



ADAPTIVE HARMONIC BALANCE METHOD
FOR UNSTEADY, NONLINEAR,
ONE-DIMENSIONAL PERIODIC FLOWS

DISSERTATION

Raymond C. Maple
Lieutenant Colonel, USAF

AFIT/DS/ENY/02-2

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE (DD-MM-YYYY) September 2002		2. REPORT TYPE Doctoral Dissertation		3. DATES COVERED (From - To) Aug 1999 - Sep 2002	
4. TITLE AND SUBTITLE ADAPTIVE HARMONIC BALANCE METHOD FOR UNSTEADY, NONLINEAR, ONE-DIMENSIONAL PERIODIC FLOWS				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Maple, Raymond, C., Lieutenant Colonel, USAF				5d. PROJECT NUMBER ENR#1999-104	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 P Street, Bldg 640 WPAFB OH 45433-7765				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/DS/ENY/02-2	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Elizabeth Downie Director, Dayton Area Graduate Studies Institute 3155 Research Blvd., Ste. 205 Kettering, Ohio 45420 Phone: (937) 781-4000 FAX: (937) 781-4005				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT A new adaptive split-domain harmonic balance computational fluid dynamics (CFD) method is developed to solve highly nonlinear time-periodic flows such as those found in turbomachinery. The basic harmonic balance CFD method transforms an unsteady time-periodic problem into a steady-state problem by assuming a solution in the form of a Fourier series in time. The new method employs a unique multi-domain split-operator solution technique to remove a large-series stability restriction present in previous harmonic balance CFD approaches. In addition, the new method adapts the frequency content to the flow, starting with a small number of Fourier frequencies and augmenting the frequency content in each cell as necessary to capture local flow physics. The method reduces compute times by allowing larger integration time steps, eliminating Fourier transforms, and reducing overall problem size. The stability and accuracy of the method are verified with solutions to the 1-D inviscid Burger's equation and 1-D Euler's equation. Accurate adapted quasi-1-D Euler solutions for a supersonic/subsonic diverging nozzle with periodic unsteady outflow conditions are generated in 86% less time than an equivalent non-adapted split-domain solution, demonstrating the performance benefit of matching frequency content to the local flow conditions.					
15. SUBJECT TERMS HARMONIC BALANCE, COMPUTATIONAL FLUID DYNAMICS, UNSTEADY FLOW, NUMERICAL METHODS AND PROCEDURES, FOURIER SERIES					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 167	19a. NAME OF RESPONSIBLE PERSON
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			Paul I. King (ENY)
					19b. TELEPHONE NUMBER (Include area code) (937) 255 3636 ext. 4628; paul.king@afit.edu

The views expressed in this dissertation are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the United States Government.

AFIT/DS/ENY/02-2

ADAPTIVE HARMONIC BALANCE METHOD
FOR UNSTEADY, NONLINEAR,
ONE-DIMENSIONAL PERIODIC FLOWS

DISSERTATION

Presented to the Faculty of the Graduate School of Engineering and Management
of the Air Force Institute of Technology
Air University
In Partial Fulfillment of the
Requirements for the Degree of
Doctor of Philosophy

Raymond C. Maple, B.S.M.E, M.S.A.E
Lieutenant Colonel, USAF

September, 2002

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

ADAPTIVE HARMONIC BALANCE METHOD
FOR UNSTEADY, NONLINEAR,
ONE-DIMENSIONAL PERIODIC FLOWS

DISSERTATION

Raymond C. Maple, B.S.M.E, M.S.A.E

Lieutenant Colonel, USAF

Approved:

/signed/

Paul I. King
Dissertation Advisor

Date

/signed/

Montgomery Hughson
Committee Member

Date

/signed/

Mark Oxley
Committee Member

Date

/signed/

J. Mitch Wolff
Committee Member

Date

/signed/

Paul D. Orkwis
Committee member

Date

/signed/

Yung Kee Yeo
Dean's Representative

Date

/signed/

Robert A. Calico, Jr
Dean

Acknowledgments

No accomplishment of any significance is achieved solely by the efforts of one person, and that is certainly true of this work. I am deeply indebted to Dr. Paul King, my advisor, who has been a true mentor and guide, and to the members of my research committee, Dr. Mitch Wolff, Dr. Paul Orkwis, Dr. Mark Oxley, and Lt Col Montgomery Hughson, for the assistance they have provided. I am also indebted to Dr. Frank Moore for partially funding this effort through Dayton Area Graduate Studies Institute grant PR-AFIT-99-07.

