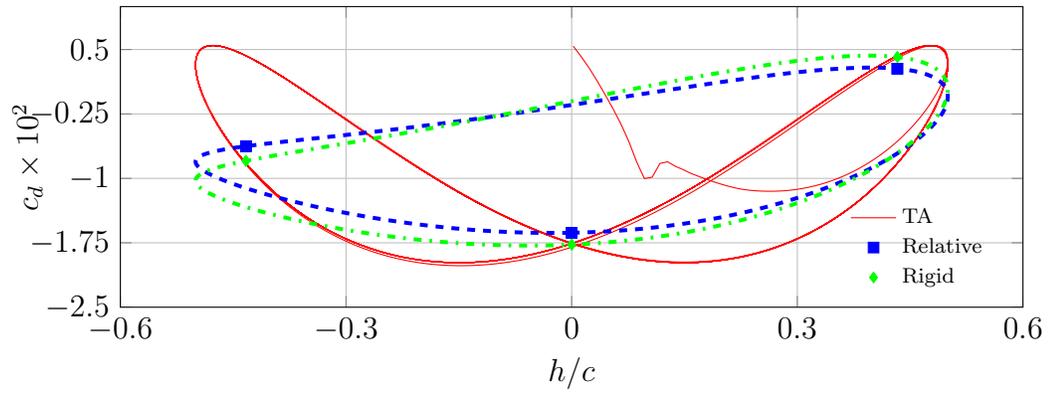
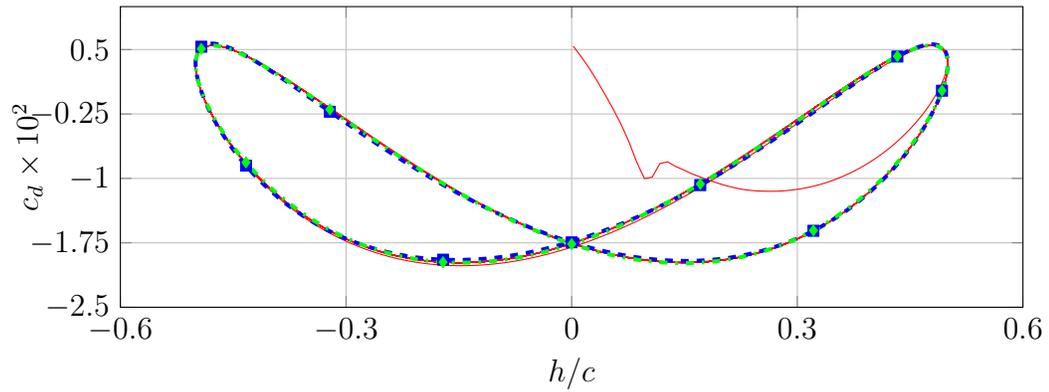
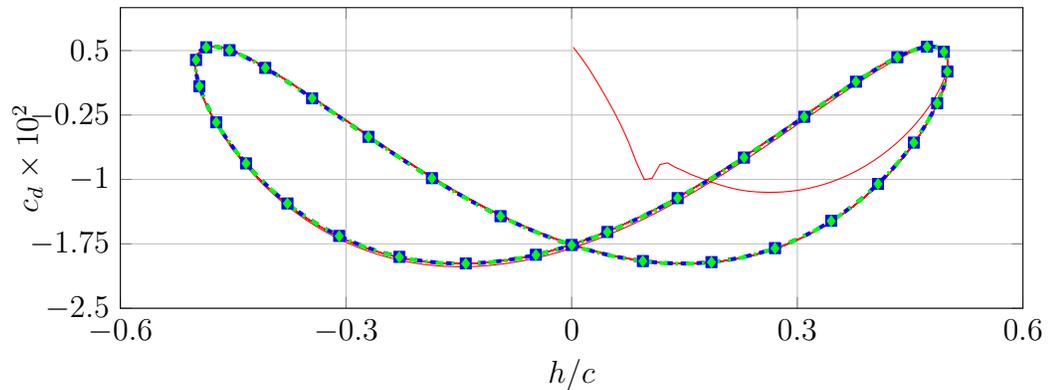
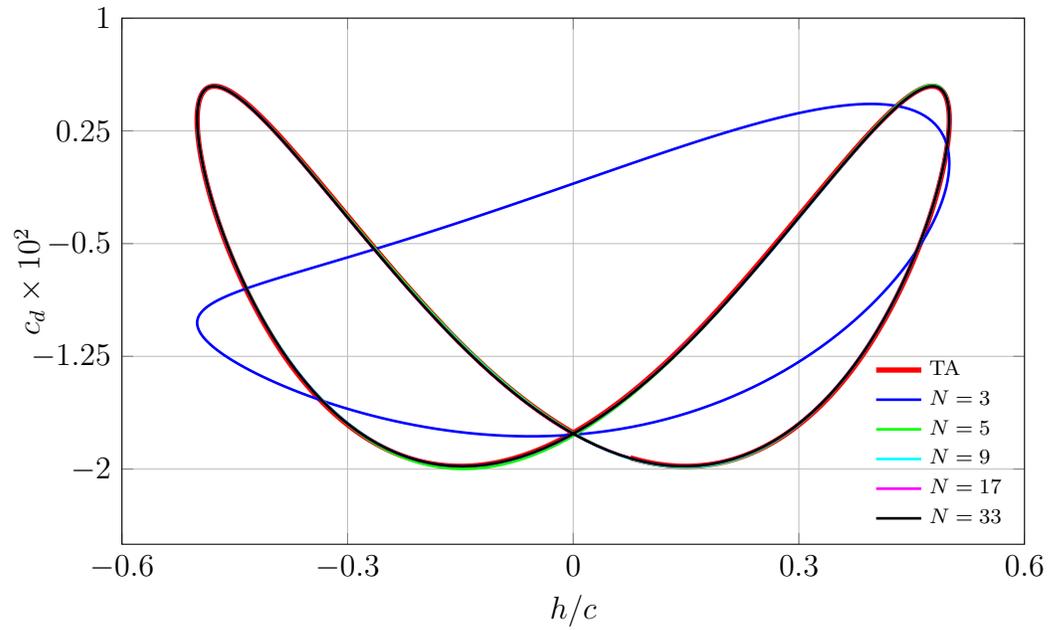
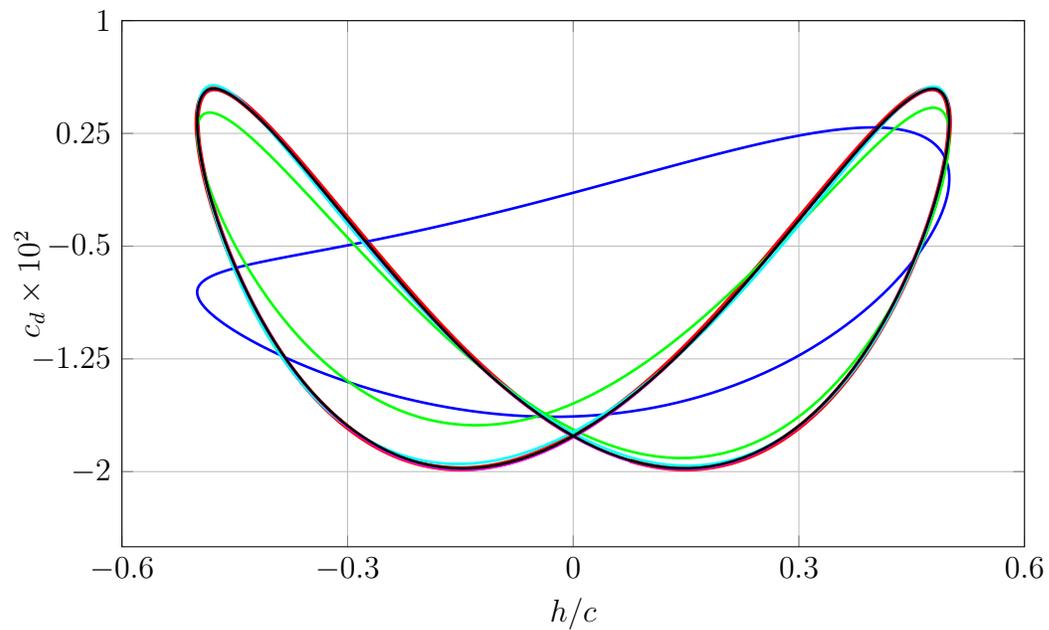
(a) $N = 3$ (b) $N = 9$ (c) $N = 33$

Figure 6.6: Inviscid Plunging Airfoil. Time-Spectral versus time-accurate drag coefficients for $N \in \{3, 9, 33\}$. Time-accurate solution is plotted in red from steady-state startup. Blue squares and green diamonds locate the drag coefficient values at the Time-Spectral collocation points for relative- and rigid-body motion, respectively. Drag coefficient signals evaluated from reconstructed Time-Spectral solutions are plotted with hashed lines.



(a) Rigid Motion



(b) Relative Motion

Figure 6.7: Inviscid Plunging Airfoil. Convergence of (a) rigid- and (b) relative-motion Time-Spectral drag coefficient signals to the periodic steady-state time-accurate calculation for $N \in \{3, 5, 9, 17, 33\}$.

the response that needs to be approximated by the discrete Fourier series. Therefore, for a given plunging amplitude and reduced frequency, more modes may need to be retained in order to resolve the relative-motion case because its solution contains higher-harmonics on the stationary grid.

Figures 6.8 and 6.9 provide the time- and frequency-response of the subsonic inviscid plunging airfoil at a node located approximately $0.35c$ downstream of the trailing edge on both the near- and off-body grids in the neutral position. For the case of rigid-body motion, the off-body grid translates with the near-body grid and the nodes remain approximately coincident for all time. For the case of relative motion, the node on the off-body grid remains in place and is therefore only approximately coincident with the node on the near-body grid at two time instances over the period of oscillation. The higher-frequency content in the solution at a node on the stationary background grid suggests that more frequencies are required to resolve the solution using relative motion than with rigid motion. This assertion is supported by Fig. 6.10 that plots the reconstruction of the off-body grid solution from the frequency components defined by a discrete Fourier transform of the time-accurate solution. Figure 6.10a demonstrates that roughly four modes are sufficient to adequately resolve the solution at the off-body grid point for the case of rigid-body motion. However, the reconstructions plotted in Fig. 6.10b suggest that thirty or more modes are required to sufficiently resolve the flow at the stationary off-body grid point. However, many fewer modes are required to match the time-accurate result with respect to integrated loads such as drag (Figure 6.7), supporting the Time-Spectral method as a potential reduced-order model (ROM). For such simple cases, rigid motion would naturally be a more efficient choice as fewer modes are required to resolve the flow. However, more complex configurations require relative motion for non-deforming grids. The transmission of low-frequency signals in one frame to high-frequency signals in another frame proves fundamental to the case of rotorcraft, which is explored in an upcoming section.

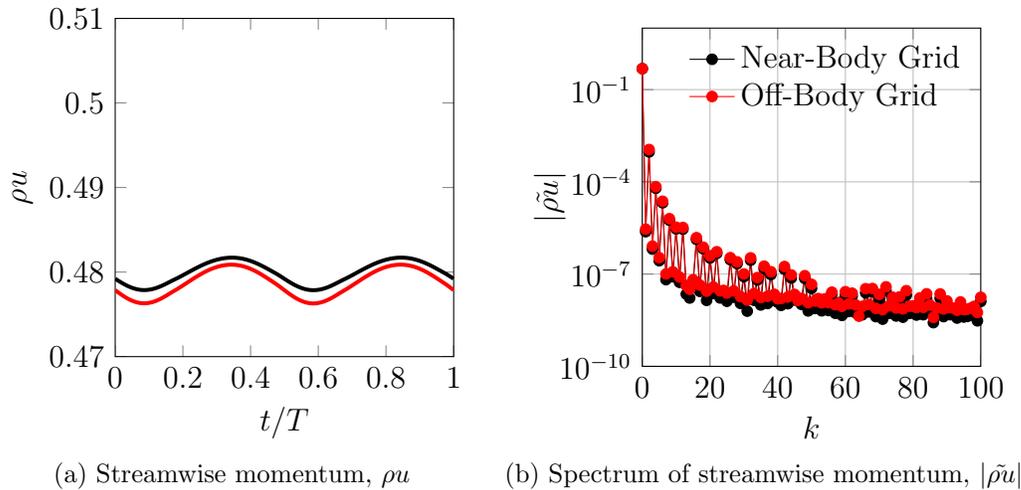


Figure 6.8: Inviscid Plunging Airfoil. Rigid motion case. Time and frequency response of a plunging airfoil with an off-body grid translating with the airfoil. The frequency components at the near- and off-body nodes are nearly identical.

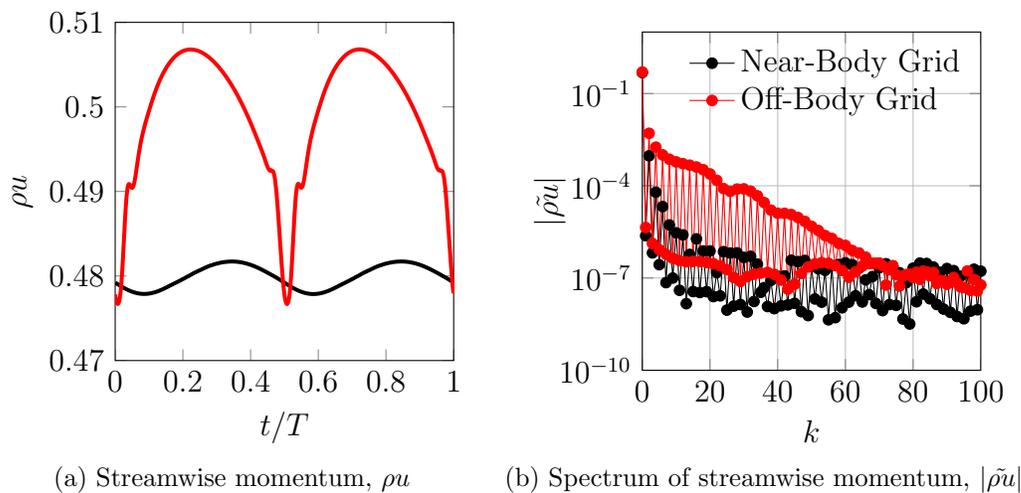
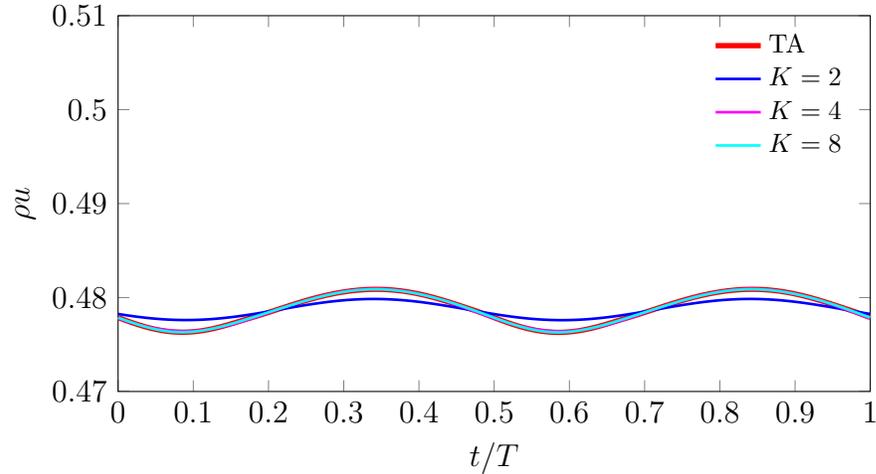
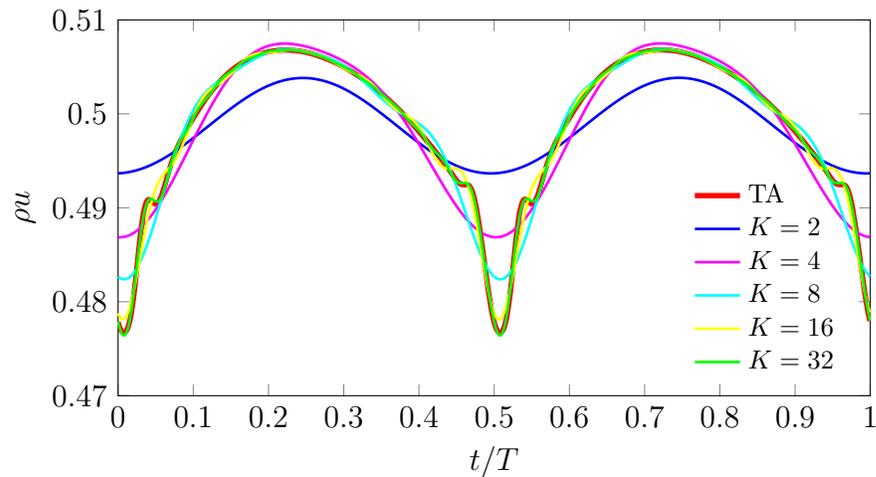


Figure 6.9: Inviscid Plunging Airfoil. Relative motion case. Time and frequency response of a plunging airfoil with a stationary background grid. The frequency components at the near- and off-body nodes are strikingly different with many more modes required to resolve the solution at the node on the stationary background grid. The node on the background grid is the same one described in Fig. 6.8a, but the background grid was translating with the airfoil in that case.



(a) Translating node on background grid for rigid-body motion.



(b) Stationary node on background grid for relative-body motion.

Figure 6.10: Inviscid Plunging Airfoil. Reconstruction of streamwise momentum at a node on the off-body grid (a) translating with the airfoil and (b) stationary relative to the translating airfoil. The solution at the stationary node requires approximately $K = 32$ modes to reasonably reconstruct the continuous signal whereas the solution at the translating node requires approximately $K = 4$ modes. This suggests an inherent inefficiency of employing relative motion. However, rigid-body motion is not universally applicable and therefore relative-body motion must be used for certain configurations.

6.1.2 Laminar Plunging NACA 0012 Airfoil

In this section, a pair of low speed, low Reynolds number laminar plunging NACA 0012 airfoil cases are investigated as a precursor to the three-dimensional rotorcraft cases. These flows exhibit some similar physical phenomena to the rotorcraft flows, most notably the convection of strong vortices through the wake. Jones et al. [100] published results from low-speed experiments that serve to highlight the complex nature of the vortical structures associated with these flows as a function of their Strouhal number, $St = kh$, which is a product of their reduced frequency, $k = \omega c/V_\infty$, and nondimensional plunging amplitude, $h = a/c$. The instantaneous vertical displacement, $y(t) = h \sin(kt)$, is defined as a single sinusoid of the reduced frequency.

At low Strouhal numbers, the plunging motion results in vortices with a momentum deficit, resulting in a time-averaged drag on the airfoil. As the frequency is increased, the vortical structure changes, eventually resulting in a thrust-producing momentum surplus. Increasing the Strouhal number even further results in both lift and thrust. This behavior has been coined the *Knoller-Betz* effect after the experimentalists who first independently documented the phenomenon [100]. Rotors in forward flight undergo a host of oscillations, either prescribed or passively via fluid-structure interaction, suggesting that analysis of the two-dimensional problem should uncover computational strategies for the full three-dimensional configurations.

Two primary cases are investigated: a drag-producing case with $St = 0.288$ and a higher-frequency thrust-producing case with $St = 0.6$. Freestream Mach number, $M_\infty = 0.2$, and Reynolds number, $Re = 1850$, are used for both cases. Each case is computed using both rigid- and relative-body motion to compare the performance between the standard and hybrid Time-Spectral methods. The grid system presented in Fig. 6.11 consists of a 589×33 C-mesh near-body grid embedded within a hierarchy of isotropic Cartesian off-body grids. The large 909×309 finest level-one (L1) off-body grid spans $8c$ downstream with $0.01c$ grid spacing to accurately resolve a significant portion of the highly-active wake. The hole-cut standoff distance is $0.075c$ and the Cartesian grids span approximately $200c$ to the farfield boundaries in both the x - and y -directions.

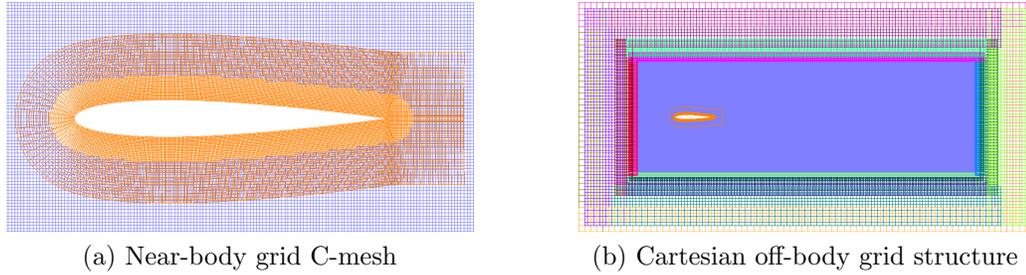


Figure 6.11: Laminar Plunging Airfoil. Grid system.

Drag-Producing Case

The lower-frequency test case of $St = 0.288$ with reduced frequency $k = 3.6$ and plunging amplitude $h = 0.08$ results in a vortical wake structure that creates a momentum deficit, producing a time-averaged drag on the airfoil. The case was simulated using time-accurate and Time-Spectral approaches for both rigid and relative motion and qualitatively compared to the experimental results offered in [100]. Figure 6.12 provides snapshots of the experimental flowfield and vorticity magnitude of the simulations. The time-accurate calculation successfully captures the primary flow features depicted in the experimental snapshot and the Time-Spectral simulations, for both rigid and relative motion, converge to a similar solution, provided a sufficient temporal resolution of around 16 modes ($N = 33$) is used. There exist significant variations for the under-resolved Time-Spectral solutions between the rigid- and relative-motion cases, but the underlying wake structure is apparent with only a few modes in either case. Figure 6.13 confirms the rapid convergence of the Time-Spectral computed drag coefficient, c_d , to the time-accurate result. Employing just four modes (See Figure 6.13b) for either Time-Spectral configuration nearly matches the time-accurate signal. In fact, the relative motion converges slightly faster as evidenced in Fig. 6.14 where the cyan line corresponding to the rigid-motion case with $N = 9$ time-samples lies off the periodic steady-state time-accurate signal for a portion of the period, but the relative-motion signal is identical to the time-accurate signal to the resolution of the figure. Thus, the scalar performance measure of drag converges more rapidly than the vortical flow features.

Thrust-Producing Case

As the Strouhal number is increased to $St = 0.6$, the vortical wake structure evolves from producing a momentum deficit, as demonstrated for $St = 0.288$, to a momentum surplus, where the time-averaged flow exhibits a thrust-inducing jet structure. Selecting a reduced frequency, $k = 6.0$ and plunging amplitude $h = 0.1$, results in the particular wake structure depicted in Fig. 6.15. Experimental results taken from [100] and depicted in Fig. 6.15a use this Strouhal number, $St = 0.6$, except that the reduced frequency is halved and the plunging amplitude doubled ($k = 3.0$ and $h = 0.2$). Simulations using these parameters resulted in significantly different vortical wake patterns. Therefore, the case using the higher reduced frequency is employed, which is consistent with the observations and calculations in [101]. Figure 6.15 demonstrates strong agreement between the time-accurate and experimental data and again demonstrates the convergence of the Time-Spectral solutions for both rigid and relative motion with increased temporal resolution. As with the drag-producing cases, the Time-Spectral calculations require sixteen modes ($N = 33$) to qualitatively match the time-accurate solution. The bulk features of the wake are resolved using half as many modes, but the solutions do not match in the vicinity of the trailing edge. Figure 6.16 confirms the convergence of the Time-Spectral computed drag coefficient, c_d , to the time-accurate result for the thrust-producing case. Seventeen time-samples are sufficient to reproduce the periodic steady-state time-accurate result for both the rigid- and relative-motion Time-Spectral calculations (Figure 6.17).

Unlike the drag-producing case using $k = 0.288$, spectral vanishing viscosity is required to enable adequate convergence of the Time-Spectral simulations for the thrust-producing configuration for both the rigid- and relative-body motion calculations. Initially, SVV is only applied to the statically-blanked (Fourier-based) nodes in the case of relative motion, leaving the dynamically-blanked nodes free of any added dissipation. For the relative motion case using eight harmonics ($N = 17$), additional dissipation is required at the undamped, dynamically-blanked nodes as described in §5.3.3.