I would like to thank my fellow students, Major John Anttonen, Major David Lucia, Major Tim Parker, and Captain Patrick Wade for taking the time to listen when I needed a sounding board, and for showing me that there is humor in any situation.

Finally I must acknowledge my wife and daughters, for without their love, patience, encouragement, and support, this work would never have been completed. I am especially indebted to my wife for many hours spent proofreading this document.

To all, a sincere thank you.

Raymond C. Maple

Table of Contents

	Page
Acknowledgments	iii
List of Figures	vii
List of Tables	xiii
Nomenclature	xiv
Abbreviations	xx
Abstract	xxi
 I. Introduction	 1-1
1.1 Overview of Harmonic Balance Method	1-3
1.2 Scope	1-6
1.2.1 Thesis Statement	1-6
1.2.2 Assumptions	1-7
1.3 Research Approach	1-7
1.4 Document Organization	1-8
 II. Theory and Implementation of Adaptive Split-Domain Harmonic Balance	 2-1
2.1 Introduction	2-1
2.2 Derivation of the Harmonic Balance Equations	2-1
2.3 Prior Implementations of Harmonic Balance for CFD	2-7
2.4 Split-Domain Harmonic Balance	2-8
2.5 Adaptive Split-Domain Harmonic Balance	2-12
2.5.1 Frequency Augmentation Criteria and Scheduling	2-14

	Page
2.5.2 Treatment of Frequency Transitions	2-18
2.6 Multigrid and Adaptive Split-Domain Harmonic Balance	2-22
III. Application of Split-Domain Harmonic Balance to Burgers' Equation	3-1
3.1 Introduction	3-1
3.2 Solver Implementation	3-1
3.3 Test Configuration	3-4
3.4 Results	3-7
3.4.1 Accuracy	3-7
3.4.2 Stability	3-14
3.4.3 Performance	3-16
3.5 Summary	3-16
IV. Application of Adaptive Split-Domain Harmonic Balance to Euler's Equation	4-1
4.1 Introduction	4-1
4.2 Solver Implementation	4-1
4.3 Test Configurations	4-7
4.4 Tuning the Adaptive Solver	4-12
4.4.1 Selection of Augmentation Threshold	4-12
4.4.2 Fringe Width	4-17
4.4.3 Pixelation Width	4-18
4.4.4 Adaptation Trigger	4-18
4.5 Results	4-19
4.5.1 Accuracy	4-19
4.5.2 Grid Density	4-22
4.5.3 Performance	4-24
4.6 Multigrid	4-25

	Page
4.7 Summary	4-27
V. Application of Adaptive Split-Domain Harmonic Balance to an Unsteady Diverging Nozzle	5-1
5.1 Introduction	5-1
5.2 Solver Implementation	5-2
5.3 Test Configuration	5-3
5.4 Results	5-6
5.4.1 Comparison with Exact Solutions	5-6
5.4.2 Unsteady Nozzle Solutions	5-8
5.5 Summary	5-19
VI. Summary and Conclusions	6-1
6.1 Summary of Research	6-1
6.2 Conclusions	6-3
VII. Recommendations for Future Research	7-1
References	REF-1
Appendix A. Derivation of Harmonic Balance Burgers' Equation for a Complex Fourier Series	A-1
Appendix B. Impact of Operator Splitting Error in the Split-Domain Harmonic Balance Solution	B-1
Appendix C. Determination of Optimal Frequency Augmentation In- crements	C-1
Appendix D. Implementation of Boundary Conditions for the 1-D and Quasi-1-D Euler Equation	D-1
Vita	VITA-1