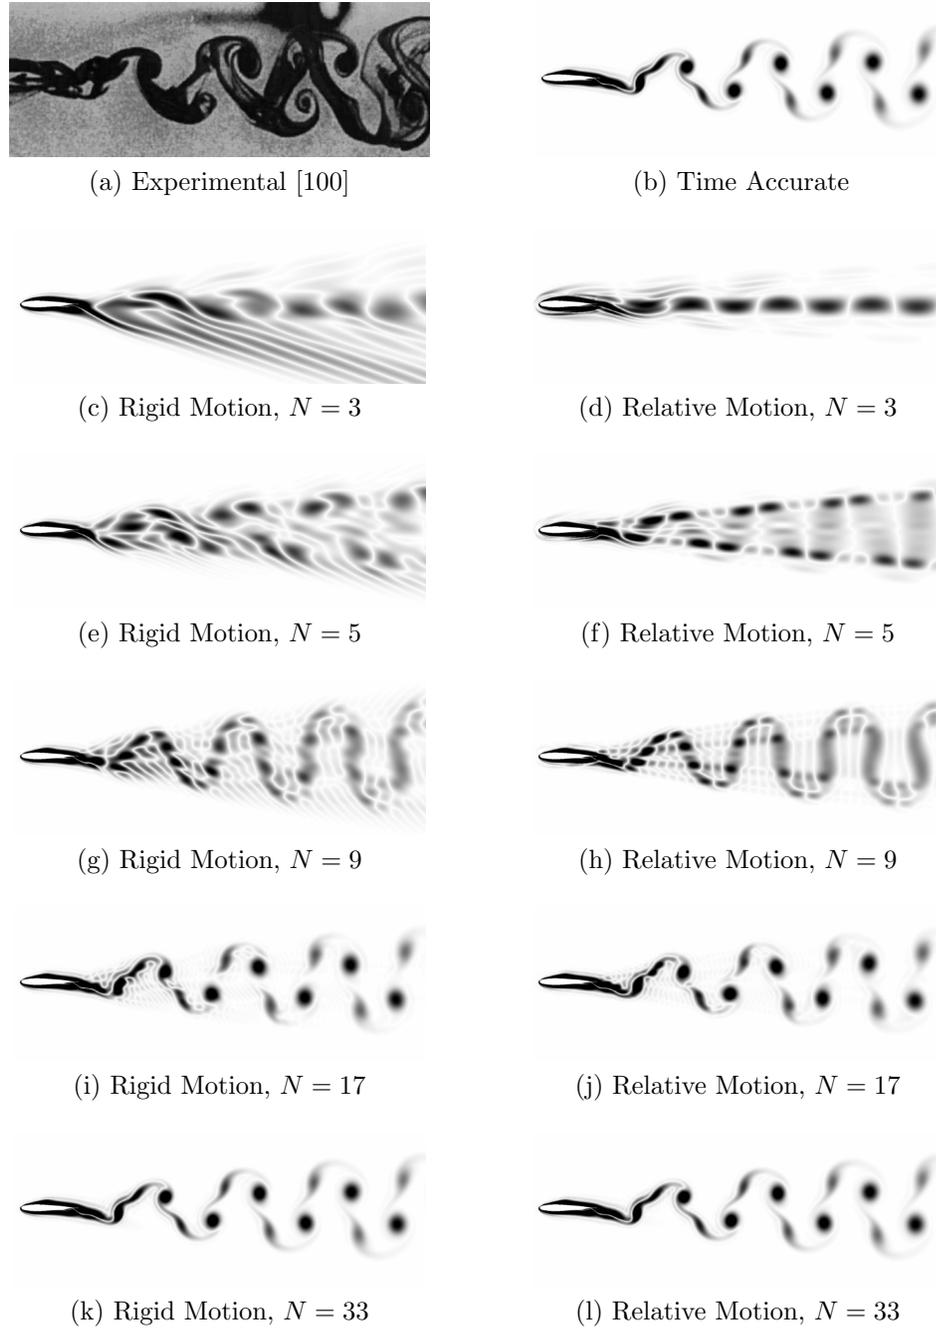


Figure 6.12: Laminar Plunging Airfoil. Drag-producing case. Visualization of vorticity magnitude. Comparison of experimental with computed time-accurate and Time-Spectral results using $N \in \{3, 5, 9, 17, 33\}$ for both rigid- and relative-body motion. Note that while certain flow features are not resolved by the lower-mode Time-Spectral cases, Fig. 6.14 suggests that using $N \geq 9$ matches the time-accurate drag-coefficient for nearly the entire period of oscillation. Vorticity magnitude varies from 0 in white to greater than or equal to 1 in black.

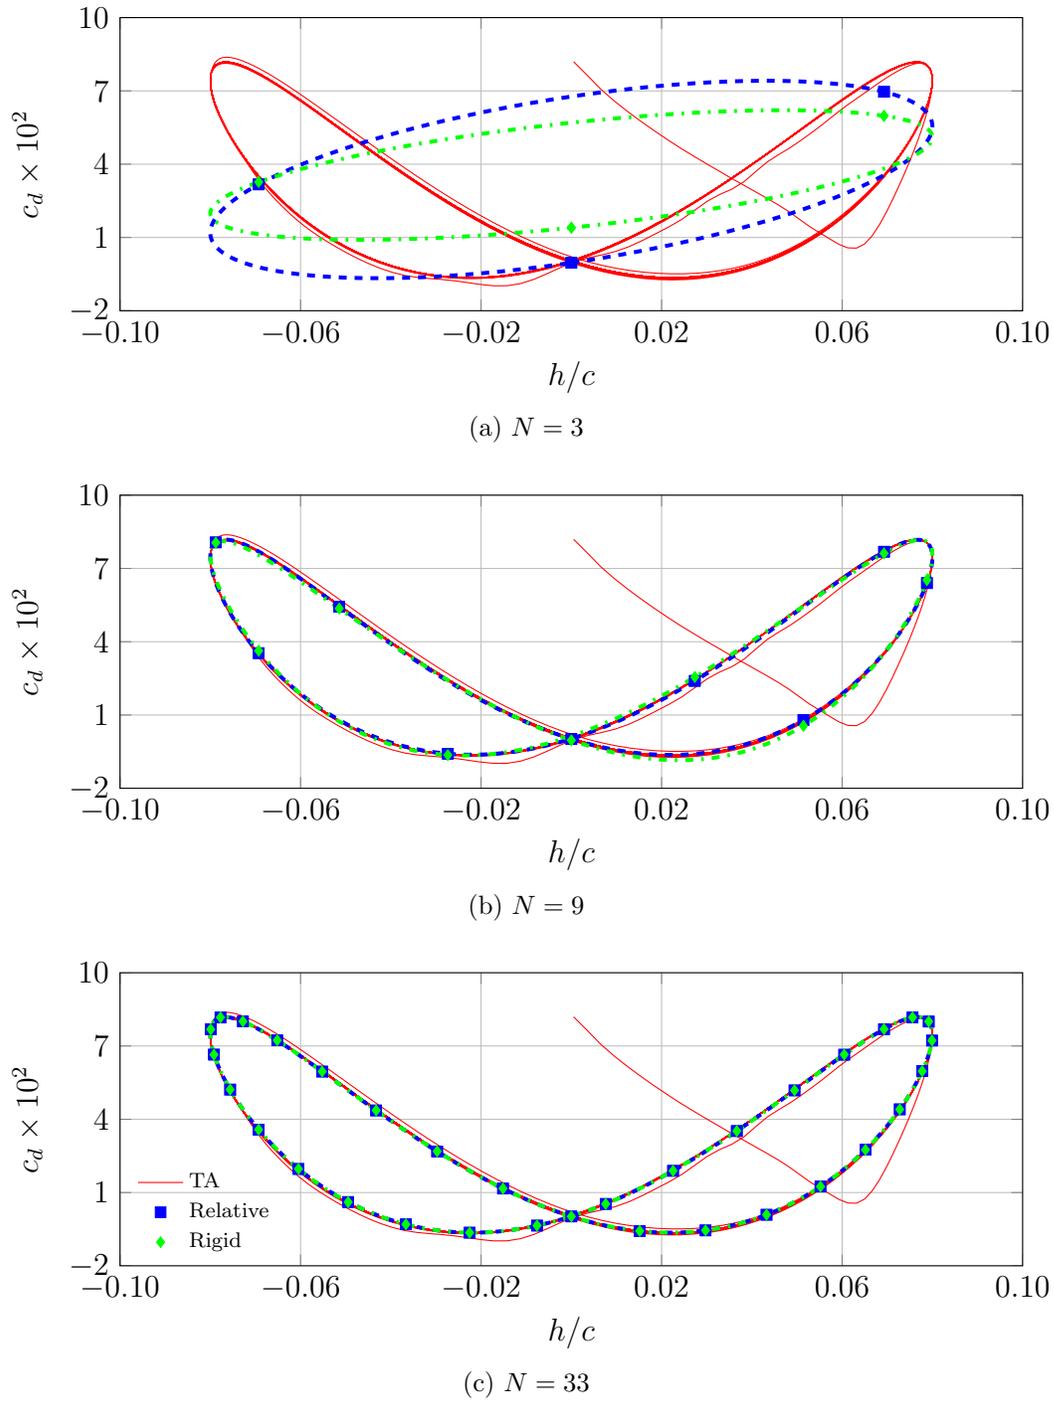
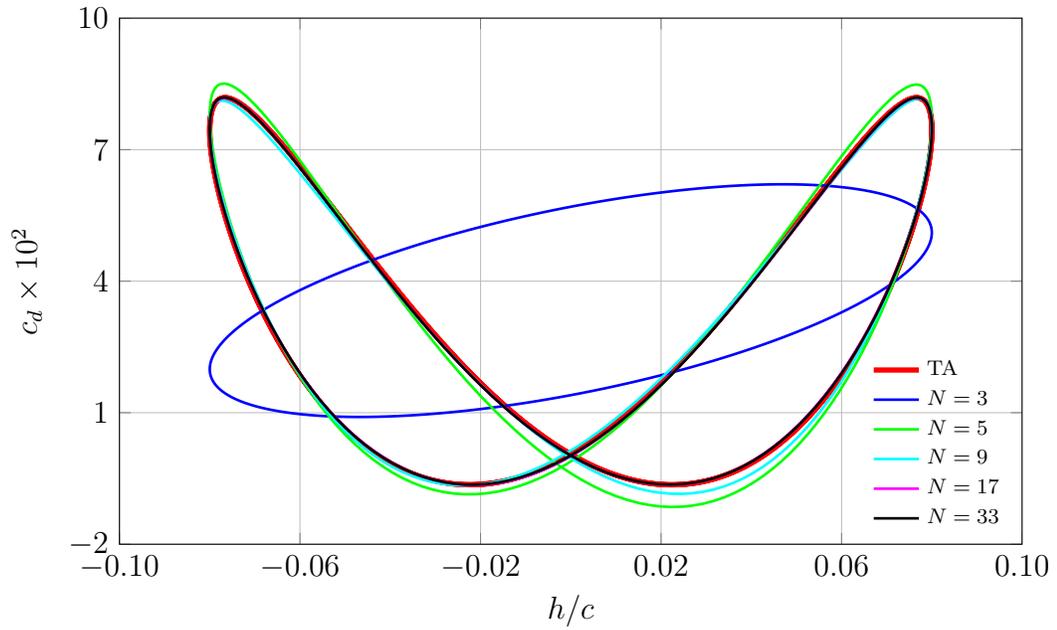
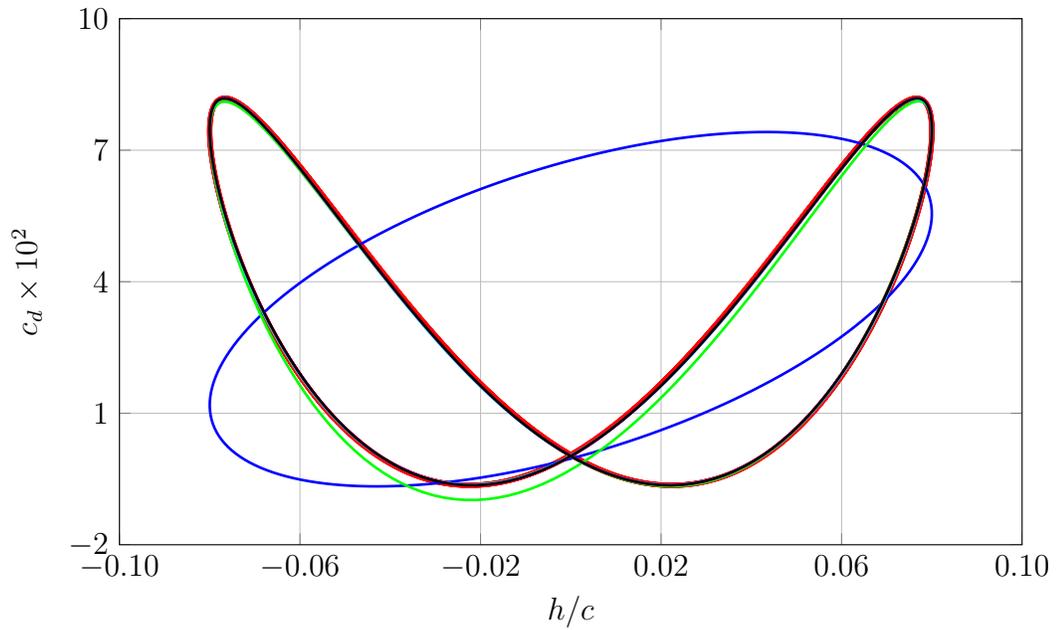


Figure 6.13: Laminar Plunging Airfoil. Drag-producing case. Time-Spectral versus time-accurate drag coefficients for $N \in \{3, 9, 33\}$. Time-accurate solution is plotted in red from steady-state startup. Blue squares and green diamonds locate the drag coefficient values at the Time-Spectral collocation points for relative- and rigid-body motion, respectively. Drag coefficient signals evaluated from reconstructed Time-Spectral solutions are plotted with hashed lines.



(a) Rigid Motion



(b) Relative Motion

Figure 6.14: Laminar Plunging Airfoil. Drag-producing case. Convergence of (a) rigid- and (b) relative-motion Time-Spectral drag coefficient signals to the periodic steady-state time-accurate calculation for $N \in \{3, 5, 9, 17, 33\}$.

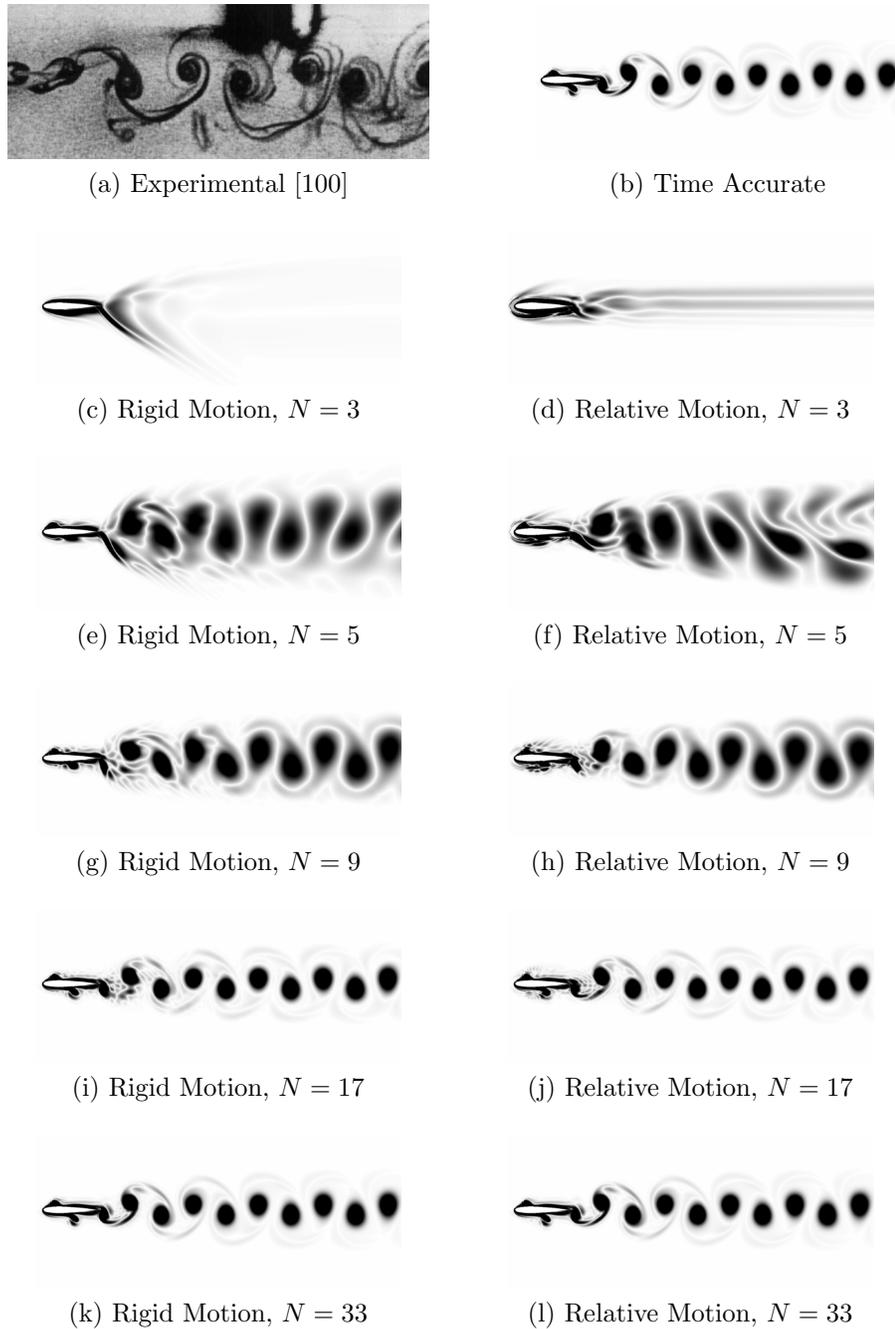


Figure 6.15: Laminar Plunging Airfoil. Thrust-producing case. Visualization of vorticity magnitude. Comparison of experimental with computed time-accurate and Time-Spectral results using $N \in \{3, 5, 9, 17, 33\}$ for both rigid- and relative-body motion. Note that while certain flow features are not resolved by the lower-mode Time-Spectral cases, Fig. 6.17 suggests that using $N \geq 9$ matches the time-accurate drag-coefficient for nearly the entire period of oscillation. Vorticity magnitude varies from 0 in white to greater than or equal to 1 in black.

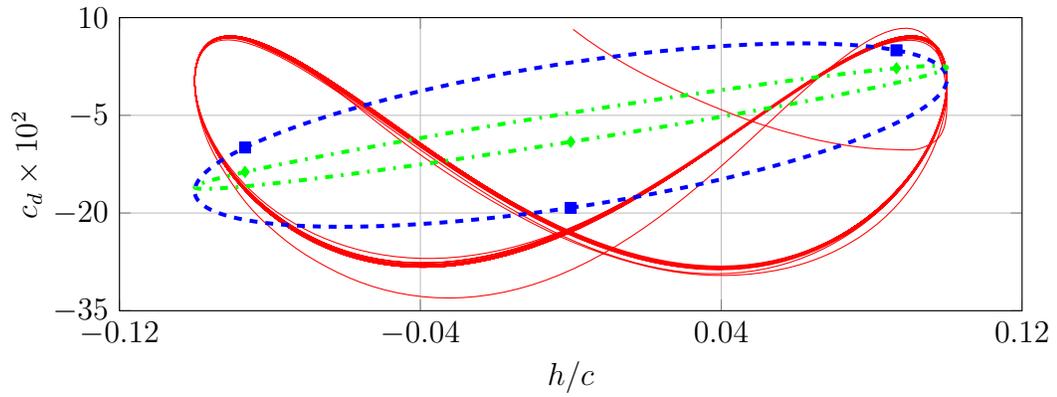
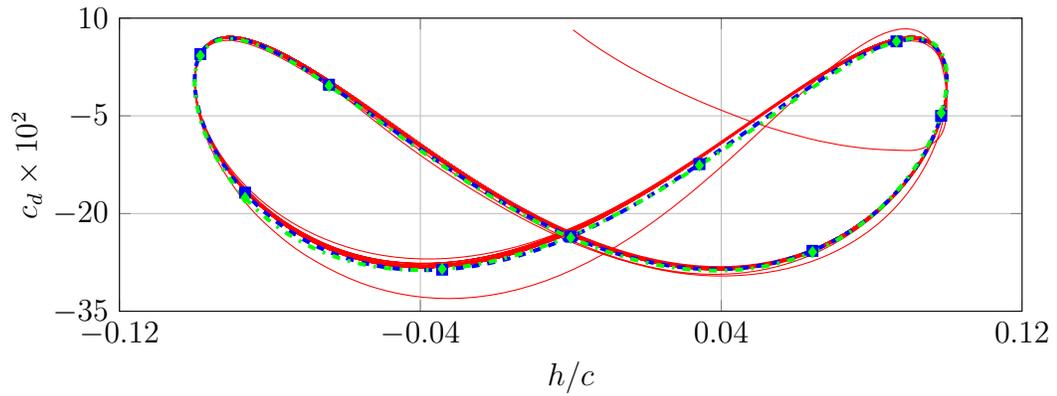
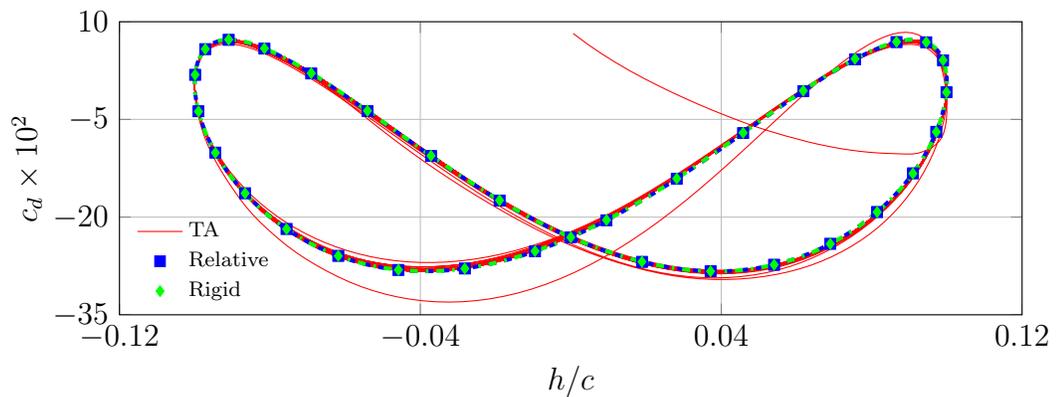
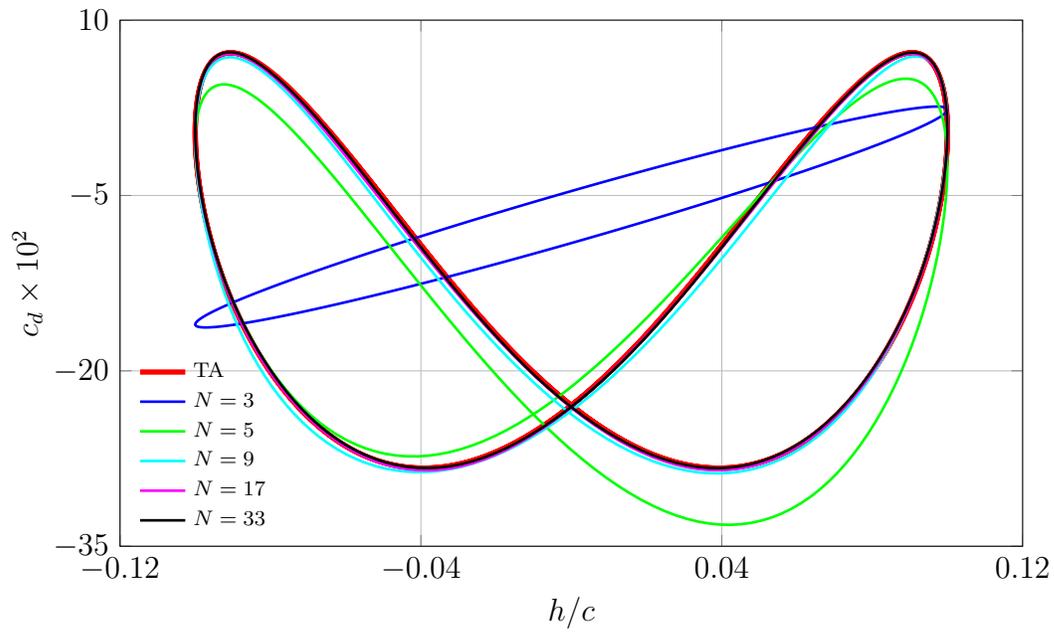
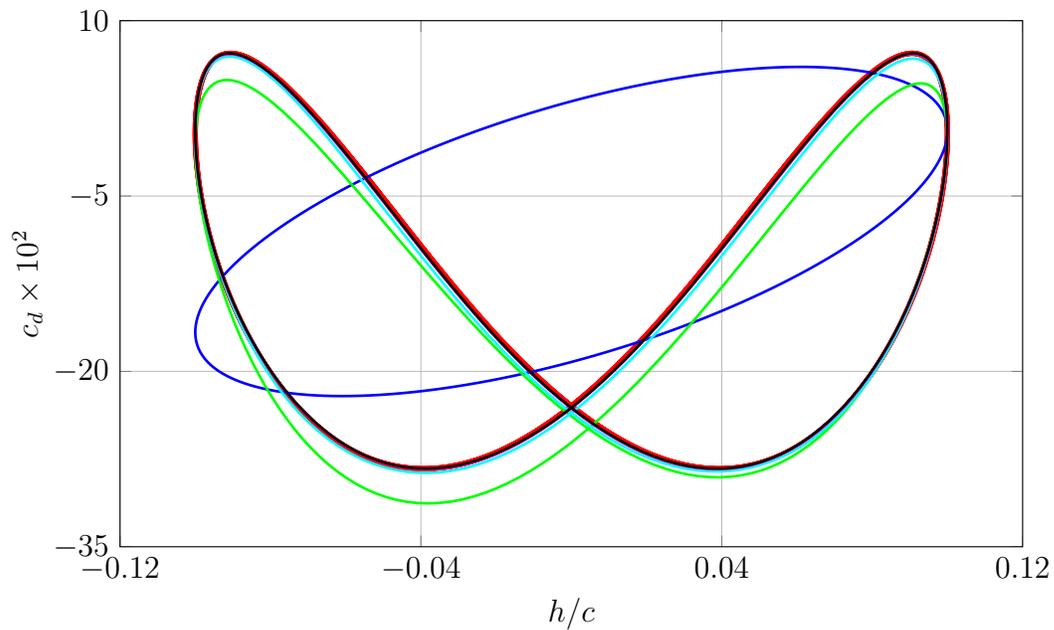
(a) $N = 3$ (b) $N = 9$ (c) $N = 33$

Figure 6.16: Laminar Plunging Airfoil. Thrust-producing case. Time-Spectral versus time-accurate drag coefficients for $N \in \{3, 9, 33\}$. Time-accurate solution is plotted in red from steady-state startup. Blue squares and green diamonds locate the drag coefficient values at the Time-Spectral collocation points for relative- and rigid-body motion, respectively. Drag coefficient signals evaluated from reconstructed Time-Spectral solutions are plotted with hashed lines.