List of Figures

Figure		Page
1.1.	Experimental Pressure Data on Inlet Guide Vane Upstream of Transonic Rotating Compressor Blade Row for 5% Span (top), and 95% Span (bottom)	1-3
2.1.	Time-sample Shifting that Occurs with Each Iteration of the Split-domain Solution Process	2-13
2.2.	Illustration of Typical Spatial Variation of a Low-frequency (top) and High-frequency (bottom) Fourier Coefficient	2-14
2.3.	Effect of Fringe Augmentation and Pixelation on a Threshold-based Frequency Distribution	2-16
2.4.	Flowchart for the Adaptive Harmonic Balance Algorithm . .	2-19
2.5.	Sample-rate Interpolation Required at a Frequency Transition from $N = 2$ to $N = 3$	2-20
2.6.	Non-physical Interpolated Values Produced by Frequency-domain Zero-padding	2-21
2.7.	Comparison of Fourier Coefficient Magnitudes at a Transition from 7 Frequencies to 16 Frequencies, for Solutions Computed with Time-domain Linear Interpolation (left) and Frequency-domain Truncation/Zero-padding (right) Resampling Methods	2-22
2.8.	Illustration of a 4-Level Multigrid V Cycle	2-25
2.9.	Illustration of How a Pixelation Width of 2^{N-1} Preserves Spatial Frequency Content and Eliminates Frequency Transitions During Restriction	2-27
3.1.	Sinusoidal (left) and Wake Function (right) Boundary Condition Waveforms	3-5

Figure		Page
3.2.	Comparison of Time-accurate (Solid Line) and 48-Frequency Harmonic Balance (Symbols) Solutions for the Sine Input for $t = 0$. Inset plots Show Dependent Variable Magnitude vs. Nondimensional Distance	3-8
3.3.	Comparison of Time-accurate (Solid line) and 48-Frequency Harmonic Balance(Symbols) Solutions for the Wake Function Input for $t = 0$. Inset Plots Show Dependent Variable Magnitude vs. Nondimensional Distance	3-9
3.4.	RMS Error Between Time-accurate and Harmonic Balance Solutions—Sine Input	3-10
3.5.	RMS Error Between Time-accurate and Harmonic Balance Solutions—Wake Function Input	3-10
3.6.	Comparison of 48-Frequency and 97-Frequency Solutions for $a = 0.5$, $f = 0.75$, on a 501 Point Grid	3-11
3.7.	Smoothing Effect of Coarse Grid on Harmonic Balance Solution	3-13
3.8.	Severe Damping Effect of Coarse Grid on Harmonic Balance Solution	3-14
3.9.	Comparison of Spatial Variation of High-frequency Coefficient on (a) 501 Point Grid and (b) 101 Point Grid	3-15
4.1.	Linear Interpolation from Coarse Grid to Fine Grid	4-7
4.2.	Static Pressure for Supersonic Test Configuration SS3, $M = 2.0 \pm 0.25$, Generated with Time-accurate Roe Solver on 1000 Cell Grid	4-8
4.3.	Flow Interactions for Simulated Oscillating Piston in an Open Tube	4-9
4.4.	Time-accurate Development of the Stationary Flow Field for the Piston-In-Tube Configuration	4-10
4.5.	Static Pressure for Subsonic Virtual Piston Test Configuration SU1, Generated with Time-accurate Roe Solver on 1000 Cell Grid	4-11

Figure		Page
4.6.	Adapted Frequency Distributions for Three Different Augmentation Thresholds, Case SS3, 500 cells	4-13
4.7.	Nondimensional Static Pressure for Adapted Solutions at Three Different Augmentation Thresholds, Case SS3, 500 Cells . . .	4-13
4.8.	Adapted Frequency Distributions for Three Different Augmentation Thresholds, Case SU1, 500 Cells	4-14
4.9.	Nondimensional Static Pressure for Adapted Solutions at Three Different Augmentation Thresholds, Case SU1, 500 Cells . . .	4-14
4.10.	Change in Average Frequency Content with Changing Augmentation Threshold for Configurations SS3 and SU1 on 500 Cell Grid	4-16
4.11.	Change in Relative Compute Time with Changing Augmentation Threshold for Configurations SS3 and SU1 on 500 Cell Grid	4-17
4.12.	Computed Nondimensional Static Pressure for Adapted (Augmentation Threshold $7.0\text{e-}8$) and Non-adapted, 49-Frequency Harmonic Balance Solutions, Configuration SS4, 601 Cells, Compared with 1000 Cell Time-accurate Solution	4-20
4.13.	Fourier Coefficient Magnitudes (Momentum Component) for Adapted (Augmentation Threshold $7.0\text{e-}8$) and Non-adapted, 49-Frequency Harmonic Balance Solutions, Configuration SS4, 601 Cells	4-20
4.14.	Computed Nondimensional Static Pressure for Adapted (Augmentation Threshold $7.0\text{e-}8$, Variable CFL) and Non-adapted, 58-Frequency Harmonic Balance Solutions, Configuration SU1, 601 Cells, Compared with 1000 Cell Time-accurate Solution .	4-21
4.15.	Fourier Coefficient Magnitudes (Momentum Component) for Adapted (Augmentation Threshold $7.0\text{e-}8$, Variable CFL) and Non-adapted, 58-Frequency Harmonic Balance Solutions, Configuration SU1, 601 Grid Cells	4-21
4.16.	Computed Pressure for Time-accurate and Harmonic Balance Solutions at the Open End of the Tube	4-23