(a) Rigid Motion



(b) Relative Motion

Figure 6.17: Laminar Plunging Airfoil. Thrust-producing case. Convergence of (a) rigid- and (b) relative-motion Time-Spectral drag coefficients to the periodic steady-state time-accurate calculation for $N \in \{3, 5, 9, 17, 33\}$.

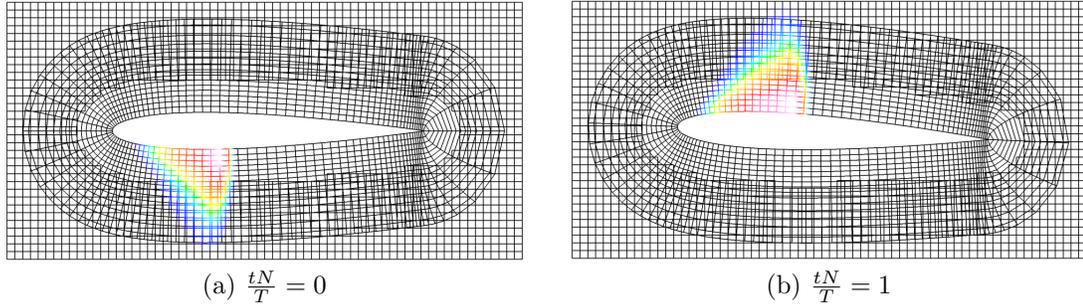


Figure 6.18: Inviscid Transonic Pitching Airfoil. The supersonic region is plotted on the grid system in the vicinity of the airfoil for two time samples of the $N = 5$ case demonstrating the shock spanning the overset interface. Subsonic regions shown in black. Every second grid point has been removed for clarity.

6.1.3 Transonic Pitching NACA 0012 Airfoil

The Time-Spectral formulation is next compared to time-accurate simulations and experimental data corresponding to the AGARD 702 3E3 oscillatory pitch test case CT5 [102]. Both inviscid and viscous turbulent simulations are performed on this transonic, $M_\infty = 0.755$, case. As with plunging, a pitching airfoil allows for simulations using either both rigid- and relative-body motion. The airfoil pitches about its quarter chord with incidence, $\alpha(t) = \alpha_0 + \alpha \sin(kt)$. This case employs mean incidence, $\alpha_0 = 0.016^\circ$ and pitching amplitude $\alpha = 2.51^\circ$, with reduced frequency $k = 0.1627$ radians per non-dimensional time unit and freestream Mach number $M_\infty = 0.755$.

A hierarchy of Time-Spectral simulations have been computed using an increasing number of modes ($N \in \{3, 5, 9, 17, 33\}$) to investigate the performance of the scheme with increased temporal resolution. The near-body grid and its overlap region are chosen such that the transient shock spans the grid interface. Figure 6.18 shows the supersonic region and the overset-grid system at two of the temporal collocation points for the $N = 5$ inviscid case.

Inviscid Case

A 241×30 O-mesh near-body grid is embedded within a 341×261 Cartesian background grid that stretches approximately $100c$ to the farfield boundaries in both

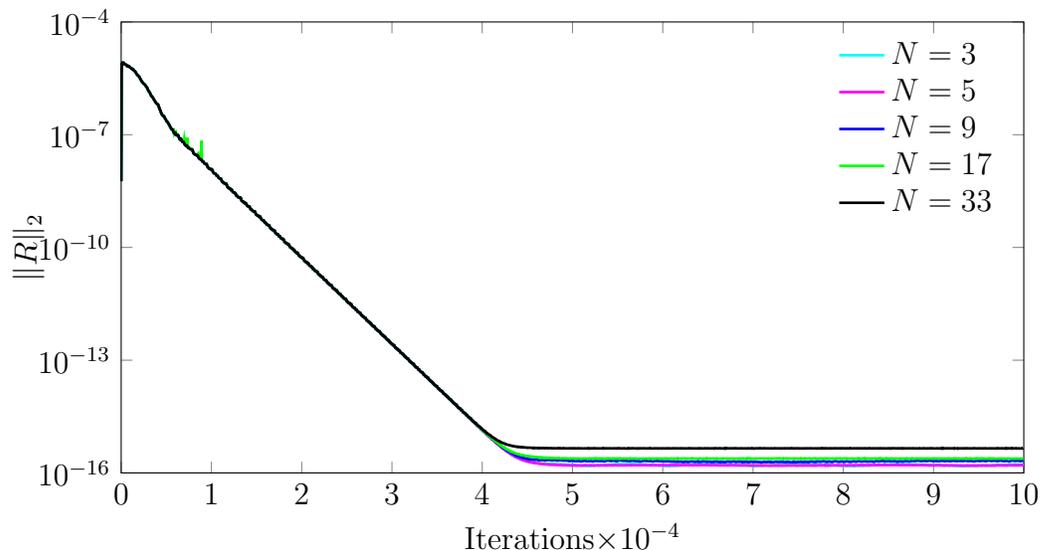


Figure 6.19: Inviscid Transonic Pitching Airfoil. Iterative convergence of the space-time residual for relative-body motion Time-Spectral calculations for $N \in \{3, 5, 9, 17, 33\}$.

directions (Figure 6.18). In the case of rigid-body motion, both the near-body and background grids rotate together. This results in a constant hole cut permitting the use of the conventional Time-Spectral approach using the Fourier-based differentiation operator. In the relative-motion case, the near-body grid rotates independently of the background grid introducing dynamic hole cutting as a result of the relative motion. The unsteady residual is converged to machine precision at each physical time step in the time-accurate case. The complete space-time residual for the Time-Spectral cases are converged to machine precision (Figure 6.19). Notably, the convergence rate is independent of the number of Time-Spectral modes in this case.

Pitching moment and drag coefficient polars are plotted in Figs. 6.20 and 6.21, respectively. The rigid- and relative-motion Time-Spectral calculations and the time-accurate calculation all converge to nearly the identical solution. Reasonable agreement with the experimental data is achieved but most importantly, monotonic convergence of the hybrid Time-Spectral solutions is achieved as the number of sample points is increased from 3 to 33 allowing the solution to resolve increasingly higher-frequency content inherent in the unsteady solution and its pitching moment and

drag coefficient signals. Note that 33 modes are required to resolve the kink in the pitching moment and drag signals at $\alpha \pm 2^\circ$.

The frequency components, amplitude and phase, of the relative-motion time-accurate pitching moment coefficient are plotted in Fig. 6.22 against those computed by the hybrid Time-Spectral scheme. Figure 6.22a reveals that the $N = 3$ case fails to resolve even the mean amplitude and the $N = 5$ case poorly resolves either the amplitude or phase corresponding to the second harmonic. Increasing the number of solution modes permits resolution of higher-frequency content, and improves resolution of the lower frequencies. Importantly, no spurious high-frequency content corrupts (or destabilizes) the solution as a result of the hybrid Time-Spectral scheme.

Turbulent RANS Case

Practical applications of the Time-Spectral method including rotorcraft and turbomachinery involve turbulent flows, and therefore the previous oscillatory airfoil test case is extended to use the Spalart-Allmaras one-equation turbulence model on a viscous grid system. The Time-Spectral implementation employs a semi-implicit approach, as described in §5.3.4, to avoid a massive overhaul of the codebase.

To illustrate a potential difficulty of a Time-Spectral application to the positivity-constrained, one-equation Spalart-Allmaras turbulence model, consider Fig. 6.23 showing contours of the (undamped) turbulent eddy viscosity, $\tilde{\nu}$, in the wake region at five temporal collocation points over the period of oscillation. This demonstrates the potentially large changes in $\tilde{\nu}$ as a function of time in the wake region. Figure 6.24 plots the time history of turbulent eddy viscosity and its temporal derivative at a point downstream of the airfoil through which the turbulent wake passes during a portion of the oscillation. Note in Fig. 6.24a, the turbulent eddy viscosity transitions from a constant value of order zero to a value of order 1000 over a small fraction of the period. While the function varies smoothly in the context of the small physical time steps associated with a time-accurate calculation, it resembles a discontinuous step function in the frame of the much coarser resolution of Time-Spectral collocation points. This leads to spurious oscillations in the Fourier expansion (Gibbs' phenomenon).

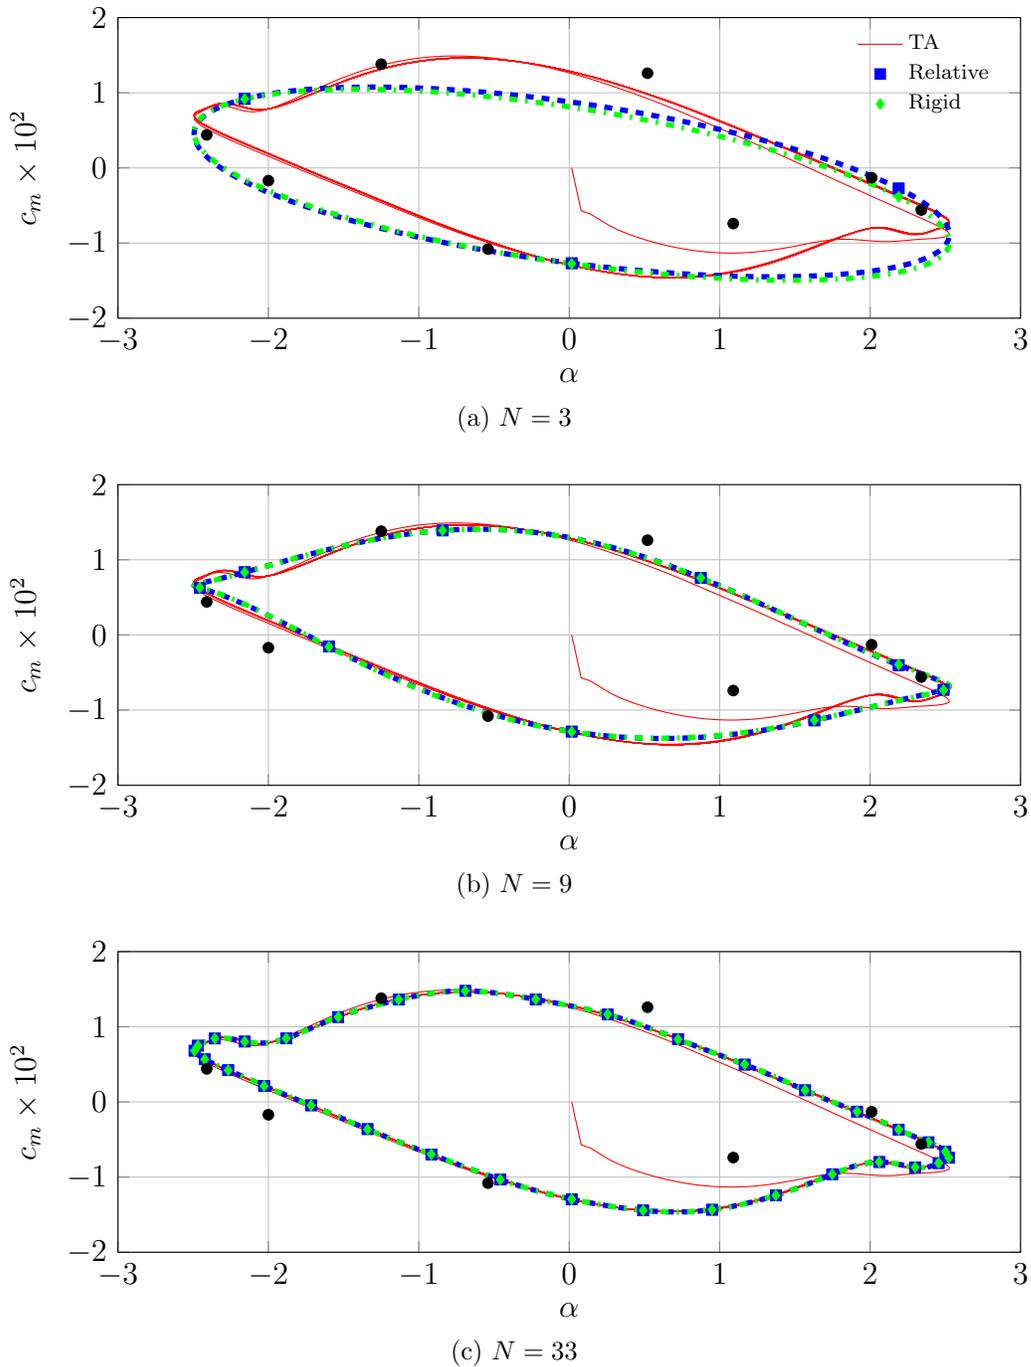


Figure 6.20: Inviscid Transonic Pitching Airfoil. Time-Spectral versus time-accurate pitching moment coefficients for $N \in \{3, 9, 33\}$. Time-accurate solution is plotted in red from steady-state startup. Blue squares and green diamonds locate the moment coefficient values at the Time-Spectral collocation points for relative- and rigid-body motion, respectively. Moment coefficient signals evaluated from reconstructed Time-Spectral solutions are plotted with hashed lines. Experimental data from the AGARD 702 Report are plotted with black dots.

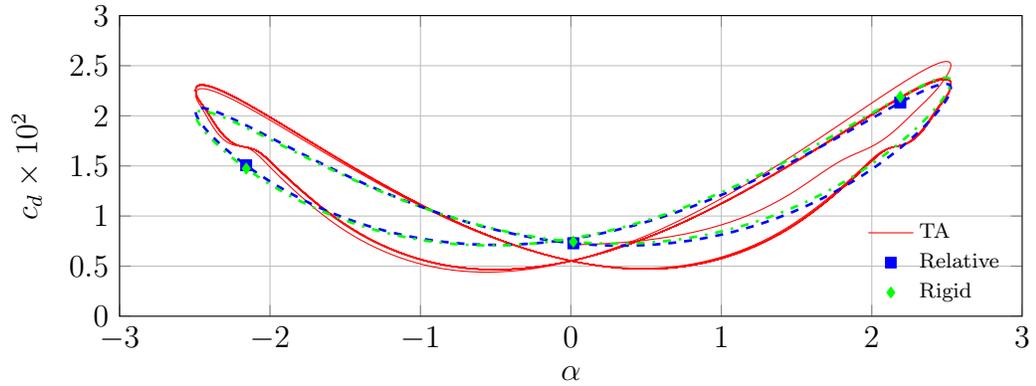
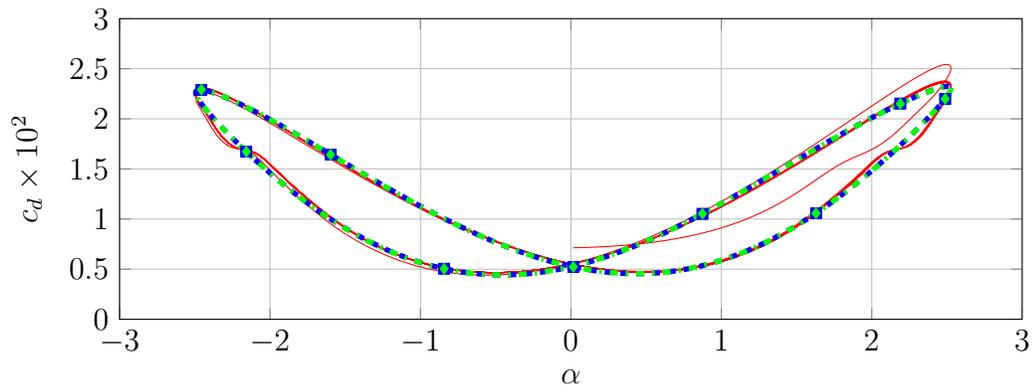
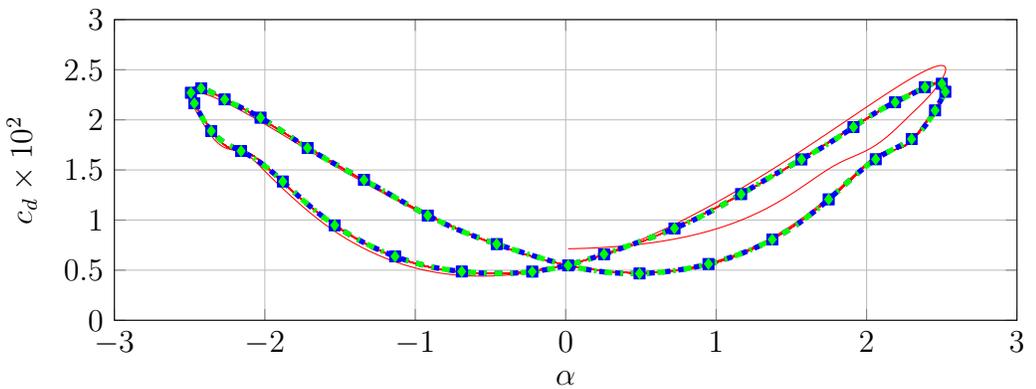
(a) $N = 3$ (b) $N = 9$ (c) $N = 33$

Figure 6.21: Inviscid Transonic Pitching Airfoil. Time-Spectral versus time-accurate drag coefficients for $N \in \{3, 9, 33\}$. Time-accurate solution is plotted in red from steady-state startup. Blue squares and green diamonds locate the drag coefficient values at the Time-Spectral collocation points for relative- and rigid-body motion, respectively. Drag coefficient signals evaluated from reconstructed Time-Spectral solutions are plotted with hashed lines. Corrected experimental drag data from the AGARD 702 Report are not available.

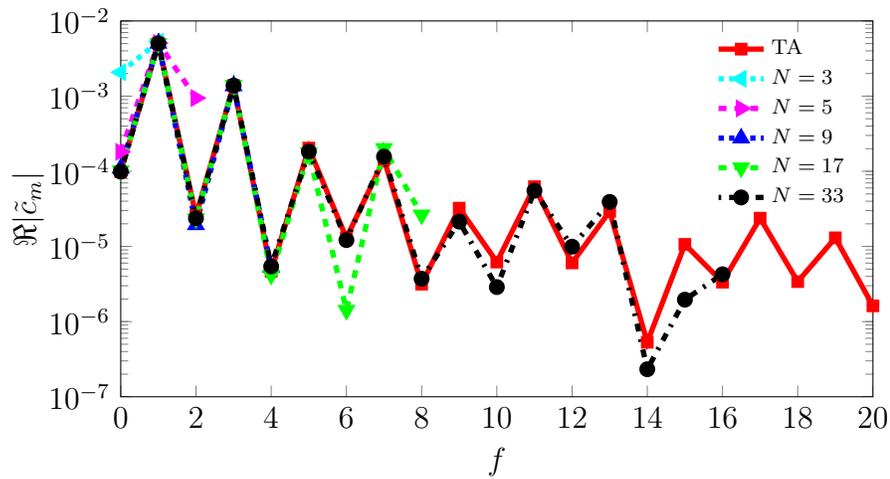
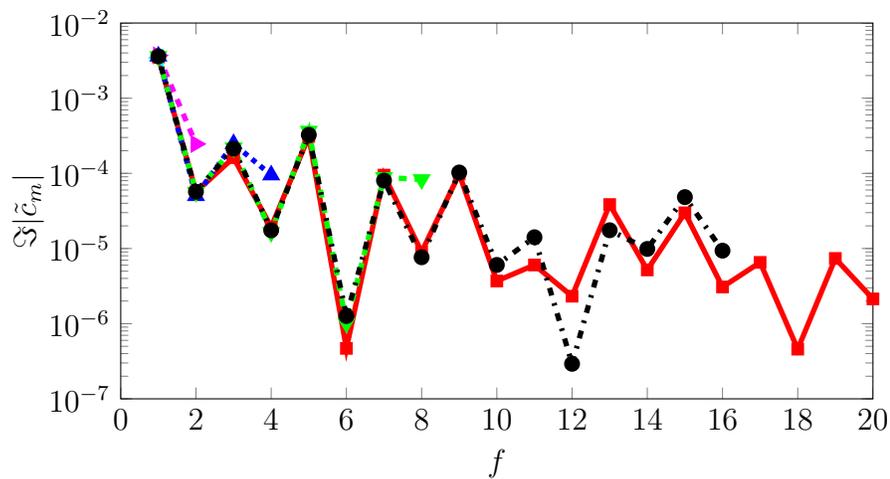
(a) Amplitude, $\Re|\tilde{c}_m|$ (b) Phase, $\Im|\tilde{c}_m|$

Figure 6.22: Inviscid Transonic Pitching Airfoil. Spectrum of Time-Spectral and time-accurate pitching moment coefficients, $\tilde{c}_m = \Phi^{-1}c_m$, for the first twenty harmonics of the relative-motion case.

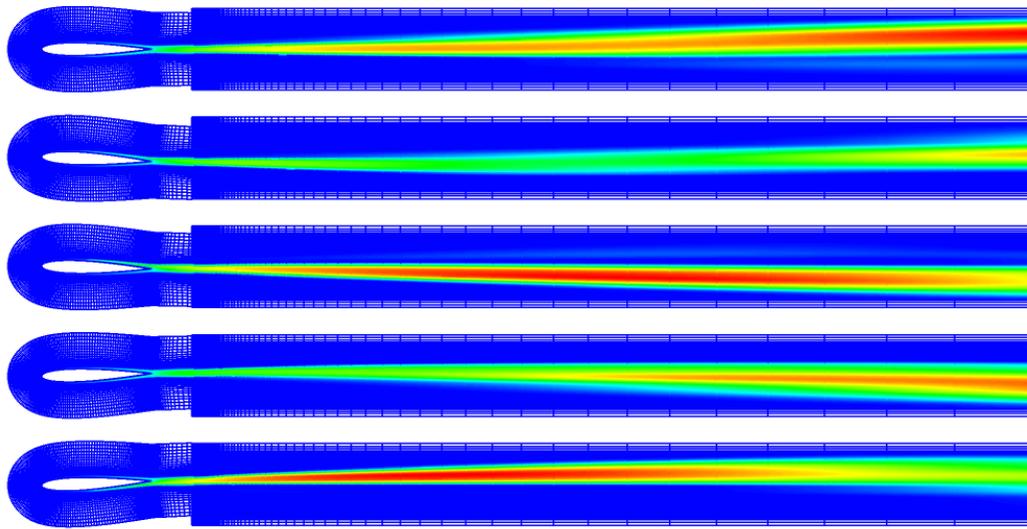


Figure 6.23: Turbulent RANS Transonic Pitching Airfoil. Visualization of undamped turbulent eddy viscosity, $\tilde{\nu}$, on the near-body grid and in the wake region at the temporal collocation points for a representative case of $N = 5$; $\tilde{\nu}$ can vary several orders of magnitude in a short time frame and must be limited in certain circumstances to avoid spurious undershoots resulting from Gibbs' phenomenon. The figure shows every second grid point removed in the normal direction for clarity.

Oscillations in the gradient are manifested at the discrete sample points which can be observed in Fig. 6.24b resulting in large overshoots. This issue is compounded by the fact that undamped eddy viscosity remains close to zero outside the wake, so that modest overshoots in the derivative may violate the positivity constraint on $\tilde{\nu}$. The current example uses eddy viscosity in the Spalart-Allmaras model, however, similar issues are encountered for turbulent kinetic energy or dissipation in other models. The same phenomenon could also occur in the fluid equations for high-speed cases. Note that this issue applies to Time-Spectral treatment of turbulence in general as it occurs downstream from the body on nodes with complete time histories.

Applying a limiter to the temporal differentiation operator maintains positivity on $\tilde{\nu}$ without sacrificing accuracy. The current limiter sets the Time-Spectral temporal residual to zero if it is estimated to drive $\tilde{\nu}$ negative. Figure 6.24a demonstrates good agreement between the Time-Spectral and time-accurate solutions for $\tilde{\nu}$ at the collocation points. Slight ringing does occur in the constant region but $\tilde{\nu}$ remains positive.

Only the relative-motion Time-Spectral scheme is applied to the viscous case as the hybrid scheme already demonstrated its ability to match the results computed with the conventional Time-Spectral scheme for the inviscid rigid-body case. The pitching moment and drag coefficients computed from the Time-Spectral simulations converge to those computed in unsteady mode with increasing numbers of temporal collocation points (Figures 6.25 and 6.26). Note that $N = 33$ samples are required to resolve the kink in the pitching moment and drag coefficient signals at $\alpha \approx \pm 2^\circ$ (Figures 6.25c and 6.26c), as with the inviscid case. Less agreement is observed with the experimental data in the pitching moment coefficient polars (Figure 6.25) when compared to the inviscid case from the previous section but this discrepancy also occurs in the time-accurate case suggesting a physical modeling issue. While experimental data are included for reference, the crucial point remains that the Time-Spectral results converge monotonically to those computed with the time-accurate solver.

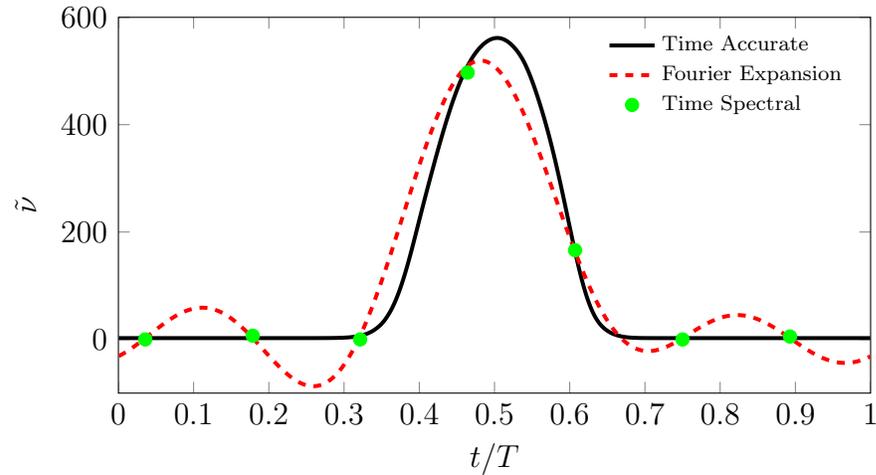
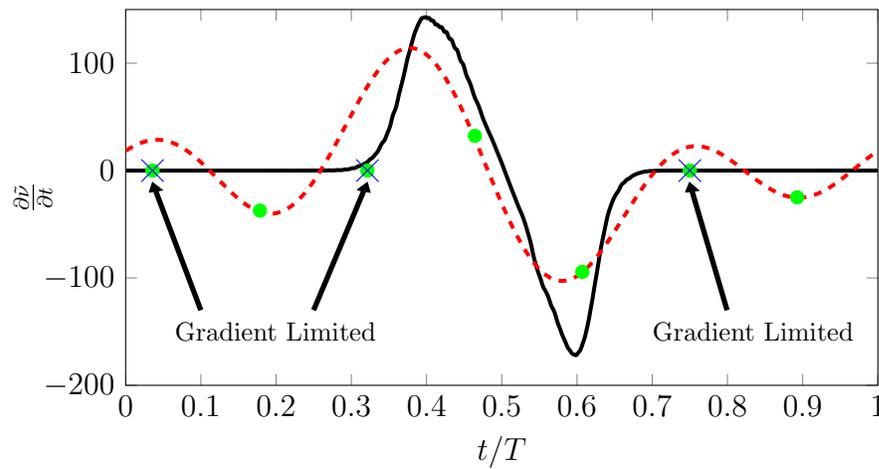
(a) Turbulent variable, $\tilde{v}(t)$ (b) Time derivative, $\frac{\partial \tilde{v}}{\partial t}(t)$

Figure 6.24: Turbulent RANS Transonic Pitching Airfoil. (a) Turbulent solution variable, $\tilde{v}(t)$, and (b) its temporal derivative, $\frac{\partial \tilde{v}}{\partial t}$, are plotted for a point in the wake over the period of oscillation. BDF2 and Fourier operators differentiated the continuous and discrete data, respectively. Note, Gibbs' phenomenon in the Fourier representation of the signal are present in the derivative where large oscillations lead to an inaccurate representation of $\frac{\partial \tilde{v}}{\partial t}$. Time-Spectral results plotted in green. Blue crosses indicate limited gradients.

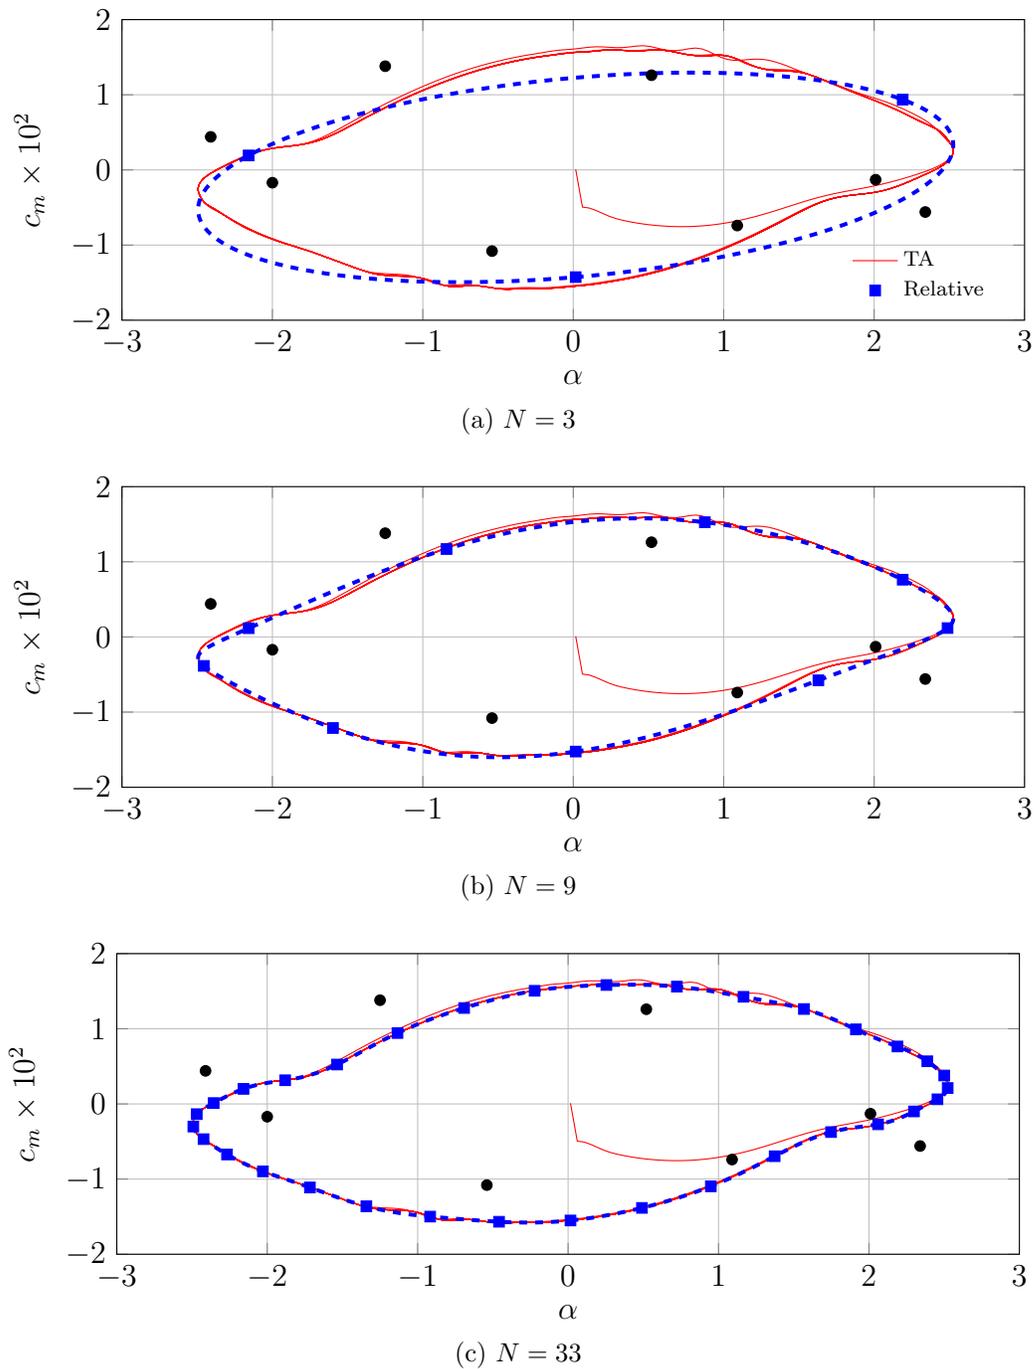


Figure 6.25: Turbulent RANS Transonic Pitching Airfoil. Time-Spectral versus time-accurate pitching moment coefficients for $N \in \{3, 9, 33\}$. Time-accurate solution is plotted in red from steady-state startup. Blue squares locate the moment coefficient values at the Time-Spectral collocation points for relative-body motion. Moment coefficient signals evaluated from reconstructed Time-Spectral solutions are plotted with hashed lines. Experimental data from the AGARD 702 Report are plotted with black dots.

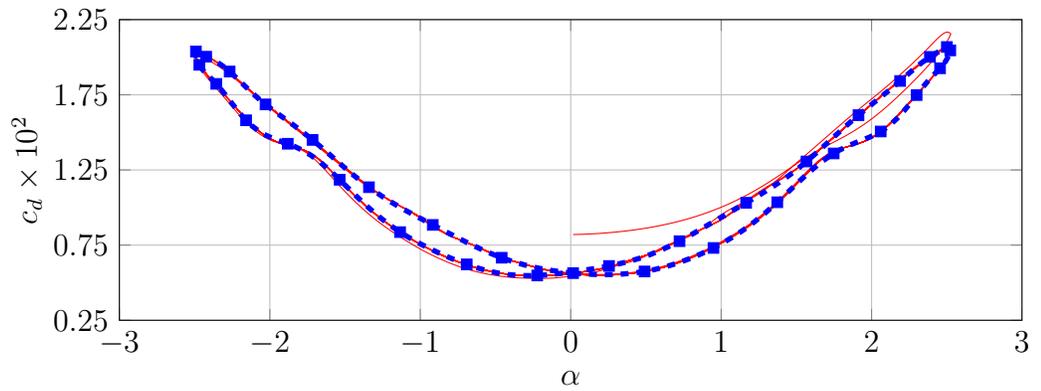
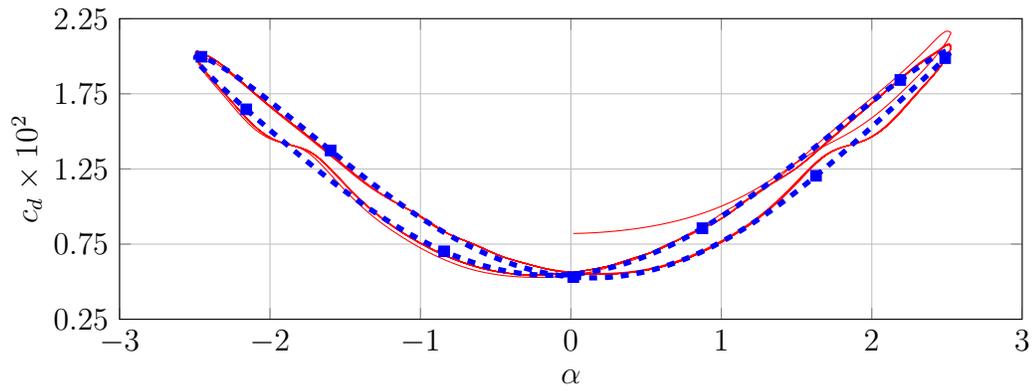
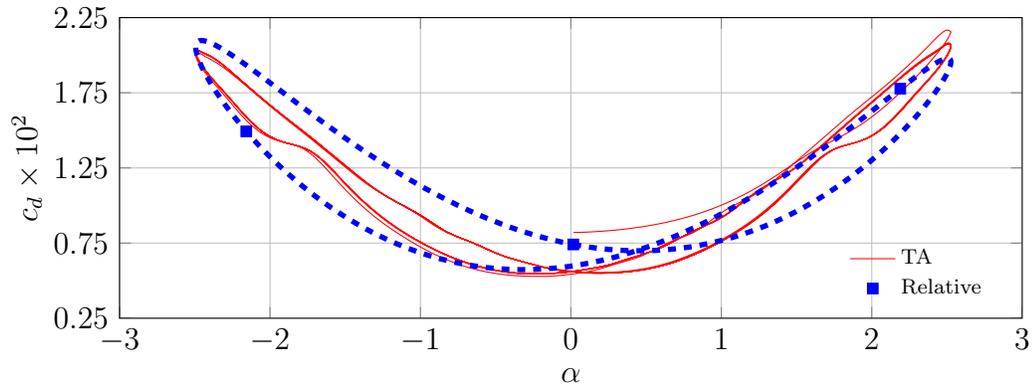


Figure 6.26: Turbulent RANS Transonic Pitching Airfoil. Time-Spectral versus time-accurate drag coefficients for $N \in \{3, 9, 33\}$. Time-accurate solution is plotted in red from steady-state startup. Blue squares locate the drag coefficient values at the Time-Spectral collocation points for relative-body motion. Drag coefficient signals evaluated from reconstructed Time-Spectral solutions are plotted with hashed lines. Corrected experimental drag data from the AGARD 702 Report are not available.

6.2 Three-Dimensional V-22 Tiltrotor

The present section demonstrates the ability of the proposed Time-Spectral approach to successfully match time-accurate predictions of a realistic three-dimensional application – the isolated V-22 Osprey tiltrotor depicted in Fig. 6.27a. Rotorcraft performance is primarily assessed on the converged periodic steady-state flowfields of both hover and forward flight, and therefore both cases are included.¹ Selection of the high-stiffness and relatively low aspect ratio V-22 rotor permits the assumption of non-deforming blades, removing the need to incorporate aeroelastic effects. However, aeroelasticity is a critical component of rotorcraft performance analysis in general and should be included either directly, in a fluid-structure formulation, or with a comprehensive coupling procedure. Research on aeroelasticity in the context of the Time-Spectral method is ongoing [25, 23, 26, 27], but including aeroelasticity is beyond the current scope. Additionally, the V-22 rotor is not typically used in forward (edgewise) flight as it operates in either hover mode or airplane (propeller) mode. Nevertheless, forward-flight Time-Spectral results are compared to those computed in an unsteady environment to evaluate the predictive capabilities of the Time-Spectral approach for that flight regime.

Previous applications of Time-Spectral solvers to rotorcraft problems include the work of Ekici et al. [105], Tatossian and Nadarajah [15, 16], Choi et al. [23, 24] and Yang et al. [38]. Ekici et al. demonstrated the ability for the HDHB method to successfully calculate three-dimensional solutions for the rotor of Caradonna and Tung [106] in hover and lifting and non-lifting forward flight. Tatossian and Nadarajah performed adjoint-based design optimization of the Caradonna and Tung rotor using an NLFD solver for both hover and forward flight. Choi and Datta [23] implemented exact fluid-structure coupling for Time-Spectral calculations and Choi et al. [24] performed adjoint-based aerodynamic shape optimization using a Time-Spectral solver for the UH-60A rotor. Yang et al. demonstrated the ability for an augmented

¹Rotorcraft performance is also gauged on transient maneuver characteristics which cannot be predicted with the standard Time-Spectral method. Mavriplis and Yang [37] extended the Time-Spectral method to treat quasi-periodic flows, such as rotorcraft maneuver, by applying polynomial subtraction.

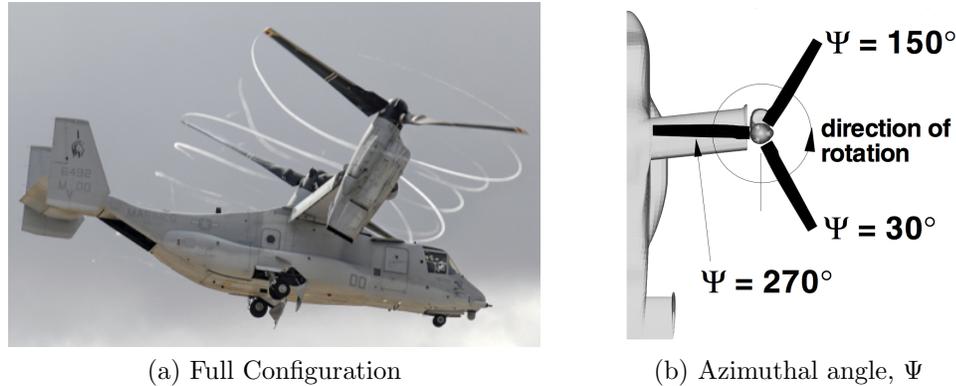


Figure 6.27: V-22 Osprey (a) full configuration [103] (b) definition of azimuthal angle, Ψ , taken from [104].

BDF/Time-Spectral algorithm to resolve quasi-periodic maneuver for the UH-60A rotor.

6.2.1 Hover

The V-22 Osprey tiltrotor has been analyzed for hover extensively – both experimentally and numerically. The quarter-scale Tilt Rotor Aeroacoustic Model (TRAM) was tested at the Duits-Nederlandse Windtunnel (DNW) with experimental details provided in [107, 108]. The isolated rotor geometry is used for all calculations presented which is consistent with the numerical experiments detailed in [104, 109, 110, 103]. With respect to overset grids, the typical solution procedure involves revolving the body-fitted curvilinear grids through a system of stationary background grids. Despite the relatively low aspect ratio of the V-22 blades, a successful Time-Spectral simulation using this approach would require retaining a large number of harmonics due to the high-frequency response induced on the fixed background grids by the relatively low-frequency signals on the near-body grids moving with the blades. This is analogous to the two-dimensional plunging airfoil result shown in Fig. 6.9 where a relatively low-frequency signal on moving near-body grids induces a much higher-frequency response on stationary background grids. Alternatively, rigid-body motion can be achieved by rotating the background grids with the rotor blades. Since hover

can be well estimated as a steady flow in the rotating frame (See [104] for non-inertial reference frame results), it is expected that only a limited number of temporal modes are required when all grids rotate in unison. However, employing rigid-body motion fails to introduce any dynamically-blanked nodes. The case of hover is used to validate the three-dimensional OVERFLOW Time-Spectral implementation. Dynamically-blanked nodes are introduced in the forward-flight case discussed in the upcoming section.

The objective of the hover test case is to validate the three-dimensional Time-Spectral OVERFLOW solver against the standard OVERFLOW time-accurate solver, by using the same grid systems and spatial discretizations to isolate the effects of varying the temporal discretization. The grid system includes three blades, each composed of an C-mesh blade grid and two cap-grids, one each for the root and tip, as depicted in Fig. 6.28a. The C-mesh consists of $181 \times 94 \times 65$ points in the chordwise, radial and body-normal directions, respectively. Each blade, root and tip grid contains 1.1, 0.48 and 0.33 million grid points, respectively, for a total of 5.7 million near-body grid points. The off-body grid system, depicted in Fig. 6.28b, is comprised of a hierarchy of forty-six isotropic Cartesian grids spanning 15 rotor radii (R) to the farfield boundaries in each direction, totaling 21.3 million mesh points. Therefore, the entire grid system is comprised of 27.1 million grid points (See Table 6.3). The finest Cartesian L1 grid spans roughly $\pm 1.2R$ in both the x - and y -directions and $-1.9R$ and $+0.38R$ in the z -direction, with spacing of $0.1c_{\text{tip}}$. While a center body is included in the experiment and previous computations, it is omitted in the current work.

Type	Grid Points	Quantity	Total
Blade	1,105,910	3	3,317,730
Root	483,405	3	1,450,215
Tip	332,995	3	998,985
Cartesian		46	21,283,543
Total		55	27,050,473

Table 6.1: Isolated V-22 Osprey Tiltrotor in Hover. Grid statistics.

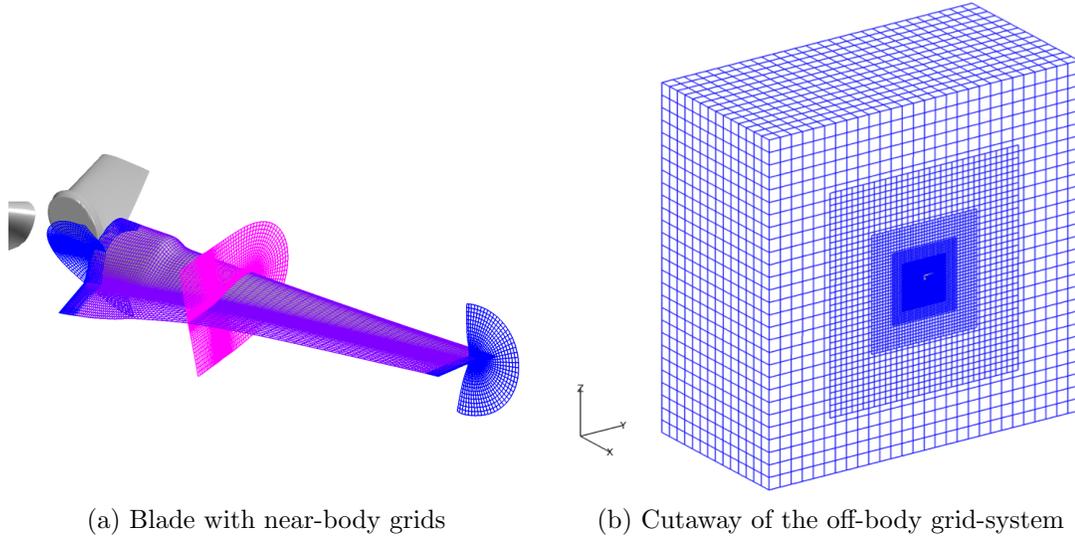


Figure 6.28: Isolated V-22 Osprey Tiltrotor in Hover. Geometry and grid system. (a) Each blade is wrapped in an C-mesh with tip and root cap grids. (b) the rotor is embedded within a hierarchy of isotropic Cartesian off-body grids. The domain of the off-body system is cut away to reveal the structure of the isotropic Cartesian grids.