Figure		Page
4.17.	Steady-state Harmonic Balance Pressure Distribution at Time Sample Just after Period of Negative Velocity	4-23
4.18.	Variation in Adapted Frequency Distributions with Changing Grid Density for Test Configuration SS3, Augmentation Threshold $5.0\text{e-}8$	4-24
4.19.	Variation in Adapted Frequency Distributions With Changing Grid Density for Test Configuration SU1, Augmentation Threshold $5.0\text{e-}8$	4-24
4.20.	Reduction in Average Frequency Content and Run Time for Adapted vs Non-adapted Harmonic Balance Solutions	4-26
4.21.	Nondimensional Static Pressure for Adapted (Augmentation Threshold $7.0\text{e-}8$) Harmonic Balance Solutions, Configuration SS1, 1025 Cells, With and Without Multigrid Acceleration	4-27
5.1.	Unsteady Diverging Nozzle Configuration	5-3
5.2.	Time-accurate Development of the Fully-developed Unsteady Nozzle Flow Field	5-6
5.3.	Comparison of 0-Frequency (Steady State) and Theoretical Pressures for 1024 Dell Grid Diverging Nozzle	5-7
5.4.	Comparison of Time-average Adaptive Harmonic Balance Mass Flux with Theoretical Solution	5-9
5.5.	Error Distribution for Time-average Adaptive Harmonic Balance Mass Flux	5-9
5.6.	Evolution of the Adapted Frequency Distribution for Each Adaptation of the Solution, 1024 Cells, Augmentation Threshold $5.0\text{e-}7$	5-10
5.7.	Comparison of Adaptive Harmonic Balance and Time-accurate Solutions at 10 Points in Time Spanning One Period of Oscillation	5-12
5.8.	Root-Mean-Square of Difference Between Adaptive Harmonic Balance and Time-accurate Solutions with Changing Adaptation Threshold	5-13

Figure		Page
5.9.	Time Snapshot of Adaptive Harmonic Balance Solution Illustrating Refinement of Oscillating Normal Shock with Decreased Augmentation Threshold	5-13
5.10.	Adapted Frequency Distributions for Different Augmentation Thresholds on Coarse Grid (top) and Fine Grid (bottom) . .	5-15
5.11.	Change in Maximum Frequency Content of Adapted Harmonic Balance Solutions with Changes in Augmentation Threshold .	5-16
5.12.	Change in Average Frequency Content of Adapted Harmonic Balance Solutions with Changes in Augmentation Threshold .	5-17
5.13.	Increase in Relative Compute Time with Decreasing Augmentation Threshold	5-17
5.14.	Reduction in Average Frequency Content for Adapted vs Non-adapted Harmonic Balance Solutions	5-18
5.15.	Reduction in Compute Time for Adapted vs Non-adapted Harmonic Balance Solutions	5-19
5.16.	Efficiency of Adaptive Harmonic Balance Solutions (Run Time Reduction/Average Frequency Reduction)	5-20
B.1.	Effect of Splitting Error on 5 th Fourier Coefficient with Changing CFL for Supersonic Flow	B-6
B.2.	Effect of Splitting Error on 15 th Fourier Coefficient with Changing CFL for Supersonic Flow	B-7
B.3.	Effect of Splitting Error on 25 th Fourier Coefficient with Changing CFL for Supersonic Flow	B-8
B.4.	Effect of Splitting Error on 32 nd Fourier Coefficient with Changing CFL for Supersonic Flow	B-9
B.5.	Effect of Splitting Error on 5 th Fourier Coefficient with Changing CFL for Subsonic Flow	B-10
B.6.	Effect of Splitting Error on 15 th Fourier Coefficient with Changing CFL for Subsonic Flow	B-11
B.7.	Effect of Splitting Error on 25 th Fourier Coefficient with Changing CFL for Subsonic Flow	B-12