Previous numerical calculations of the V-22 TRAM hover case employed a variety of turbulence model treatments, but the current results were computed following the *laminar off-body* procedure outlined by Potsdam and Pulliam [109]. The one-equation SA turbulence model is solved on all body-conforming near-body grids. However, the turbulence model production source terms are deactivated on the Cartesian background grids to prevent the eddy viscosity from growing unbounded in the vortex-dominated wake and feeding back to the flow in the vicinity of the rotor.

The hover test case uses a uniform collective of 14° with a tip speed of $M_{\text{tip}} = 0.625$ and tip Reynolds number of $Re_{\text{tip}} = 2.1$ million, which is consistent with the primary TRAM test case in [109]. Qualitative results from both time-accurate and single-harmonic Time-Spectral V-22 TRAM hover calculations are presented in Figs. 6.29-6.30. Figure 6.29 plots iso-contours of vorticity magnitude, as well as a cutting plane of the same functional. Figure 6.30 plots the same iso-contours of vorticity magnitude over a cutting plane of undamped eddy viscosity. The time-accurate and

Time-Spectral calculations demonstrate strong agreement in both cases.

In general, dozens of periods are required to converge certain functionals of interest in a time-accurate calculation. The Figure of Merit, FM , is an established scalar performance measure for rotorcraft in hover that measures rotor efficiency; it is a ratio of the ideal power, derived from momentum theory, to the actual (computed) power.

$$FM = \frac{1}{\sqrt{2}} \frac{C_T^{3/2}}{C_Q} \quad (6.1)$$

The Time-Spectral simulation is steady in nature and converges directly to the space-time solution.

	C_T	C_Q	FM
TS	0.11379×10^{-1}	0.15320×10^{-2}	0.5602
TA	0.11382×10^{-1}	0.15320×10^{-2}	0.5605

Table 6.2: Isolated V-22 Osprey Tiltrotor in Hover. Force and moment coefficients and Figure of Merit for the single-harmonic ($N = 3$) Time-Spectral (TS) and time-accurate (TA) calculations.

Table 6.2 provides the force and moment coefficients and the Figure of Merit for the time-accurate and single harmonic Time-Spectral calculations. The Figure of Merit computed using the Time-Spectral method with a single harmonic agrees with the time-accurate result to five percent of a percent. However, the Figure of Merit is substantially lower than those presented in [104, 109, 110, 103]. This discrepancy is attributable to the use of only second-order central differencing of the convective terms, as opposed to higher-order schemes in the references cited. Preliminary results do indicate an improvement with increasing the spatial discretization order of accuracy. Nevertheless, the case of hover has validated the three-dimensional Time-spectral implementation.

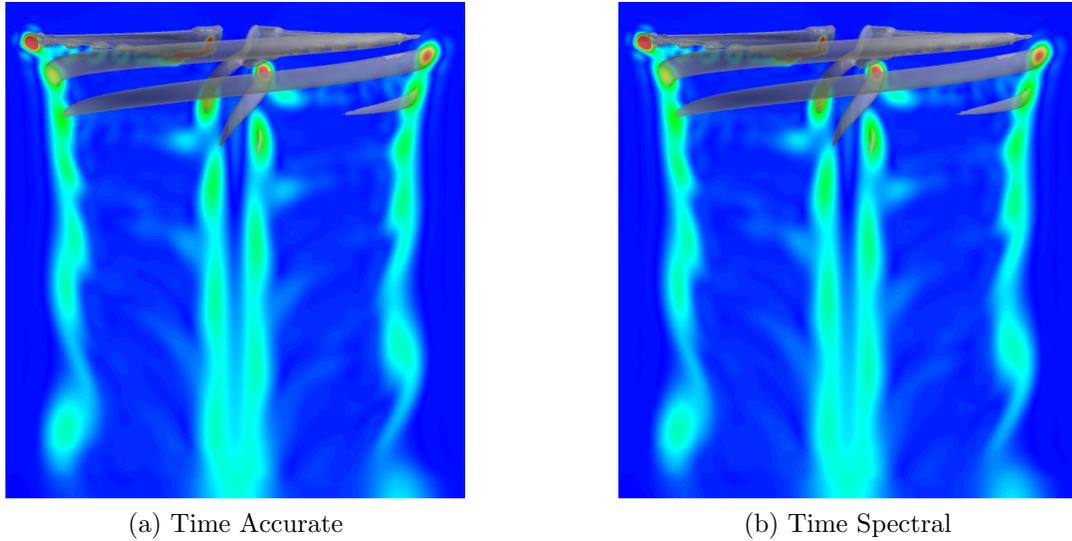


Figure 6.29: Isolated V-22 Osprey Tiltrotor in Hover. Iso-contours of vorticity for time-accurate and Time-Spectral simulations. Contours of vorticity magnitude are plotted in a constant y -plane, ranging from a value of zero in blue to 0.01 in red.

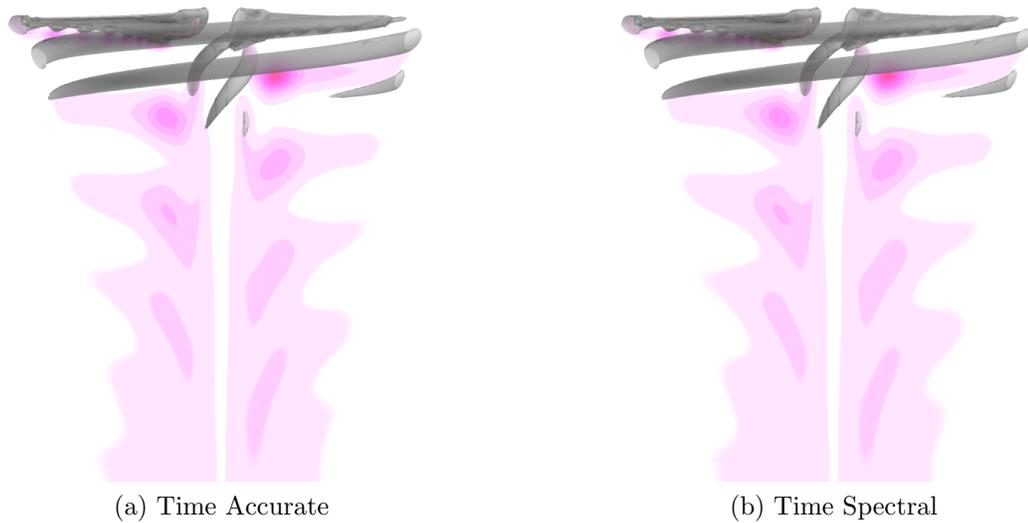


Figure 6.30: Isolated V-22 Osprey Tiltrotor in Hover. Iso-contours of vorticity for time-accurate and Time-Spectral simulations. Contours of undamped eddy viscosity, $\tilde{\nu}$, are plotted in a constant y -plane, ranging from a value of zero in white to 900 in black.

6.2.2 Forward (Edgewise) Flight

The hover case of the previous section prescribed the rigid motion of the isolated V-22 rotor with its off-body grid system, resulting in an absence of dynamically-blanked nodes. Rotors in forward (edgewise) flight experience an imbalance of dynamic pressure on the advancing ($0 \leq \Psi \leq 180$) and retreating ($180 \leq \Psi \leq 360$) segments of a rotation (See Figure 6.27b). Left untreated, this leads to an undesired rolling moment about the longitudinal axis of the vehicle. Cyclic pitch is used to mitigate the force imbalance by dynamically changing the pitching amplitude of the blade as it rotates. On the advancing side, where blades experience a higher dynamic pressure, pitch is decreased to reduce the resulting normal force. Conversely, blades pitch upwards on the retreating side to increase the normal force in the presence of lower dynamic pressure. While the background grids are still prescribed to rotate with the rotor, the relative pitching motion between the blades and background grids results in dynamically-blanked nodes, thereby requiring the hybrid Time-Spectral scheme.

While the V-22 Osprey is not typically employed in forward (edgewise) flight because its thrust is derived by tilting its rotors forward into propellor mode, the edgewise-flight case provides an opportunity to gauge the three-dimensional hybrid Time-Spectral scheme for a rigid blade. A reasonable selection of parameters were chosen for the values of collective and cyclic pitch (estimated from Young et al. [107]). Additionally, the rotor shaft axis is chosen to be vertical. While the rotor is not trimmed (e.g. with a comprehensive analysis tool), the prescription of fixed collective and cyclic pitch serve as a representative estimate of rotor trim. Additionally, the rotors move in prescribed fashion, ignoring the force-induced blade dynamics typically associated with high-fidelity rotorcraft calculations. The current objective is to compare Time-Spectral results to those computed with the standard time-accurate flow solver. Aeroelasticity is a complicating factor that obscures the current objective, but should be included in full-scale rotorcraft simulations.

Time-accurate and Time-Spectral forward-flight V-22 TRAM calculations are presented for an advance ratio, $\mu = 0.2$, where $M_\infty = \mu M_{\text{tip}}$, using the same tip Mach and Reynolds numbers as the hover case. The test case employs a fixed collective, $\theta_0 = 10^\circ$,

and longitudinal and lateral cyclic pitch amplitudes, $\theta_{C1} = 3^\circ$ and $\theta_{S1} = -5^\circ$, respectively, where, the pitch amplitude, $\theta(\Psi)$, is prescribed by the following relationship.

$$\theta(\Psi) = \theta_0 + \theta_{C1} \cos(\Psi) + \theta_{S1} \sin(\Psi) \quad (6.2)$$

Azimuthal angle, Ψ , is taken to be zero in the freestream direction as depicted in Fig. 6.27b.

The off-body grid system is modified for the forward flight case because vorticity shedding off of the rotor is convected downstream, as opposed to straight down as for the case of hover. Thus, the Cartesian off-body grid system is compressed to provide a wider L1 grid near the rotor to capture vortical content as it convects downstream. The off-body grid system is still comprised of a hierarchy of 46 isotropic Cartesian meshes combining for a total of 20.3 million grid points. Therefore, the entire grid system is comprised of 26.1 million grid points (See Table 6.3), which is nearly equivalent to the amount used for hover. The L1 grid spans roughly $\pm 1.65R$ in both the x - and y -directions and $-0.9R$ and $+0.28R$ in the z -direction, with the same spacing of $0.1c_{tip}$ as the hover case. Additionally, spectral vanishing viscosity is required to satisfactorily converge the Time-Spectral calculations.

Type	Grid Points	Quantity	Total
Blade	1,105,910	3	3,317,730
Root	483,405	3	1,450,215
Tip	332,995	3	998,985
Cartesian		46	20,283,716
Total		55	26,050,646

Table 6.3: Isolated V-22 Osprey Tiltrotor in Forward Flight. Grid statistics.

Qualitative comparisons between the time-accurate and Time-Spectral forward-flight results are provided in Figs. 6.31-6.32. Figure 6.31 plots iso-contours of vorticity magnitude and a cutting plane of undamped eddy viscosity. The Time-Spectral

solution using only $N = 11$ time samples shows some agreement with the time-accurate result but demonstrates marked disagreement in the turbulent variable. The Time-Spectral solution using $N = 21$ time samples offers significant improvement. The $N = 31$ Time-Spectral solution resolves some of the finer features exhibited in the time-accurate result and shows improved agreement of the turbulent eddy viscosity field. The side views of the four cases, provided in Fig. 6.32, demonstrate similar trends. Unlike the case of hover, the flow is no longer steady in the rotating frame. The *tip* vortices being convected downstream pose a resolution problem for unsteady simulation as a point on the rotating background mesh moves quickly through the vortices. This is yet another example of a case where a relatively low-frequency disturbance in one frame is transformed to a much higher-frequency signal in another frame, which may explain the poor resolution of the tip vortices in the Time-Spectral calculations.

Finally, Figure 6.33 demonstrates quantitative convergence of the Time-Spectral calculations to the time-accurate solution for the instantaneous thrust coefficient, C_T , versus azimuthal angle, Ψ . Despite rather poor qualitative agreement, the $N = 11$ Time-Spectral case resolves the bulk C_T signal. Naturally, agreement improves as the temporal resolution is refined. The $N = 31$ Time-Spectral result matches the time-accurate signal over nearly the entire period.

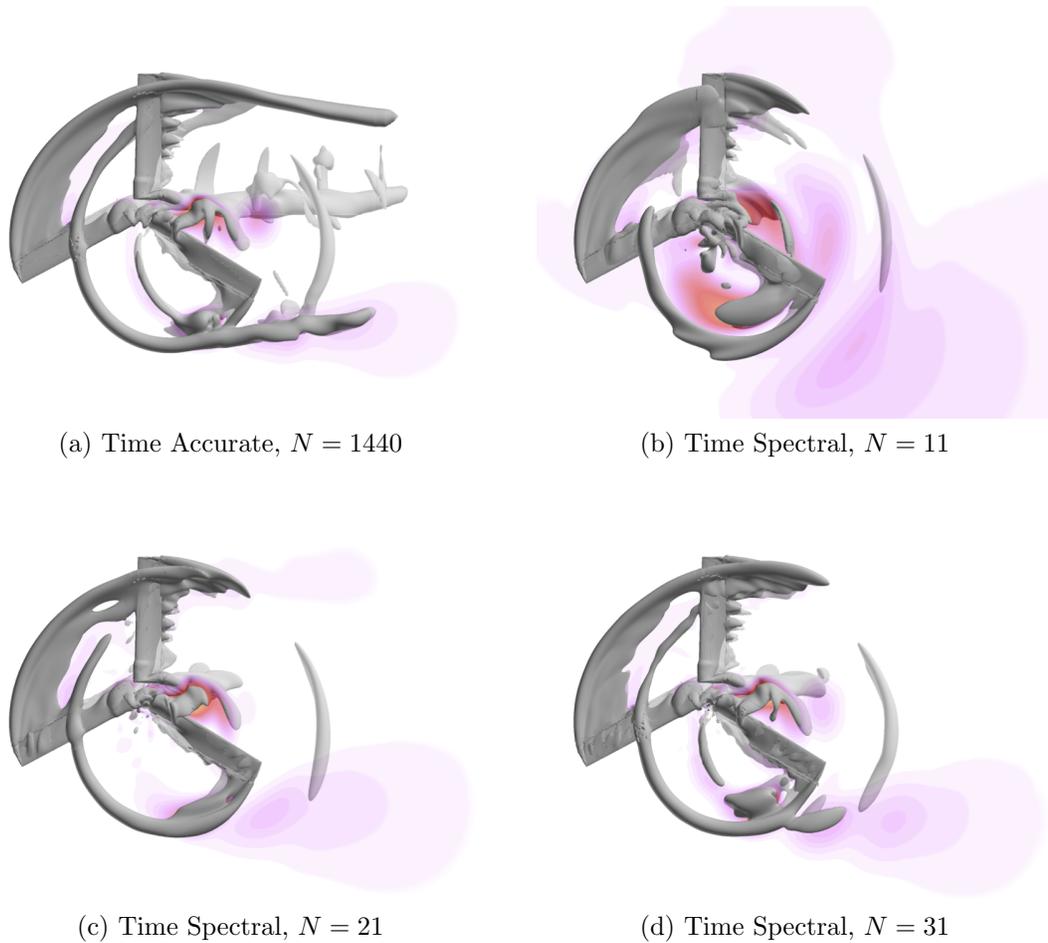


Figure 6.31: Isolated V-22 Osprey Tiltrotor in Forward Flight. Top view of iso-contours of vorticity for time-accurate and Time-Spectral calculations with advance ratio $\mu = 0.2$ and tip Mach number $M_{\text{tip}} = 0.625$. Time-Spectral simulations using $N \in \{11, 21, 31\}$ are presented to demonstrate the convergence towards the time-accurate result with the addition of temporal modes. Contours of undamped eddy viscosity are plotted in a constant z plane.

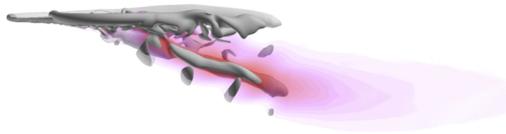
(a) Time Accurate, $N = 1440$ (b) Time Spectral, $N = 11$ (c) Time Spectral, $N = 21$ (d) Time Spectral, $N = 31$

Figure 6.32: Isolated V-22 Osprey Tiltrotor in Forward Flight. Side view of iso-contours of vorticity for time-accurate and Time-Spectral calculations with advance ratio $\mu = 0.2$ and tip Mach number $M_{\text{tip}} = 0.625$. Time-Spectral simulations using $N \in \{11, 21, 31\}$ are presented to demonstrate the convergence towards the time-accurate result with the addition of temporal modes. Contours of undamped eddy viscosity are plotted in a constant y plane.

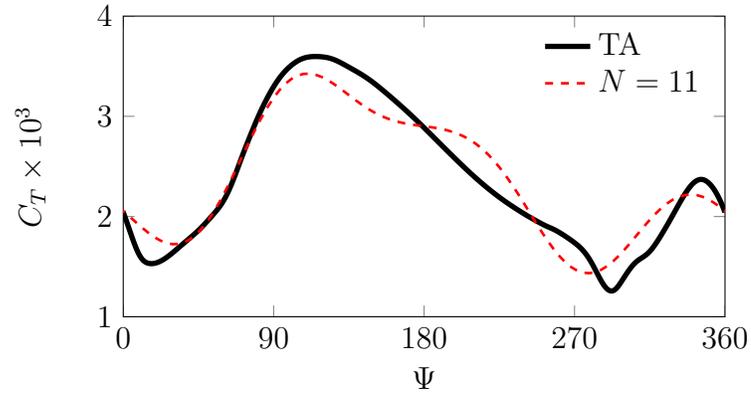
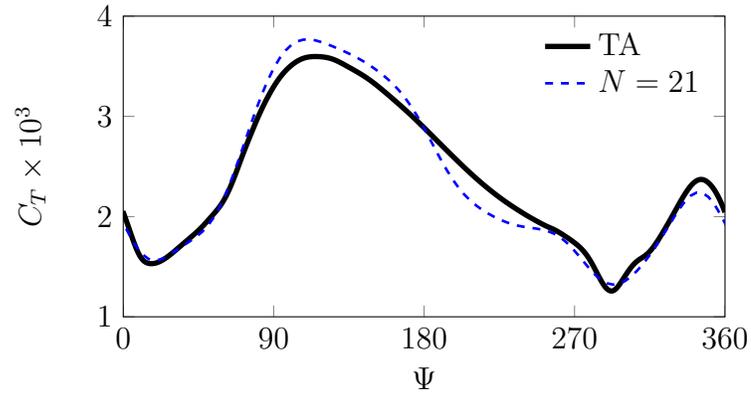
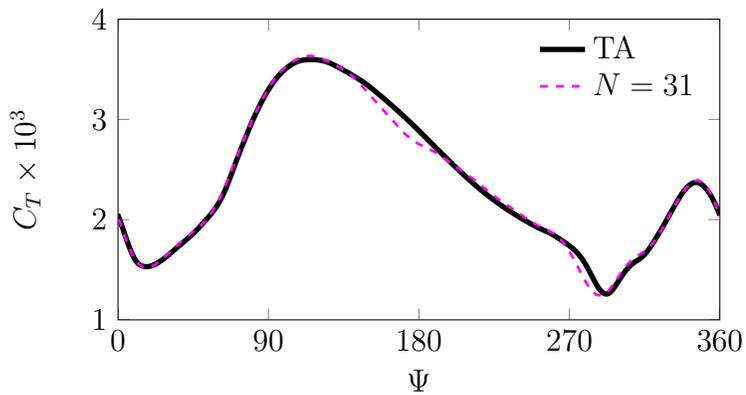
(a) Five harmonics, $N = 11$ (b) Ten Harmonics, $N = 21$ (c) Fifteen harmonics, $N = 31$

Figure 6.33: Isolated V-22 Osprey Tiltrotor in Forward Flight. Reconstruction of Time-Spectral thrust coefficient, C_T , for $N \in \{11, 21, 31\}$ versus the time-accurate result using a time step of equivalent to 0.25° ($\Delta t = T/1440$).

Chapter 7

Conclusions

This chapter summarizes the dissertation and discusses some of the general conclusions of the research including some of the observed limitations. The chapter closes with a discussion of potential areas of future work.

7.1 Summary

Development of the Hybrid Time-Spectral Scheme

A novel overset hybrid Time-Spectral scheme has been developed to accurately treat dynamically-blanked nodes arising from relative motion. A key feature of Fourier pseudospectral methods, and the Time-Spectral method in particular, is the relative ease with which such schemes can be incorporated into existing flow solvers. With the desire of leveraging a mature codebase, it was determined that the temporal collocation points at dynamically-blanked nodes would be required to share the same evenly-spaced time samples as their statically-blanked, Fourier-based counterparts. Using any other distribution of time samples would require costly and potentially ill-conditioned temporal interpolation on a node-to-node basis. For this reason, the temporal collocation points were a priori constrained to the standard uniform distribution.

In Chapter 4, three primary approaches were considered to represent the solution

at such nodes: a global Fourier least-squares approach, a localized approach expanding the solution within partitioned intervals of consecutively-defined nodes with independent bases, and a mixed approach that combines the two. It was demonstrated on both simple test-cases and numerical experiments in §6.1.1 that the local description of the solution was the most promising option; the barycentric rational interpolant was established as the default treatment for the remainder of the thesis. However, the hybrid Time-Spectral scheme has been developed in a general manner that can incorporate future treatments provided the availability of a linear differentiation operator.

Time-Spectral Implementation within the OVERFLOW Solver

The implicit overset Reynolds-averaged Navier-Stokes (RANS) OVERFLOW solver has been augmented with Time-Spectral capability as described in Chapter 5. This approach leveraged the existing codebase of a mature flow solver, thus enabling the calculation of complex three-dimensional configurations, and facilitating a direct comparison between the Time-Spectral and time-accurate schemes.

The OVERFLOW implementation applies the hybrid Time-Spectral scheme developed in Chapter 4 to cases involving relative motion and reverts to the standard Time-Spectral method in the absence of relative motion. The incorporation of the Time-Spectral method within OVERFLOW's approximate-factorization (AF) scheme enabled a reduction in computational complexity from $\mathcal{O}(N^3)$ to $\mathcal{O}(N \log N)$, which can provide significant cost savings if a large number of temporal modes is desired. The existing spatial multigrid algorithm has been extended in the Time-Spectral implementation to facilitate convergence acceleration and a semi-implicit Time-Spectral approach has been incorporated within OVERFLOW's loosely-coupled turbulence model. A temporal dissipation operator derived from barycentric rational interpolants was developed to stabilize calculations involving dynamically-blanked nodes by damping the solution at interval boundaries.

Numerical Results

A number of numerical experiments were conducted to validate the proposed hybrid Time-Spectral method. An array of two-dimensional test cases was selected that could be computed using either rigid or relative motion in order to compare the hybrid Time-Spectral approach to the well-established, standard Time-Spectral method. Three-dimensional calculations of the isolated V-22 Osprey tiltrotor demonstrated the ability of the Time-Spectral implementation to resolve realistic three-dimensional flows on overset grids. Accuracy of the Time-Spectral calculations was assessed by comparisons to well-resolved time-accurate calculations. All spatial and algorithmic parameters were held constant between the time-accurate and Time-Spectral calculations to isolate the effect of varying the temporal discretization.

Computational efficiency has not been explicitly addressed. Fourier pseudospectral schemes have demonstrated a dramatic reduction in the computational cost for time-periodic flows [13]. Here, the objective was to assess the capability of the hybrid Time-Spectral scheme (for the general case of overset relative motion) to match the performance of the traditional Time-Spectral method under the assumption that the standard Time-Spectral method was a competitive approach. As such, demonstrating cost effectiveness was not included in the current scope. However, the upcoming section describing possible areas of future work includes a number of algorithmic strategies to exploit untapped sources of additional computational efficiency. A comprehensive cost analysis would be more relevant upon the implementation of such algorithmic improvements and addressing programming inefficiencies.