Figure		Page
B.8.	Effect of Splitting Error on 35 th Fourier Coefficient with Changing CFL for Subsonic Flow	B-13
B.9.	Effect of Splitting Error on 45 th Fourier Coefficient with Changing CFL for Subsonic Flow	B-14
B.10.	Effect of Splitting Error on Pressure with Changing CFL for Supersonic Flow	B-15
B.11.	Effect of Splitting Error on Pressure with Changing CFL for Subsonic Flow	B-15
C.1.	Effect of Series Length on Split Domain Solver Run Time . .	C-1

List of Tables

Table		Page
4.1.	Summary of Test Configuration Parameters	4-10
4.2.	Maximum Fourier Frequency and Stable CFL for Configuration SU1, 500 Cells, for Decreases in Augmentation Threshold . .	4-15
4.3.	Run-time Performance Comparison for Adapted vs Non-adapted Harmonic Balance Solutions, With and Without Multigrid Ac- celeration, Case SS1, 1025 Cells, in Seconds	4-27
5.1.	Candidate Numbers of Frequencies Available to the Adaptive Harmonic Balance Solver for the Unsteady Nozzle Test Case .	5-5
C.1.	Candidate Numbers of Fourier Frequencies, N , for an Efficient Adaptive Harmonic Balance Solution	C-2

Nomenclature

<i>Symbol</i>	<i>Description</i>
a	Unsteady disturbance amplitude
a	Nondimensional speed of sound, a_d/u_{ref}
a	Wave speed in linear model equation
a, b, c, d, e	Integer exponents to the prime numbers 2, 3, 5, 7, and 11, ideal factors for an FFT
a_n, b_n	n^{th} real Fourier cosine and sine coefficients
A	Area (length)
b	Source term in linear model equation
c_n	n^{th} complex Fourier coefficient
D	Jameson-Schmidt-Turkel artificial dissipation term
e	Base of the natural logarithm, approximately 2.718281828
e	Nondimensional specific internal energy, e_d/u_{ref}^2
E_N	Fraction of spectral energy in N th Fourier frequency
E_t	Nondimensional total energy, $(E_t)_d/(\rho_{\text{ref}} u_{\text{ref}}^2)$
E_{thresh}	Frequency augmentation threshold
f	Unsteady disturbance frequency (1/time)
f	Forcing function for a general nonlinear operator equation
$\hat{f}$	Forcing function for full approximation storage multigrid harmonic balance implementation
F	1-D Euler flux vector
F'	Quasi-1-D Euler flux vector
H	Nondimensional stagnation enthalpy, H_d/u_{ref}^2

$\mathbf{H}$	Quasi-1-D Euler equation source vector
i	$\sqrt{-1}$
imaginary number k	Fourier frequency index
k	Number of grids in a multigrid calculation
$K^{(2)}$	User-specified second order artificial dissipation parameter
$K^{(4)}$	User-specified Fourth order artificial dissipation parameter
$K^{(2g)}$	User-specified global second order artificial dissipation parameter
L_2 norm, R	Root-mean-square vector norm (residual) of a solution
M	Summation limit
M	Mach number
n	Fourier frequency index
$\hat{\mathbf{n}}$	Surface normal vector
N	Number of frequencies in a truncated Fourier series
p	Nondimensional static pressure, $p_d/(\rho_{\text{ref}} u_{\text{ref}}^2)$
P	Period of unsteady time-periodic flow (time)
$\mathbf{Q}$	Vector of conserved variables for 1-D Euler equation
$\mathbf{Q}'$	Vector of conserved variables for the quasi-1-D Euler equation
$\hat{\mathbf{Q}}$	Solution vector for the harmonic balance 1-D Euler equation
R	Residual of a discretized general conservation equation
$\hat{R}$	Residual of a discretized harmonic balance equation (frequency domain)
$\overline{R}$	Residual of a discretized harmonic balance equation (time domain)