Two-Dimensional Oscillating Airfoils

The variety of relatively inexpensive two-dimensional test cases presented in §6.1 explored several relevant areas of interest. The inviscid plunging NACA 0012 airfoil case (§6.1.1) was used to differentiate the various treatments for expanding the solution at dynamically-blanked nodes. It appeared that the bounded interval approach was the most successful treatment in general. The barycentric rational interpolant (BRI) was therefore selected as the default basis. Splines also demonstrated success,

but the rational interpolant was selected due to the ability to explicitly generate its differentiation operator for an arbitrary approximation order, d . As previously mentioned, the general construction of the hybrid Time-Spectral can easily incorporate improved treatments for dynamically-blanked nodes, provided there exists a corresponding linear differentiation operator.

The fundamental issue of the transmission of relatively low-frequency signals in one frame, to high-frequency signals in another frame was also demonstrated for the inviscid plunging case when the grids moved relative to each other. Because of the large plunging motion, the smooth transitions of the solution at nodes on the near-body grid were experienced more rapidly on a node on the stationary background grid. The higher-frequency content therefore required the retention of more temporal modes on the background grid.

The high-frequency laminar plunging NACA 0012 airfoil test cases demonstrated the ability for vanishing viscosity (both spectral- and BRI-based) to adequately damp otherwise unstable calculations (thrust-producing case with $St = 0.6$). Additional two-dimensional test cases qualified the Time-Spectral implementation to resolve transonic and turbulent flows that are relevant to the three-dimensional rotorcraft calculations.

Three-Dimensional V-22 Tiltrotor

The V-22 Osprey tiltrotor served as the three-dimensional validation case for the Time-Spectral OVERFLOW implementation. The hover test case was computed using rigid motion where all grids rotated in unison. This was done to avoid the highly-disparate grid velocities between the blades and the stationary background grids which is the typical solution procedure for hover in the context of overset grids. A large number of harmonics would be required to resolve the relative-motion hover case, despite the fact that the flow can be well-estimated as steady in the rotating frame. The rigid-motion hover calculation required only a single harmonic to match the time-accurate result nearly identically, but did not address dynamically-blanked nodes.

The forward (edgewise) flight test case employed cyclic pitch to simulate rotor

trim, resulting in relative motion between the blades and the spinning background grids. Hybrid Time-Spectral calculations required SVV to adequately converge the solutions, but demonstrated convergence to the time-accurate calculation with increased temporal resolution.

7.2 Limitations

While the hybrid Time-Spectral method has demonstrated the ability to mirror the performance of the standard Time-Spectral scheme, there exists a primary observed limitation. Numerical experiments with large or high-speed relative motion have exhibited oscillations in the vicinity of overset interfaces. A lack of resolution is the most likely explanation for this; nodes on the near-body grid transmit a relatively low-frequency signal to stationary background grids which experiences a much higher-frequency response. Perhaps the inability of the background grid to resolve the higher-frequency response results in an under-resolved solution that aliases into the lower modes, thereby generating spurious oscillations that collect on the boundaries. This issue was introduced in §6.1.1 but the potential for aliasing issues was not addressed. Aliasing was, in fact, an issue for the thrust-producing case of the laminar plunging NACA 0012 airfoil in §6.1.2 for both the rigid- and relative- motion calculations. The large Strouhal number, $St = 0.6$, results in a large plunging velocity, $\dot{y}(t) = kh \cos(kt)$, which attains its maximum value of $St = 0.6$ in the neutral position. There exists a large disparity between grid velocities on the overset boundaries between the plunging near-body grid and stationary background grid when employing relative motion. Density contours plotted in Fig. 7.1 reveal spurious oscillations in density that occur on the overset boundaries for the relative-motion case. While these oscillations diminish with increased temporal resolution, they are wholly absent from the rigid-motion calculations where the grid velocities on the overset boundaries are equivalent. Interestingly, drag was well-predicted by the under-resolved relative-motion Time-Spectral calculations, despite the obvious errors apparent in Fig. 7.1.

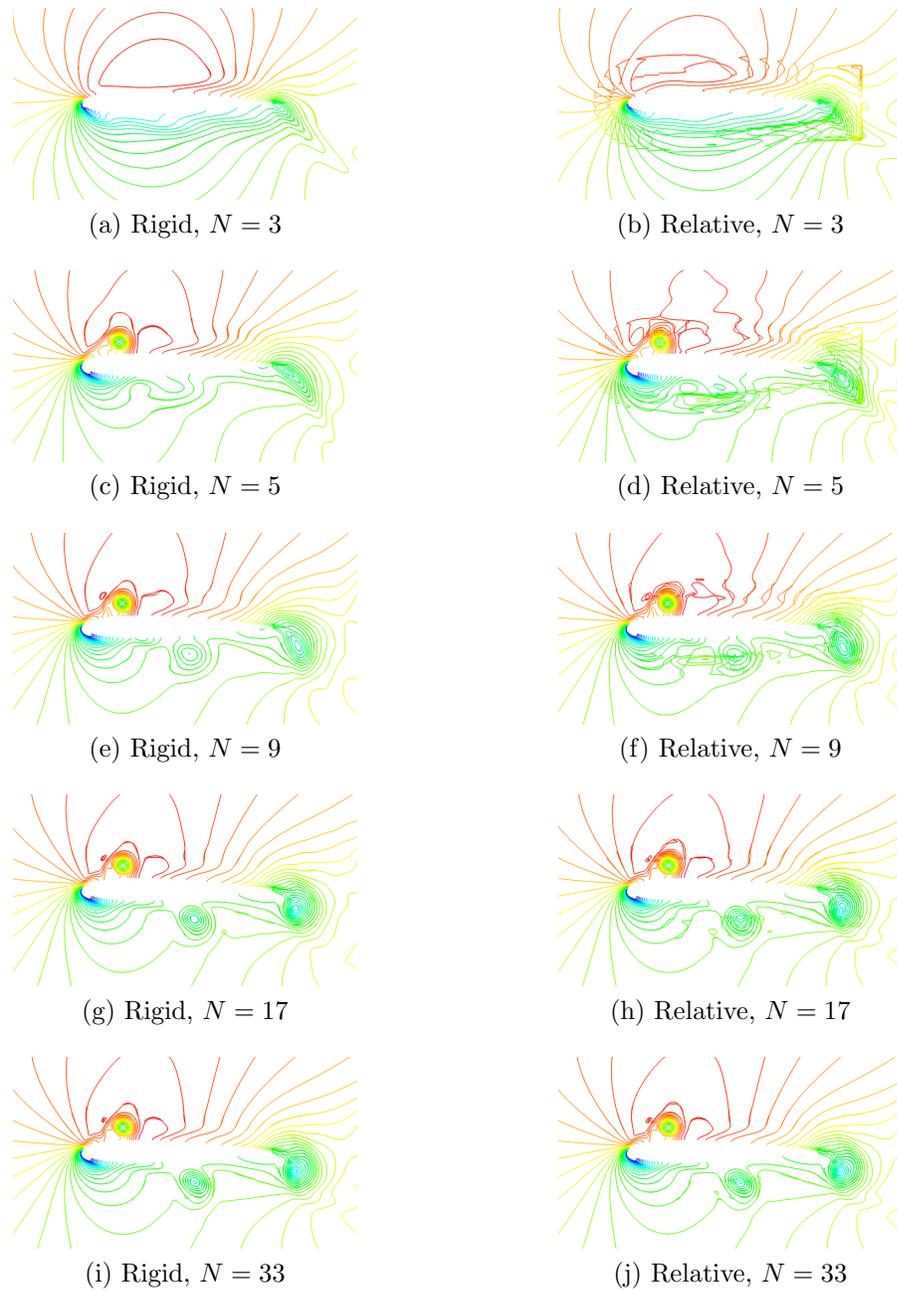


Figure 7.1: Laminar Plunging NACA 0012 Airfoil. Thrust-producing case. Density contours reveal spurious oscillations on overset boundaries for relative motion in the neutral position corresponding to maximum grid velocity. Increased temporal resolution mitigates the oscillations.

The disparate grid velocities between overset interfaces may be a source of considerable conservation error for relative motion. The deformation caused by the transition of blanked nodes moving in and out of the computational domain is not accounted for and OVERFLOW is independently a non-conservative scheme. Under normal operation in the time-accurate mode, the physical time steps are small, which may mitigate the conservation error as it scales with the time step. However, large physical time steps are typical for the Time-Spectral approach. When similarly large time steps are used for both the time-accurate and Time-Spectral approaches, spurious oscillations are observed for both cases.

Increased temporal resolution resolves the issue. However, universal temporal refinement may prove too costly if only a small region experiences high frequencies. To balance the cost of temporal Fourier collocation calculations requiring a relatively large number of modes in a confined region, frequency adaptive methods have been developed and will be expanded upon in the following discussion of potential areas of future work.

7.3 Future Work

This dissertation has demonstrated the ability for the hybrid Time-Spectral scheme to accurately simulate forced periodic flows involving relative motion on overset grids. The Time-Spectral solutions were shown to converge the respective time-accurate solutions provided sufficient resolution, however, efficiency considerations were not explicitly addressed.

In addition to the potential benefits of improved boundary conditions for the bounded interval approach discussed in §4.3.2, the subsequent sections describe various ways to improve the computational efficiency of overset Time-Spectral calculations. Adaptation of the temporal domain can more efficiently allocate the temporal degrees of freedom to more dynamic regions of the flowfield. A spatio-temporal Time-Spectral multigrid approach should improve the convergence acceleration already demonstrated by the spatial algorithm. Exploiting the parallelism in the temporal dimension has been demonstrated by Mavriplis and Yang [37].

Improved Temporal Boundary Conditions

As described in §4.3.2, the bounded interval approach at dynamically-blanked nodes lacks either explicit or implicit boundary conditions. Despite the potential ill-posedness of such a construction, calculations using this approach demonstrate both stability and convergence. However, future work may uncover accuracy and stability improvements resulting from the derivation and application of suitable boundary conditions for the local or mixed approach.

Adaptive Mesh Refinement in Time

It has been demonstrated that calculations involving large-amplitude relative motion often require increased resolution in comparison to their rigid-motion analogues where all grids move in unison. In extreme cases, with highly disparate grid velocities on overset interfaces, spurious oscillations can be generated; temporal refinement has been shown to resolve this issue (See §7.2). Currently, the number of time samples, N , is applied universally throughout the domain. However, the increased temporal resolution required to accurately resolve flows may only be required in a relatively small region (e.g. in the vicinity of the trailing edge of an airfoil).

Here, a strategy is proposed to enable more efficient solution of such cases in the context of overset Time-Spectral solvers. Mosahebi and Nadarajah [33, 34, 35, 36] developed a frequency-adaptation approach for the Nonlinear Frequency Domain (NLFD) method, whereby the solution could be expanded into a different number of frequencies based on the dynamic content of the local solution. Adaptation was accommodated on a point-to-point basis as depicted in Fig. 7.2 for the same thrust-producing $St = 0.6$ plunging airfoil case presented in §6.1.2.

This approach is not applicable to OVERFLOW if the spatial operators are to be leveraged. Changing the number of temporal collocation points on a node-to-node basis would require a total overhaul of the codebase. OVERFLOW performs tasks (e.g. residual evaluation) independently on grid-based zones. Therefore, the number of time instances assigned to each grid is not constrained to be constant across all grids. Grids are typically coarsened in space as their location moves farther from the body as

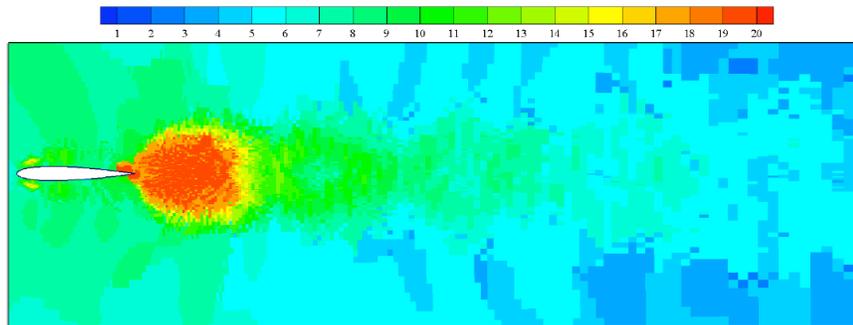


Figure 7.2: Modal distribution of frequency-adaptive solution for the laminar plunging NACA 0012 airfoil at $Re = 1850$ and $St = 0.6$ taken from Mosahebi and Nadarajah [35]. This case was demonstrated in §6.1.2. A large number of modes are required to resolve the solution in the vicinity of the trailing edge of the airfoil. However, very few modes are needed outside this small region suggesting significant savings could be realized from a frequency-adaptive approach. The NLFD adaptive approach refines on a node-to-node basis. A similar technique could be implemented in the overset framework by applying adaptive mesh refinement in time where each grid would have an associated number of time samples. New grids with finer or coarser temporal resolution could be generated while either maintaining the spatial resolution or adapting both in time and space.

they do not have to resolve as sharp spatial gradients in the flow. There is no reason to not apply a similar approach in time; the flow is much less dynamic in the regions away from the body. It would also be beneficial to provide automatic refinement or coarsening enabling the space-time mesh to completely adapt to features of the flow as they evolve. If a particular region is tagged for refinement, a new grid can be generated with more time samples that will then blank out the less temporally-resolved grid. Alternatively, an existing grid can uniformly add or subtract modes. This approach only requires restriction or prolongation at the small fraction of grid points located on overset interfaces, i.e. the fringe points. The user can assign each grid a value of N a priori or elect to implement an adaptive temporal refinement strategy. This may significantly enhance applicability of the hybrid Time-Spectral method by (a) enabling the use of greater temporal resolution in regions whose solution is expanded in a sub-spectral representation (dynamically-blanked nodes) and by (b) supplying enough temporal resolution to sufficiently handle the issues addressed in §7.2; both

scenarios would otherwise require costly uniform refinement.

Spatio-Temporal Multigrid

As discussed in §5.3.2, only spatial multigrid has been implemented within the OVERFLOW Time-Spectral solver. Additional convergence acceleration may be afforded by the addition of temporal multigrid resulting in a spatio-temporal FAS scheme.

Temporal Parallelization

Mavriplis and Yang [37] exploited the parallelism inherent in the Time-Spectral approach to vastly accelerate simulation time by distributing each time sample to a different process; assuming iterative convergence rates are independent of N , Time-Spectral calculations employing different numbers of modes can be computed in the same amount of real time. Additionally, temporal parallelization could help facilitate the increased temporal resolution required to mitigate the spurious oscillations observed for under-resolved relative-motion calculations involving highly-disparate grid velocities.

Appendix A

The Navier-Stokes Equations

A.1 Cartesian Coordinates

The Navier-Stokes equations are statements of the conservation of mass, momentum and energy that can be expressed in Cartesian coordinates in strong conservation-law form.

$$\frac{\partial Q}{\partial t} + \nabla \cdot \mathbf{F}(Q) = 0 \quad (\text{A.1})$$

The state vector, $Q = \{\rho, \rho \mathbf{u}, e\}^T$, consists of the conserved quantities of mass, momentum, and energy. The flux vector, \mathbf{F} , is composed of both convective and viscous contributions, $\mathbf{F} = \mathbf{F}^c - \mathbf{F}^v$. For brevity, only two spatial dimensions are considered, but an extension to three spatial dimensions is straightforward.

$$Q = \begin{bmatrix} \rho \\ \rho u_1 \\ \rho u_2 \\ e \end{bmatrix}, \quad F_i^c = \begin{bmatrix} \rho u_i \\ \rho u_1 u_i + \delta_{i1} p \\ \rho u_2 u_i + \delta_{i2} p \\ (e + p) u_i \end{bmatrix}, \quad F_i^v = \begin{bmatrix} 0 \\ \tau_{i1} \\ \tau_{i2} \\ u_j \tau_{ij} + q_i \end{bmatrix}$$

The viscous stress tensor, $\boldsymbol{\tau}$, and heat flux, \mathbf{q} , are defined as

$$\begin{aligned}\tau_{ij} &= \mu \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) + \lambda \delta_{ij} \frac{\partial u_k}{\partial x_k} \\ q_i &= \kappa \frac{\partial T}{\partial x_i} = \frac{\gamma}{\gamma - 1} \frac{\mu}{Pr} \frac{\partial}{\partial x_i} \left(\frac{p}{\rho} \right)\end{aligned}$$

where the absolute viscosity, μ , is defined as a function of temperature, T , according to Sutherland's law. The equation of state for an ideal gas closes the system.

$$p = (\gamma - 1) \left[e - \frac{1}{2} \rho (u_i u_i) \right]$$

The ratio of specific heats, γ , is taken to be 1.4, the Prandtl number, Pr , is fixed at 0.72 and the speed of sound is defined as $a = \sqrt{\gamma p / \rho}$.

In order to solve the discretized equations on arbitrarily-oriented moving meshes, the equations must be transformed into generalized curvilinear coordinates. The strong conservation-law form of the governing equations in the physical domain

$$\frac{\partial Q}{\partial t} + \frac{\partial F}{\partial x} + \frac{\partial G}{\partial y} = 0 \quad (\text{A.2})$$

is transformed into strong conservation-law form in the generalized curvilinear coordinates where $F = F_1$ and $G = F_2$ are the flux vectors corresponding to the Cartesian x - and y -directions, respectively.

A.2 Generalized Curvilinear Coordinates

The current section follows the approach developed in [111] and [112]. The Navier-Stokes equations can be mapped from strong conservation-law form in the physical domain, (t, x, y) ,

$$\frac{\partial Q}{\partial t} + \frac{\partial F}{\partial x} + \frac{\partial G}{\partial y} = 0 \quad (\text{A.3})$$

into strong conservation-law form in computational domain, (τ, ξ, η) , using the following generalized curvilinear coordinates.

$$\tau = t, \quad \xi = \xi(t, x, y), \quad \eta = \eta(t, x, y)$$

The following differential relationships hold.

$$\frac{\partial}{\partial t} = \frac{\partial}{\partial \tau} + \frac{\partial \xi}{\partial t} \frac{\partial}{\partial \xi} + \frac{\partial \eta}{\partial t} \frac{\partial}{\partial \eta} \quad (\text{A.4})$$

$$\frac{\partial}{\partial x} = \frac{\partial \xi}{\partial x} \frac{\partial}{\partial \xi} + \frac{\partial \eta}{\partial x} \frac{\partial}{\partial \eta} \quad (\text{A.5})$$

$$\frac{\partial}{\partial y} = \frac{\partial \xi}{\partial y} \frac{\partial}{\partial \xi} + \frac{\partial \eta}{\partial y} \frac{\partial}{\partial \eta} \quad (\text{A.6})$$

Their substitution into Eq. A.3 results in

$$\frac{\partial Q}{\partial \tau} + \frac{\partial Q}{\partial \xi} \frac{\partial \xi}{\partial \tau} + \frac{\partial Q}{\partial \eta} \frac{\partial \eta}{\partial \tau} + \frac{\partial F}{\partial \xi} \frac{\partial \xi}{\partial x} + \frac{\partial F}{\partial \eta} \frac{\partial \eta}{\partial x} + \frac{\partial G}{\partial \xi} \frac{\partial \xi}{\partial y} + \frac{\partial G}{\partial \eta} \frac{\partial \eta}{\partial y} = 0 \quad (\text{A.7})$$

Equations A.4-A.6 can be expressed in matrix form, where T is a differential transformation operator from the computational domain into the physical domain.

$$\begin{bmatrix} \frac{\partial}{\partial t} \\ \frac{\partial}{\partial x} \\ \frac{\partial}{\partial y} \end{bmatrix} = \overbrace{\begin{bmatrix} 1 & \frac{\partial \xi}{\partial t} & \frac{\partial \eta}{\partial t} \\ 0 & \frac{\partial \xi}{\partial x} & \frac{\partial \eta}{\partial x} \\ 0 & \frac{\partial \xi}{\partial y} & \frac{\partial \eta}{\partial y} \end{bmatrix}}^T \begin{bmatrix} \frac{\partial}{\partial \tau} \\ \frac{\partial}{\partial \xi} \\ \frac{\partial}{\partial \eta} \end{bmatrix}$$

Similarly, T^{-1} can be expressed as the inverse transformation.

$$\begin{bmatrix} \frac{\partial}{\partial \tau} \\ \frac{\partial}{\partial \xi} \\ \frac{\partial}{\partial \eta} \end{bmatrix} = \overbrace{\begin{bmatrix} 1 & \frac{\partial x}{\partial \tau} & \frac{\partial y}{\partial \tau} \\ 0 & \frac{\partial x}{\partial \xi} & \frac{\partial y}{\partial \xi} \\ 0 & \frac{\partial x}{\partial \eta} & \frac{\partial y}{\partial \eta} \end{bmatrix}}^{T^{-1}} \begin{bmatrix} \frac{\partial}{\partial t} \\ \frac{\partial}{\partial x} \\ \frac{\partial}{\partial y} \end{bmatrix}$$

The mappings between the two coordinate systems are displayed in Fig. A.1. The following metric identities are found by relating corresponding entries between T and

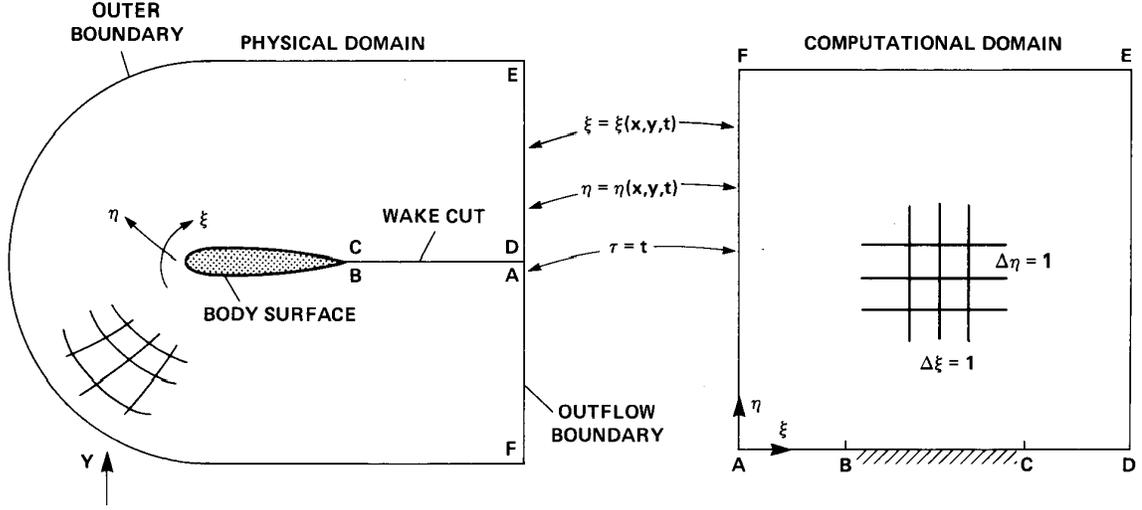


Figure A.1: Generalized curvilinear coordinate transformations. Figure taken from Pulliam [111]

the inverse of its inverse, $(T^{-1})^{-1}$.

$$\begin{aligned} \frac{\partial \xi}{\partial t} &= J \left(\frac{\partial y}{\partial \tau} \frac{\partial x}{\partial \eta} - \frac{\partial x}{\partial \tau} \frac{\partial y}{\partial \eta} \right) & \frac{\partial \eta}{\partial t} &= J \left(\frac{\partial x}{\partial \tau} \frac{\partial y}{\partial \xi} - \frac{\partial y}{\partial \tau} \frac{\partial x}{\partial \xi} \right) \\ \frac{\partial \xi}{\partial x} &= J \frac{\partial y}{\partial \eta} & \frac{\partial \eta}{\partial x} &= -J \frac{\partial y}{\partial \xi} \\ \frac{\partial \xi}{\partial y} &= -J \frac{\partial x}{\partial \eta} & \frac{\partial \eta}{\partial y} &= J \frac{\partial x}{\partial \xi} \end{aligned}$$

The Jacobian, J , is defined as the determinant of T , which is also inverse of the volume, V .

$$J = \frac{\partial \xi}{\partial x} \frac{\partial \eta}{\partial y} - \frac{\partial \eta}{\partial x} \frac{\partial \xi}{\partial y} = \left(\frac{\partial x}{\partial \xi} \frac{\partial y}{\partial \eta} - \frac{\partial y}{\partial \xi} \frac{\partial x}{\partial \eta} \right)^{-1} = V^{-1} = |T| \quad (\text{A.8})$$

This result can be substituted into Eq. A.7 to arrive at the curvilinear form of the Navier-Stokes equations in differential conservation-law form by defining the following

transformed solution and flux vectors, \mathcal{Q} , \mathcal{F} and \mathcal{G} .

$$\mathcal{Q} = \frac{Q}{J} \quad (\text{A.9})$$

$$\mathcal{F} = \frac{1}{J} \left(Q \frac{\partial \xi}{\partial t} + F \frac{\partial \xi}{\partial x} + G \frac{\partial \xi}{\partial y} \right) \quad (\text{A.10})$$

$$\mathcal{G} = \frac{1}{J} \left(Q \frac{\partial \eta}{\partial t} + F \frac{\partial \eta}{\partial x} + G \frac{\partial \eta}{\partial y} \right) \quad (\text{A.11})$$

Metric terms can be substituted for evaluation in the computational domain.

$$\mathcal{F} = Q \left(\frac{\partial y}{\partial \tau} \frac{\partial x}{\partial \eta} - \frac{\partial x}{\partial \tau} \frac{\partial y}{\partial \eta} \right) + F \frac{\partial y}{\partial \eta} - G \frac{\partial x}{\partial \eta} \quad (\text{A.12})$$

$$\mathcal{G} = Q \left(\frac{\partial x}{\partial \tau} \frac{\partial y}{\partial \xi} - \frac{\partial y}{\partial \tau} \frac{\partial x}{\partial \xi} \right) - F \frac{\partial y}{\partial \xi} + G \frac{\partial x}{\partial \xi} \quad (\text{A.13})$$

Algebraic manipulation and application of differential identities such as $\frac{\partial^2 x}{\partial \eta \partial \xi} = \frac{\partial^2 x}{\partial \xi \partial \eta}$ results in the final expression where the conservation-law form of the governing equations is mapped from the physical domain (t, x, y) into the computational domain (τ, ξ, η) .

$$\frac{\partial \mathcal{Q}}{\partial t} + \frac{\partial \mathcal{F}}{\partial \xi} + \frac{\partial \mathcal{G}}{\partial \eta} = 0 \quad (\text{A.14})$$

However, OVERFLOW stores the vector of conserved quantities in the physical domain (not normalized by J) and thus Eq. A.14 is expressed as follows.

$$\frac{\partial}{\partial t} \left(\frac{Q}{J} \right) + \frac{\partial \mathcal{F}}{\partial \xi} + \frac{\partial \mathcal{G}}{\partial \eta} = 0 \quad (\text{A.15})$$

The transformed solution, \mathcal{Q} , and flux vectors, \mathcal{F} and \mathcal{G} , in the ξ - and η -directions, respectively, are defined as follows, where the flux terms multiplied by the solution vector, Q , are a result of mesh motion.

$$\begin{bmatrix} \mathcal{Q} \\ \mathcal{F} \\ \mathcal{G} \end{bmatrix} = \begin{bmatrix} V & 0 & 0 \\ \frac{\partial y}{\partial \tau} \frac{\partial x}{\partial \eta} - \frac{\partial x}{\partial \tau} \frac{\partial y}{\partial \eta} & \frac{\partial y}{\partial \eta} & -\frac{\partial x}{\partial \eta} \\ \frac{\partial x}{\partial \tau} \frac{\partial y}{\partial \xi} - \frac{\partial y}{\partial \tau} \frac{\partial x}{\partial \xi} & -\frac{\partial y}{\partial \xi} & \frac{\partial x}{\partial \xi} \end{bmatrix} \begin{bmatrix} Q \\ F \\ G \end{bmatrix}$$

Appendix B

Linear Stability Analysis

Iterative stability is critical for all numerical schemes to ensure a contractive solution that will not grow unbounded. Analysis of a particular numerical scheme reveals its degree of stability by examining the eigenspectrum of the discrete evolution operator. This chapter examines stability for both the standard and hybrid Time-Spectral methods. OVERFLOW employs an approximately-factored scheme that applies a series of left-hand side inversions in each spatial direction, reducing the bandwidth of the implicit operators and the dimension of the linear systems involved. The Time-Spectral augmentation of OVERFLOW employs the same directional approach, where time is treated like an additional spatial dimension. Therefore, a stability analysis is performed to uncover the stability characteristics of the approximately-factored Time-Spectral scheme and determine if additional stabilization techniques are required.

While approximate factorization enables more rapid solution of the implicit system of equations, it is prone to instabilities in more than two dimensions. Approximate factorization is unconditionally stable in two dimensions, but is unconditionally unstable in three or more dimensions for the case of periodic central difference operators; adding artificial dissipation to the spatial operators stabilizes three-dimensional calculations. This proves critical in the analysis of the approximately-factored Time-Spectral scheme applied within OVERFLOW, as instabilities are to be expected for systems involving two- and three-spatial dimensions because of the additional temporal dimension. Custer [41] and Thomas et al. [42, 43] demonstrated unconditional

stability for both the unfactored and approximately-factored form of the equations. However, this stability analysis was performed in only two dimensions - one spatial and one temporal. A more-detailed linear stability analysis should uncover whether application of artificial dissipation in the N_{sd} spatial dimensions alone is sufficient to either unconditionally or conditionally stabilize the $(N_{sd} + 1)$ -factor approximate-factorization scheme, or whether smoothing in the temporal dimension, in the form of spectral vanishing viscosity [50, 51, 52, 113, 55] or filtering is required. Haphazard application of artificial dissipation in the temporal dimension can greatly degrade the spectral accuracy of the temporal derivative, so care must be taken if dissipation is to be employed.

B.1 Standard Time-Spectral Factored Scheme

The linear advection equation in strong conservation-law form serves as the model problem for the stability analysis of the approximately-factored Time-Spectral scheme.

$$\frac{\partial u}{\partial t} + \nabla \cdot (\mathbf{a}u) = 0 \quad (\text{B.1})$$

The vector $\mathbf{a} = \{a_x, a_y, a_z\}^T$ contains the advective speeds in each coordinate direction. Equation B.1 is discretized implicitly using the backward Euler time advance scheme in pseudotime with periodic central-difference operators (with purely imaginary eigenvalues) in the spatial dimensions and the spectrally accurate Fourier differentiation operator in the temporal dimension.

$$\frac{u^{s+1} - u^s}{\Delta\tau} + (A_t + A_x + A_y + A_z) u^{s+1} = 0 \quad (\text{B.2})$$

Stability analysis of the discrete system relies on its eigenspectrum. All circulant matrices, including the periodic central-difference and Fourier differentiation operators, share the same complete set of orthogonal eigenvectors, X [73]. Equation B.2 can be diagonalized to isolate the eigenvalues of the system, Λ , which is demonstrated by

the substitution of $u = Xw$.

$$X \frac{w^{s+1} - w^s}{\Delta\tau} + (A_t + A_x + A_y + A_z) X w^{s+1} = 0 \quad (\text{B.3})$$

Premultiplication by X^{-1} diagonalizes the system,

$$X^{-1} X \frac{w^{s+1} - w^s}{\Delta\tau} + X^{-1} (A_t + A_x + A_y + A_z) X w^{s+1} = 0 \quad (\text{B.4})$$

because $\Lambda = X^{-1}AX$.

$$\frac{w^{s+1} - w^s}{\Delta\tau} + (\Lambda_t + \Lambda_x + \Lambda_y + \Lambda_z) w^{s+1} = 0 \quad (\text{B.5})$$

Each component can therefore be analyzed independently.

$$\frac{w^{s+1} - w^s}{\Delta\tau} + (\lambda_t + \lambda_x + \lambda_y + \lambda_z) w^{s+1} = 0 \quad (\text{B.6})$$

The discretization of Eq. B.6 in *delta* form, where $\Delta w = w^{s+1} - w^s$, results in an expression for the model problem in terms of the purely imaginary eigenvalues, λ_x , λ_y and λ_z , of the spatial differencing operators A_x , A_y and A_z , respectively. The eigenvalues $\lambda_t = i\omega k$ of the Fourier temporal-differentiation operator, $A_t = \mathcal{D}_N$, are also purely imaginary.

$$[1 + h(\lambda_t + \lambda_x + \lambda_y + \lambda_z)]\Delta w = -h[\lambda_t + \lambda_x + \lambda_y + \lambda_z]w^s \quad (\text{B.7})$$

The unfactored and factored schemes are analyzed for both three- and four-dimensional cases in two and three spatial dimensions, respectively.

Three-dimensional Unfactored Scheme (x, y, t)

Starting with Eq. B.7 and considering only two spatial dimensions results in the following expression for the unfactored scheme in two spatial dimensions (the differentiation operators are of dimension $N \times N_x \times N_y$ and applied to the complete

space-time solution vector).

$$[1 + h(\lambda_t + \lambda_x + \lambda_y)]\Delta w = -h[\lambda_t + \lambda_x + \lambda_y]w^s \quad (\text{B.8})$$

The amplification factor $\sigma = w^{s+1}/w^s$ for the unfactored scheme is expressed by substituting $\Delta w = w^{s+1} - w^s$ into Eq. B.8

$$\sigma = \frac{1}{1 + h(\lambda_t + \lambda_x + \lambda_y)} \quad (\text{B.9})$$

which can be written in terms of its real and imaginary components by representing general combinations of $\Im(\lambda_x + \lambda_y + \lambda_t)$ as α .

$$\sigma = \frac{1}{1 + ih\alpha} \quad (\text{B.10})$$

Taking the magnitude of Eq. B.10 confirms that the amplification factor is less than or equal to unity for all eigenvalue combinations.

$$|\sigma|^2 = \frac{1}{1 + h^2\alpha^2} \leq 1 \quad \forall \alpha \quad (\text{B.11})$$

Despite the unconditional stability of this approach, we are interested in its factored form to avoid solving the complete space-time system and leveraging the existing OVERFLOW source which employs approximate factorization.

Three-dimensional Factored Scheme (x, y, t)

The factored form is achieved by approximately factoring the left hand side of Eq. B.8 thereby introducing a second-order factorization error of that does not degrade the accuracy of the original discretization.

$$[1 + h\lambda_t][I + h\lambda_x][I + h\lambda_y]\Delta w = -h[\lambda_t + \lambda_x + \lambda_y]w^s + \mathcal{O}(h^2) \quad (\text{B.12})$$

Again, we can analyze the stability for the scheme by examining the magnitude of the amplification factor, σ , for all combinations of the eigenvalues of the spatial

differencing operators. The magnitude of the amplification factor must be less than or equal to unity for all combinations of eigenvalues to guarantee unconditional stability. An expression for σ can be written in terms of the eigenvalues, λ .

$$\sigma = \frac{1 + h^2 (\lambda_t \lambda_x + \lambda_t \lambda_y + \lambda_x \lambda_y) + h^3 \lambda_t \lambda_x \lambda_y}{1 + h (\lambda_t + \lambda_x + \lambda_y) + h^2 (\lambda_t \lambda_x + \lambda_t \lambda_y + \lambda_x \lambda_y) + h^3 \lambda_t \lambda_x \lambda_y} \quad (\text{B.13})$$

The terms can be separated into real and imaginary parts by representing general combinations of λ_x , λ_y and λ_t as α , β and γ in order to express the magnitude of σ [73].

$$\sigma = \frac{(1 - h^2 \beta) + i(-h^3 \gamma)}{(1 - h^2 \beta) + i(h\alpha - h^3 \gamma)} \quad (\text{B.14})$$

The magnitude of Eq. B.14 can be expressed as follows.

$$|\sigma|^2 = \frac{(1 - h^2 \beta)^2 + (h^3 \gamma)^2}{(1 - h^2 \beta)^2 + (h\alpha - h^3 \gamma)^2} \not\leq 1 \quad \forall \alpha, \beta, \gamma \quad (\text{B.15})$$

This expression for the magnitude of the amplification factor is not guaranteed to be less than unity for all eigenvalue combinations. Therefore, the factored Time-Spectral scheme in two spatial dimensions is *unconditionally* unstable. By extension, the standard factored scheme in three spatial dimensions is also unconditionally unstable but can be stabilized by application of artificial dissipation [92]. We want to now investigate whether the addition of artificial dissipation in the spatial dimensions alone (retaining purely imaginary λ_t and spectral temporal accuracy) provides unconditional or conditional stability for the three-factor Time-Spectral scheme in two spatial dimensions.

The amplification factor, $\sigma \geq 1$, implies unconditional instability. This is overcome with a modest addition of fourth-difference artificial dissipation which is the default for OVERFLOW with dissipation coefficient $\epsilon_4 = 0.04$. The stability characteristics of the Time-Spectral system in two spatial dimensions mirrors that of the non Time-Spectral system in three spatial dimensions using purely convective operators, which has been shown to be unconditionally unstable. As mentioned, in the

absence of any artificial dissipation the scheme is unconditionally unstable, however applying artificial dissipation in the spatial dimensions alone results in a *conditionally* stable scheme dependent on the number of temporal degrees of freedom, N , the CFL number and the fundamental frequency of the solution, ω . Therefore, the spectral accuracy of the temporal differentiation operator can be maintained over its entire spectrum.

Figure B.1 plots the amplification spectra for a discretization of 51 spatial nodes in each of the x - and y -directions and $K = 3$ harmonics ($N = 7$). Neither scheme is stable absent artificial dissipation. While the addition of fourth-difference dissipation to the spatial operators using $\epsilon_4 = 0.04$ (standard value for OVERFLOW) stabilizes the case of $\omega = 2\pi$ (Figure B.1a), the case for $\omega = 20\pi$ (Figure B.1b) remains unstable; this result suggests conditional stability for the scheme dependent on the fundamental frequency. Figure B.2 delineates the stability regions for the cases of $N = 3$ and $N = 15$ for a discretization of 51 spatial nodes in each of the x - and y -directions. The stability region contracts with increased temporal resolution. Thus, conditional stability for the three-dimensional approximately-factored Time-Spectral scheme is dependent on N , CFL and ω .

Four-dimensional Unfactored Scheme (x, y, z, t)

Time-Spectral simulations in three spatial dimensions employ four differentiation operators. Adding a fourth term to Eq. B.8 results in the same expression of the magnitude of the amplification factor as all eigenvalues are imaginary.

$$|\sigma|^2 = \frac{1}{1 + h^2\alpha^2} \leq 1 \quad \forall \alpha \quad (\text{B.16})$$

Therefore, the four-dimensional unfactored scheme is *unconditionally* stable.

Four-dimensional Factored Scheme (x, y, z, t)

Not surprisingly, similar analysis demonstrates that the four-dimensional approximately-factored scheme is unconditionally unstable. Figures B.3 and B.4 are similar to their three-dimensional analogs in the previous section, however, the stability regions in

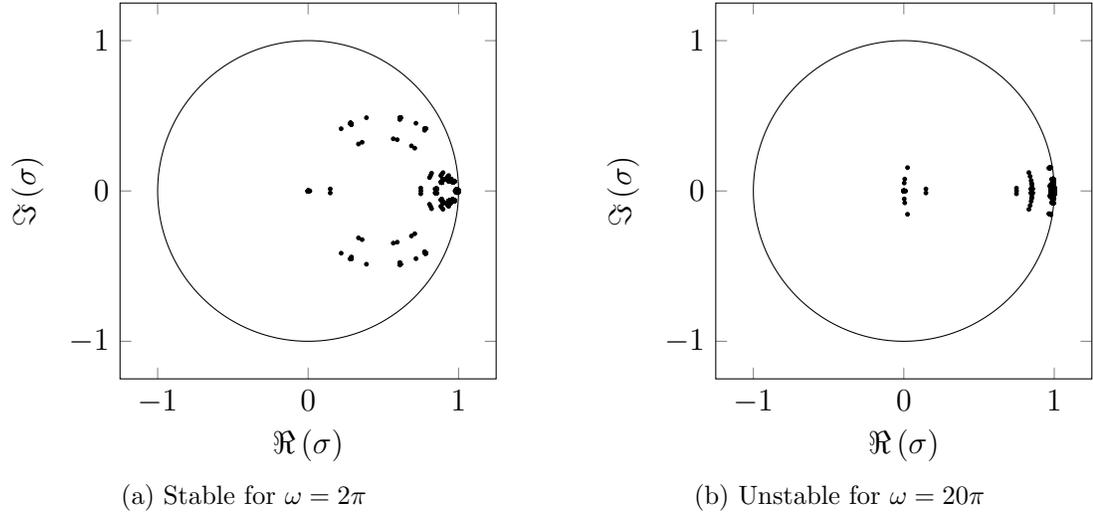


Figure B.1: Amplification spectra, σ , with $\epsilon_4 = 0.04$ for fundamental frequencies 2π and 20π and CFL = 5 with $N = 7$ and 51 spatial nodes in both the x - and y -directions.

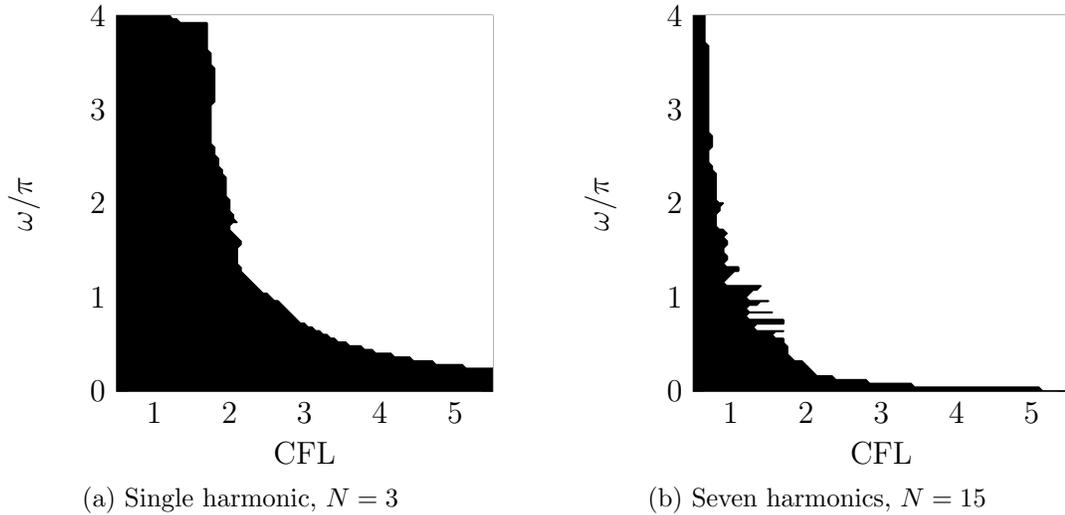


Figure B.2: Stability regions (solid black) for the approximately-factored Time-Spectral scheme in three dimensions, (x, y, t) , with 51 spatial nodes in both the x - and y -directions for $N = 3$ and 15.

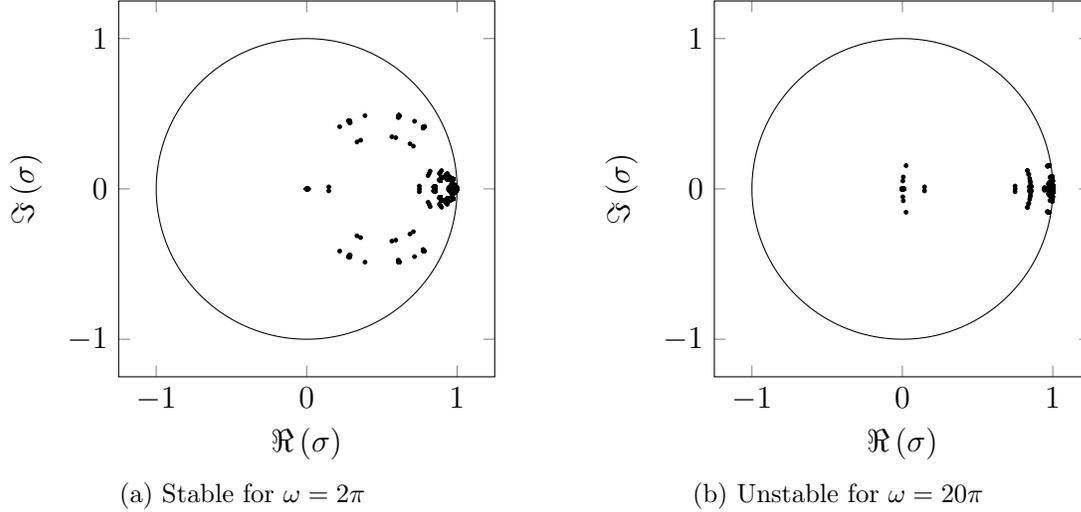


Figure B.3: Amplification spectra, σ , with $\epsilon_4 = 0.04$ for fundamental frequencies 2π and 20π and CFL = 5 with $N = 7$ and 51 spatial nodes in the x -, y - and z -directions.

Fig. B.4 are considerably more restrictive for the four-dimensional case. Nevertheless, the addition of artificial dissipation in the spatial dimensions alone, is sufficient to adequately control the conditional stability of the approximately-factored Time-Spectral scheme in both three and four dimensions.

$$|\sigma|^2 = \frac{(1 - h^2\beta + h^4\delta)^2 + (h^3\gamma)^2}{(1 - h^2\beta + h^4\delta)^2 + (h\alpha - h^3\gamma)^2} \not\leq 1 \quad \forall \alpha, \beta, \gamma, \delta \quad (\text{B.17})$$

B.2 Hybrid Time-Spectral Factored Scheme

The hybrid Time-Spectral scheme employs both the Fourier interpolation differentiation operator and an alternative differentiation operator corresponding the selected basis used to expand the solution at dynamically-blanked nodes. The various treatments for dynamically-blanked nodes introduced in Chapter 4 are compared in Chapter 6. The bounded interval approach emerged as the strongest candidate due to its superior accuracy demonstrated on the inviscid plunging NACA 0012 airfoil case in

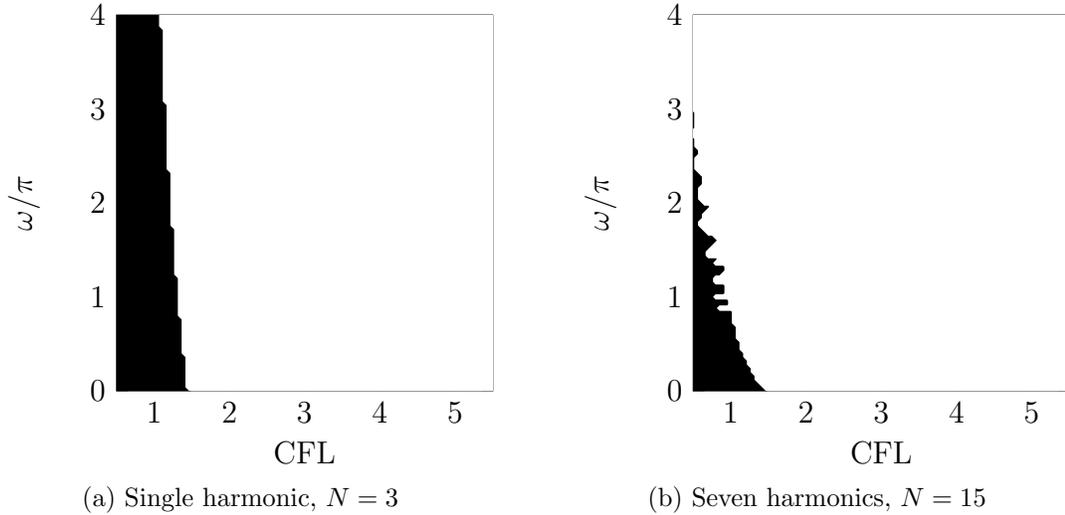


Figure B.4: Stability regions (solid black) for the approximately-factored Time-Spectral scheme in four dimensions, (x, y, z, t) , with 51 spatial nodes in the x -, y - and z -directions for $N = 3$ and 15.

§6.1.1, and will be the focus of the current section.

Linear stability analysis is typically performed on periodic boundaries, as was demonstrated in the previous section. However, this treatment would disregard the true structure of the temporal differentiation operator at dynamically-blanked nodes because it operates on bounded intervals, which are aperiodic by definition. A periodized BRI differentiation operator constructed from circulating the central entry of the aperiodic BRI operator demonstrates stability characteristics similar to the traditional Time-Spectral method, but this fact says nothing about the stability of applying the aperiodic BRI operator to bounded intervals. Finite matrix analysis can offer some insight into the stability of the hybrid BRI-Fourier Time-Spectral system. Stability would be based on a combination of the total number of degrees of freedom in time and space in addition to the number (and likely location) of dynamically-blanked nodes. For a given discretization, the global system can be constructed and approximately factored. The spectral radius of the implicit system can then be analyzed. However, infinitely many combinations are possible so further analysis is not

pursued here. However, a potential destabilizing factor to the bounded interval approach may concern the lack of physical boundary conditions. This issue was raised in Chapter 4, but Gottlieb and Orszag [10] assert that:

“[s]pectral methods are extremely sensitive to the proper formulation of boundary conditions. When proper boundary conditions are imposed, the methods yield very accurate results; when improper boundary conditions are mistakenly applied, the methods are likely to be explosively unstable” (p. 101).

While no stability proof or analysis has been offered, the fact that the approximately-factored hybrid Time-Spectral scheme has successfully and repeatedly calculated flows in two and three spatial dimensions is a strong endorsement of its stability. However, a conditional stability limit has been observed; a practical stability limit of $d_{\max} = 3-4$ has been exposed for a number of test cases.

Bibliography

- [1] K.C. Hall, J.P. Thomas, and W.S. Clark. Computation of Unsteady Nonlinear Flows in Cascades using a Harmonic Balance Technique. *9th International Symposium on Unsteady Aerodynamics, Aeroacoustics and Aeroelasticity of Turbomachines*, Lyon, France, September 2000.
- [2] Kenneth C. Hall, Jeffrey P. Thomas, and W.S. Clark. Computation of Unsteady Nonlinear Flows in Cascades Using a Harmonic Balance Technique. *AIAA Journal*, 40:879–886, May 2002.
- [3] Matthew McMullen, Antony Jameson, and Juan J. Alonso. Acceleration of Convergence to a Periodic Steady State in Turbomachinery Flows. AIAA Paper 0152, Reno, Nevada, January 2001.
- [4] M.S. McMullen and A. Jameson. The computational efficiency of non-linear frequency domain methods. *Journal of Computational Physics*, 212:637–661, 2006.
- [5] Kenneth C. Hall and Edward F. Crawley. Calculation of Unsteady Flows in Turbomachinery Using the Linearized Euler Equations. *AIAA Journal*, 27(6):777–787, June 1989.
- [6] K.C. Hall and C. B. Lorence. Calculation of Three-Dimensional Unsteady Flows in Turbomachinery Using the Linearized Harmonic Euler Equations. *Journal of Turbomachinery*, 115:800–809, October 1993.
- [7] Kenneth C. Hall and William S. Clark. Linearized Euler Predictions of Unsteady Aerodynamic Loads in Cascades. *AIAA Journal*, 31(3):540–550, March 1993.

- [8] L. He and W. Ning. Nonlinear Harmonic Analysis of Unsteady Transonic Inviscid and Viscous Flows. *Proceedings of the 8th International Symposium on Unsteady Aerodynamics and Aeroelasticity of Turbomachines*, Stockholm, Sweden, September 1998.
- [9] 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.
- [10] David Gottlieb and Steven A. Orszag. *Numerical Analysis of Spectral Methods: Theory and Applications*. Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania, 1977.
- [11] Matthew Scott McMullen. *The Application of Non-Linear Frequency Domain Methods to the Euler and Navier-Stokes Equations*. PhD thesis, Stanford University, Stanford, CA, March 2003.
- [12] Matthew McMullen, Antony Jameson, and Juan J. Alonso. Application of a Non-Linear Frequency Domain Solver to the Euler and Navier-Stokes Equations. AIAA Paper 0120, Reno, Nevada, January 2002.
- [13] Matthew McMullen, Antony Jameson, and Juan Alonso. Demonstration of Nonlinear Frequency Domain Methods. *AIAA Journal*, 44(7):1428–1435, July 2006.
- [14] Siva K. Nadarajah, Matthew S. McMullen, and Antony Jameson. Optimum Shape Design for Unsteady Flows Using Time Accurate and Non-Linear Frequency Domain Methods. AIAA Paper 3875, Orlando, Florida, June 2003.
- [15] Siva K. Nadarajah and Charles A. Tatossian. Adjoint-Based Aerodynamic Shape Optimization of Rotorcraft Blades. AIAA Paper 6730, Honolulu, Hawaii, August 2008.
- [16] Charles A. Tatossian, Siva K. Nadarajah, and Patrice Castonguay. Aerodynamic Shape Optimization of Hovering Rotor Blades Using a Non-Linear Frequency Domain Approach. AIAA Paper 0322, Reno, Nevada, January 2008.

- [17] Scott M. Murman. Reduced-Frequency Approach for Calculating Dynamic Derivatives. *AIAA Journal*, 45(6):1161–1168, June 2007.
- [18] Arathi K. Gopinath and Antony Jameson. Time Spectral Method for Periodic Unsteady Computations over Two- and Three-Dimensional Bodies. AIAA Paper 1220, Reno, Nevada, January 2005.
- [19] Arathi K. Gopinath and Antony Jameson. Application of the Time Spectral Method to Periodic Unsteady Vortex Shedding. AIAA Paper 0449, January 2006.
- [20] Edwin van der Weide, Arathi K. Gopinath, and Antony Jameson. Turbomachinery Applications with the Time Spectral Method. AIAA Paper 4905, Toronto, Ontario, June 2005.
- [21] Arathi K. Gopinath, Edwin van der Weide, Juan J. Alonso, Antony Jameson, Kivanc Ekici, and Kenneth C. Hall. Three-Dimensional Unsteady Multi-stage Turbomachinery Simulations using the Harmonic Balance Technique. AIAA Paper 0892, Reno, Nevada, January 2007.
- [22] Arathi Kamath Gopinath. *Efficient Fourier-Based Algorithms for Time-Periodic Unsteady Problems*. PhD thesis, Stanford University, Stanford, CA, April 2007.
- [23] Seongim Choi and Anubhav Datta. CFD Prediction of Rotor Loads using Time-Spectral Method and Exact Fluid-Structure Interface. AIAA Paper 7325, Honolulu, Hawaii, August 2008.
- [24] Seongim Choi, Mark Potsdam, Kihwan Lee, Gianluca Iaccarino, and Juan J. Alonso. Helicopter Rotor Design Using a Time-Spectral and Adjoint-Based Method. AIAA Paper 5810, Victoria, British Columbia, September 2008.
- [25] Florian Blanc, François-Xavier Roux, Jean-Christophe Jouhaud, and Jean-François Boussuge. Numerical Methods for Control Surfaces Aerodynamics with Flexibility Effects. IFASD 2009, Seattle, WA, June 2009.

- [26] Nathan L. Mundis and Dimitri J. Mavriplis. Quasi-periodic Time Spectral Method for Aeroelastic Flutter Analysis. AIAA Paper 0638, Grapevine, Texas, January 2013.
- [27] Nathan L. Mundis and Dimitri J. Mavriplis. An Efficient Flexible GMRES Solver for the Fully-coupled Time-Spectral Aeroelastic System. AIAA Paper 1427, National Harbor, Maryland, January 2014.
- [28] Vinod K. Lakshminarayan, Karthikeyan Duraisamy, and Juan J. Alonso. Adjoint-Based Estimation and Control of Spatial, Temporal and Stochastic Approximation Errors for Aerodynamic Applications. AIAA Paper 0518, Grapevine, Texas, January 2013.
- [29] Kedar R. Naik, Thomas D. Economon, Michael R. Colonno, Francisco Palacios, and Juan J. Alonso. A Time-Spectral Adjoint Formulation for Shape Optimization of Helicopter Rotors in Forward Flight. AHS Paper 082, Phoenix, Arizona, May 2013.
- [30] Raymond C. Maple. *Adaptive Harmonic Balance Method for Unsteady, Non-linear, One-Dimensional Periodic Flows*. PhD thesis, The Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio, September 2002.
- [31] Raymond C. Maple, Paul I King, J. Mitch Wolff, and Paul D. Orkwis. Split-Domain Harmonic Balance Solutions to Burger’s Equation for Large-Amplitude Disturbances. *AIAA Journal*, 41(2):206–212, February 2003.
- [32] Raymond C. Maple, Paul I King, and Mark E. Oxley. Adaptive Harmonic Balance Solutions to Euler’s Equation. *AIAA Journal*, 41(9):1705–1714, September 2003.
- [33] Ali Mosahebi and Siva K. Nadarajah. An Adaptive Non-linear Frequency Domain Method for Viscous Periodic Steady State Flows. AIAA Paper 1267, Orlando, Florida, January 2010.

- [34] Ali Mosahebi and Siva K. Nadarajah. An Implicit Adaptive Non-Linear Frequency Domain Method (p NLFD) for Viscous Periodic Steady State Flows on Deformable Grids. AIAA Paper 0775, Orlando, Florida, January 2011.
- [35] Ali Mosahebi and Siva K. Nadarajah. Algorithmic Advanced for the Adaptive Non-Linear Frequency Domain Method. AIAA Paper 0380, Grapevine, Texas, January 2013.
- [36] A. Mosahebi and S. Nadarajah. An adaptive Non-Linear Frequency Domain method for viscous flows. *Computers & Fluids*, 75:140–154, 2013.
- [37] Dimitri J Mavriplis and Zhi Yang. Time Spectral Method for Periodic and Quasi-Periodic Unsteady Computations on Unstructured Meshes. *Mathematical Modeling of Natural Phenomenon*, 6(3):213–236, 2011.
- [38] Zhi Yang, Dimitri J Mavriplis, and Jayanarayanan Sitaraman. Prediction of Helicopter Maneuver Loads Using BDF/Time Spectral Method on Unstructured Meshes. AIAA Paper 1122, Orlando, Florida, January 2011.
- [39] Dimitri J. Mavriplis, Zhi Yang, and Nathan Mundis. Extensions of Time Spectral Methods for Practical Rotorcraft Problems. AIAA Paper 0423, Nashville, Tennessee, January 2012.
- [40] Nathan L. Mundis and Dimitri J. Mavriplis. GMRES applied to the Time-spectral and Quasi-periodic Time-spectral Methods. AIAA Paper 3084, San Diego, California, June 2013.
- [41] Chad H. Custer. *A Nonlinear Harmonic Balance Solver for an Implicit CFD Code: OVERFLOW 2*. PhD thesis, Duke University, 2009.
- [42] Jeffrey P. Thomas, Chad H. Custer, Earl H. Dowell, and Kenneth C. Hall. Unsteady Flow Computation Using a Harmonic Balance Approach Implemented about the OVERFLOW 2 Flow Solver. AIAA Paper 4270, San Antonio, Texas, June 2009.

- [43] Jeffrey P. Thomas, Chad H. Custer, Earl H. Dowell, Kenneth C. Hall, and Christophe Corre. Compact Implementation Strategy for a Harmonic Balance Method Within Implicit Flow Solvers. *AIAA Journal*, 51(6):1374–1381, June 2013.
- [44] Olivier Soucy and Siva K. Nadarajah. A NLFD method for the simulation of periodic unsteady flows for overset meshes. AIAA Paper 4271, San Antonio, Texas, June 2009.
- [45] Arthur D. Snider. An Improved Estimate of the Accuracy of Trigonometric Interpolation. *SIAM Journal of Numerical Analysis*, 9(3):505–508, September 1972.
- [46] Roger Peyret. *Spectral Methods for Incompressible Viscous Flow*. Applied Mathematical Sciences. Springer, New York, 2002.
- [47] Jan S. Hesthaven, Sigal Gottlieb, and David Gottlieb. *Spectral Methods for Time-Dependent Problems*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, Cambridge, 2007.
- [48] C. Canuto, M.Y. Hussaini, A. Quarteroni, and T.A. Zang. *Spectral Methods: Fundamentals in Single Domains*. Scientific Computation. Springer, Berlin, 2006.
- [49] Steven A. Orszag. On the Elimination of Aliasing in Finite-Difference Schemes by Filtering High-Wavenumber Components. *Journal of the Atmospheric Sciences*, 28(6):1074, September 1971.
- [50] Eitan Tadmor. Convergence of Spectral Methods for Nonlinear Conservation Laws. *SIAM Journal of Numerical Analysis*, 26(1):30–44, February 1989.
- [51] Yvon Maday and Eitan Tadmor. Analysis of the Spectral Vanishing Viscosity Method for Periodic Conservation Laws. *SIAM Journal of Numerical Analysis*, 26(4):854–870, August 1989.

- [52] Yvon Maday, Sidi M. Ould Kaber, and Eitan Tadmor. Legendre Pseudospectral Viscosity Method for Nonlinear Conservation Laws. *SIAM Journal of Numerical Analysis*, 30(2):321–342, April 1993.
- [53] L. Liu, J.P. Thomas, E.H. Dowell, P. Attar, and K.C. Hall. A comparison of classical and high dimensional harmonic balance approaches for a Duffing oscillator. *Journal of Computational Physics*, 215:298–320, 2006.
- [54] A. LaBryer and P.J. Attar. High dimensional harmonic balance dealiasing techniques for a Duffing oscillator. *Journal of Sound and Vibration*, 324:1016–1038, 2009.
- [55] Huang Huang and Kivanc Ekici. Stabilization of High-Dimensional Harmonic Balance Solvers Using a Temporal Spectral Viscosity Operator. AIAA Paper 0354, Grapevine, Texas, January 2013.
- [56] T. H. Pulliam and J. L. Steger. On Implicit Finite-Difference Simulations of Three Dimensional Flow. AIAA Paper 10, Huntsville, Alabama, January 1978.
- [57] Christopher L. Rumsey, Robert T. Biedron, and James L. Thomas. *CFL3D, Its History and Some Recent Applications*. National Aeronautics and Space Administration, Langley Research Center, Hampton, Virginia, 1997.
- [58] J. A. Benek, J. L. Steger, and F. C. Dougherty. A Flexible Grid Embedding Technique with Application to the Euler Equations. AIAA-83-1944-CP, 1983.
- [59] J. A. Benek, P. G. Buning, and J. L. Steger. A 3-D Chimera Grid Embedding Technique. AIAA-85-1523-CP, 1985.
- [60] Joseph L. Steger and John A. Benek. On the Use of Composite Grid Schemes in Computational Aerodynamics. *Computer Methods in Applied Mechanics and Engineering*, 64:301–320, 1987.
- [61] Robert L. Meakin. A New Method for Establishing Intergrid Communication Among Systems of Overset Grids. AIAA-91-1586-CP, 1991.

- [62] Z.J. Wang, P. Buning, and J. Benek. Critical Evaluation of Conservative and Non-Conservative Interface Treatment for Chimera Grids. AIAA Paper 0077, Reno, Nevada, January 1995.
- [63] Robert L. Meakin. Object X-Rays for Cutting Holes in Composite Overset Structured Grids. AIAA Paper 2537, Anaheim, CA, June 2001.
- [64] William M. Chan. Overset grid technology development at NASA Ames Research Center. *Computers & Fluids*, 38:496–503, 2009.
- [65] Emmanuel Candès. Compressive sampling. Madrid, Spain, 2006. International Congress of Mathematicians.
- [66] Daan Huybrechs. On the Fourier Extension of Nonperiodic Functions. *SIAM Journal of Numerical Analysis*, 47(6):4326–4355, 2010.
- [67] John P. Boyd. A Comparison of Numerical Algorithms for Fourier Extension of the First, Second, and Third Kinds. *Journal of Computational Physics*, 178:118–160, 2002.
- [68] Oscar P. Bruno, Youngae Han, and Matthew M. Pohlman. Accurate, high-order representation of complex three-dimensional surfaces via Fourier continuation analysis. *Journal of Computational Physics*, 227:1094–1125, 2007.
- [69] Oscar P. Bruno and Mark Lyon. High-order unconditionally stable FC-AD solvers for general smooth domains I. Basic Elements. *Journal of Computational Physics*, 229:2009–2033, 2010.
- [70] Mark Lyon and Oscar P. Bruno. High-order unconditionally stable FC-AD solvers for general smooth domains II. Elliptic, parabolic and hyperbolic PDEs; theoretical considerations. *Journal of Computational Physics*, 229:3358–3381, 2010.
- [71] Nathan Albin and Oscar P. Bruno. A spectral FC solver for the compressible Navier-Stokes equations in general domains I: Explicit time-stepping. *Journal of Computational Physics*, 230:6248–6270, 2011.

- [72] John P. Boyd. *Chebyshev and Fourier Spectral Methods*. Dover, Mineola, New York, Second (revised) edition, 2001.
- [73] Harvard Lomax, Thomas H. Pulliam, and David W. Zingg. *Fundamentals of Computational Fluid Dynamics*. Scientific Computation. Springer, Berlin, 2001.
- [74] Parviz Moin. *Fundamentals of Engineering Numerical Analysis*. Cambridge University Press, New York, 2001.
- [75] Ingrid Daubechies. Orthonormal Bases of Compactly Supported Wavelets. *Communications on Pure and Applied Mathematics*, XLI:909–996, 1988.
- [76] Ingrid Daubechies. Orthonormal Bases of Compactly Supported Wavelets II. Variations on a Theme. *SIAM Journal of Mathematical Analysis*, 24(2):499–519, March 1993.
- [77] J. Ko, A.J. Kurdila, and M.S. Pilant. A class of finite element methods based on orthonormal, compactly-supported wavelets. *Computational Mechanics*, 16:235–244, 1995.
- [78] C. H. Romine and B. W. Peyton. Computing Connection Coefficients of Compactly Supported Wavelets on Bounded Intervals. Mathematical Sciences Section ORNL/TM-13413, Oak Ridge National Laboratory, April 1997.
- [79] Alfredo Eisinberg and Giuseppe Fedele. Discrete Orthogonal Polynomials on Equidistant Nodes. *International Mathematical Forum*, 2(21):1007–1020, 2007.
- [80] Rodrigo B. Platte, Lloyd N. Trefethen, and Arno B.J. Kuijlaars. Impossibility of Fast Stable Approximation of Analytic Functions from Equispaced Samples. *SIAM Review*, 53(2):308–318, 2011.
- [81] Len Bos, Stefano De Marchi, Kai Hormann, and Georges Klein. On the Lebesgue constant of barycentric rational interpolation at equidistant nodes. *Numerische Mathematik*, 121(3):461–471, 2012.

- [82] Richard Baltensperger and Jean-Paul Berrut. The linear rational collocation method. *Journal of Computational and Applied Mathematics*, 134:243–258, 2001.
- [83] Jean-Paul Berrut and Richard Baltensperger. The Linear Rational Pseudospectral Method for Boundary Value Problems. *BIT*, 41(5):868–879, 2001.
- [84] Georges Klein. An Extension of the Floater-Hormann Family of Barycentric Rational Interpolants. *Mathematics of Computation*, 82(284):2273–2292, October 2013.
- [85] T. W. Tee and Lloyd N. Trefethen. A Rational Spectral Collocation Method with Adaptively Transformed Chebyshev Grid Points. *SIAM Journal of Scientific Computing*, 28(5):1798–1811, 2006.
- [86] Richard Baltensperger, Jean-Paul Berrut, and Benjamin Noël. Exponential Convergence of a Linear Rational Interpolant Between Transformed Chebyshev Points. *Mathematics of Computation*, 68(227):1109–1120, February 1999.
- [87] Michael S. Floater and Kai Hormann. Barycentric rational interpolation with no poles and high rates of approximation. *Numerische Mathematik*, 107(2):315–331, August 2007.
- [88] Qiqi Wang, Parviz Moin, and Gianluca Iaccarino. A Rational Interpolation Scheme with Superpolynomial Rate of Convergence. *SIAM Journal of Numerical Analysis*, 47(6):4073–4097, 2010.
- [89] R.H. Nichols, R.W. Tramel, and P.G. Buning. Solver and Turbulence Model Upgrades to OVERFLOW 2 for Unsteady and High-Speed Applications. AIAA Paper 2824, San Francisco, CA, June 2006.
- [90] Richard 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.

- [91] A. Jameson. Time Dependent Calculations Using Multigrid, with Applications to Unsteady Flows Past Airfoils and Wings. AIAA Paper 1596, Honolulu, Hawaii, June 1991.
- [92] T. H. Pulliam and D. S. Chaussee. A Diagonal Form of an Implicit Approximate-Factorization Algorithm. *Journal of Computational Physics*, 39:347–363, 1981.
- [93] A. Jameson, W. Schmidt, and E. Turkel. Numerical Solution of the Euler Equations by Finite Volume Methods Using Runge-Kutta Time-Stepping Schemes. AIAA-81-1259, Palo Alto, CA, June 1981.
- [94] Thomas H. Pulliam. Artificial Dissipation Models for the Euler Equations. *AIAA Journal*, 24(12):1931–1940, December 1986.
- [95] Thomas H. Pulliam and David W. Zingg. *Fundamental Algorithms in Computational Fluid Dynamics*. Scientific Computation. Springer, 2014.
- [96] P. Wesseling. *An Introduction to Multigrid Methods*. John Wiley & Sons, Chichester, 1992.
- [97] Dennis Jespersen, Thomas Pulliam, and Pieter Buning. Recent Enhancements to OVERFLOW (Navier-Stokes code). AIAA Paper 0644, Reno, Nevada, January 1997.
- [98] P. R. Spalart and S. R. Allmaras. A One-Equation Turbulence Model for Aerodynamic Flows. AIAA Paper 0439, Reno, NV, January 1992.
- [99] Joshua I. Leffell, Scott M. Murman, and Thomas H. Pulliam. An Extension of the Time-Spectral Method to Overset Solvers. AIAA Paper 0637, Grapevine, Texas, January 2013.
- [100] K. D. Jones, C. M. Dohring, and M. F. Platzer. Experimental and Computational Investigation of the Knoller-Betz Effect. *AIAA Journal*, 36(7):1240–1246, July 1998.

- [101] Yves Allaneau and Antony Jameson. Direct Numerical Simulations of Plunging Airfoils. AIAA Paper 0728, Orlando, Florida, January 2010.
- [102] R. H. Landon. Compendium of Unsteady Aerodynamic Measurements. AGARD Report No. 702, August 1982.
- [103] Neal M. Chaderjian. Advances in Rotor Performance and Turbulent Wake Simulation using DES and Adaptive Mesh Refinement. Seventh International Conference of Computational Fluid Dynamics ICCFD7-3506, Big Island, Hawaii, July 2012.
- [104] Mark A. Potsdam and Roger C. Strawn. CFD Simulations of Tiltrotor Configurations in Hover. *Journal of the American Helicopter Society*, 50(1):82–94, 2005. Also presented at the American Helicopter Society 58th Annual Forum, Montreal, Canada, June 11-13, 2002.
- [105] Kivanc Ekici, Kenneth C. Hall, and Earl H. Dowell. Computationally Fast Harmonic Balance Methods for Unsteady Aerodynamic Predictions of Helicopter Rotors. AIAA Paper 1439, Reno, Nevada, January 2008.
- [106] F.X. Caradonna and C. Tung. *Experimental and Analytical Studies of a Model Helicopter Rotor in Hover*. NASA Technical Manual 81232, Ames Research Center, Moffett Field, CA, September 1981.
- [107] Larry A. Young, Jr. Earl R. Booth, Gloria K. Yamauchi, Gavin Botha, and Seth Dawson. Overview of the Testing of a Small-Scale Proprotor. American Helicopter Society 55th Annual Forum, Montreal, Canada, May 1999.
- [108] Stephen M. Swanson, Megan S. McCluer, Gloria K. Yamauchi, and Alexandra A. Swanson. Airloads Measurements From a 1/4-Scale Tiltrotor Wind Tunnel Test. 25th European Rotorcraft Forum, Rome, Italy, September 1999.
- [109] Mark Potsdam and Thomas Pulliam. Turbulence Modeling Treatment for Rotorcraft Wakes. Presented at the AHS Specialist’s Conference on Aeromechanics, San Francisco, CA, January 2008.

- [110] Neal M. Chaderjian and Pieter G. Buning. High Resolution Navier-Stokes Simulations of Rotor Wakes. Presented at the American Helicopter Society 67th Annual Forum, Virginia Beach, VA, May 2011.
- [111] Thomas H. Pulliam. *Efficient Solution Methods for The Navier-Stokes Equations*. Lecture Notes for the von Kármán Institute For Fluid Dynamics Lecture Series: Numerical Techniques for Viscous Flow Computation In Turbomachinery Bladings, von Kármán Institute, Rhode-St-Genese, Belgium, 1985.
- [112] Robert MacCormack. Numerical Methods for Compressible Flows. Class notes provided as part of the AA214 series at Stanford University, 2009.
- [113] Anne Gelb and Eitan Tadmor. Enhanced spectral viscosity approximations for conservation laws. *Applied Numerical Mathematics*, 33:3–21, 2000.