$\hat{S}$	Harmonic balance source term
t	Nondimensional time variable, $t_d (\text{unit length})/u_{\text{ref}}$
u	Burgers' equation dependent variable
u	Nondimensional x velocity, u_d/u_{ref}
u	Approximate solution to a general nonlinear operator equation
U	Exact solution to a general nonlinear operator equation
V	Error in approximate solution to a general nonlinear operator equation, $U - u$
V	Volume
x	Nondimensional space variable, (unit reference length)
α	Magnitude of sampled data shift
γ	Ratio of specific heats, 1.4
$\delta_{i,j}$	Kronecker delta function for integers i, j
ΔA	Change in nozzle cross sectional area across a cell
Δt	Time-sample increment
Δx	Computational grid spacing (cell size)
$\Delta \hat{\xi}$	Change in the general vector of Fourier coefficients in one solver iteration
$\Delta \tau$	Computational pseudo-time step
$\varepsilon^{(2)}$	Second order JST artificial dissipation scaling factor
$\varepsilon^{(4)}$	Fourth order JST artificial dissipation scaling factor
κ_1	Initial residual-based adaptation trigger
κ_2	Repeat residual-based adaptation trigger
λ	Maximum eigenvalue of the Euler flux Jacobian
ν	Pressure scaling factor for JST artificial dissipation

ξ	General dependent variable
$\hat{\xi}$	General vector of Fourier coefficients
$\overline{\xi}$	General time-sampled dependent variable
π	Ratio of the circumference of a circle to its diameter, approximately 3.14159265
ρ	Nondimensional density, ρ_d/ρ_{ref}
τ	Pseudo-time variable
ϕ	Element of harmonic balance flux vector
Φ	General conservative flux function
$\hat{\Phi}$	General harmonic balance flux function (frequency-domain)
$\overline{\Phi}$	General harmonic balance flux function (time domain)
ω	Fundamental Fourier frequency (radians/time)
<i>Operators</i>	
$\mathbf{A}$	Shift-difference operator matrix
$\mathbf{D}^2$	Second order JST artificial dissipation operator
$\mathbf{D}^4$	Fourth order JST artificial dissipation operator
$\mathcal{F}$	Discrete Fourier transform
$\mathcal{F}^{-1}$	Inverse discrete Fourier transform
I	Identity operator
I_h^{2h}	Multigrid restriction operator
I_{2h}^h	Multigrid prolongation operator
L	General nonlinear operator
L_f	Split-domain frequency-domain ODE update operator
L_t	Split-domain time-domain PDE update operator

$\mathcal{L}$	Second order difference operator for an ordinary differential equation
$\mathcal{M}$	Second order difference operator for a partial differential equation
O	Computational complexity upper bound
$\underline{\Delta}$	First order forward difference
$\overline{\nabla}$	First order backward difference
Ω	Computational complexity lower bound
$\sim$	Complex conjugation
<i>Subscripts</i>	
adj	Interior cell adjacent to boundary
b	On boundary
d	Dimensional quantity
g	Boundary ghost cell
i	General grid cell index
l	Left cell interface
max	Maximum value
n	Fourier frequency number
r	Right cell interface
ref	Dimensional reference value
x	Derivative with respect to x
<i>Superscripts</i>	
$*$	Intermediate time level
h	Fine grid level
$2h$	Coarse grid level

k	Multigrid grid level
n	Time (iteration) level

Abbreviations

<i>Abbreviation</i>	<i>Description</i>
1-D	One-dimensional
2-D	Two-dimensional
CFD	Computational Fluid Dynamics
CFL	Courant, Friedrichs, Lewy stability criteria
DFT	Discrete Fourier Transform
FAS	Full Approximation Storage
FFT	Fast Fourier Transform
IDFT	Inverse Discrete Fourier Transform
JST	Jameson, Schmidt, Turkel
ODE	Ordinary Differential Equation
PDE	Partial Differential Equation
RK	Runge-Kutta
RMS	Root-Mean-Square

Abstract

A new adaptive split-domain harmonic balance computational fluid dynamics (CFD) method is developed to solve highly nonlinear time-periodic flows such as those found in transonic turbomachinery. The basic harmonic balance CFD method transforms an unsteady time-periodic problem into a steady-state problem by assuming a solution in the form of a Fourier series in time. The new method employs a unique multi-domain split-operator solution technique to remove a large-series stability restriction present in previous harmonic balance CFD approaches. The new method also minimizes the computational work required to obtain a harmonic balance solution by adapting the frequency content to the flow, starting with a small number of Fourier frequencies and augmenting the frequency content in each cell as necessary to capture local flow physics. The method reduces compute times by allowing larger integration time steps, reducing Fourier transform calculations, and reducing overall problem size. Split-domain solutions to the 1-D inviscid Burgers' equation are computed with up to 97 frequencies, demonstrating improved stability. Differences between harmonic balance solutions and time-accurate solutions are found to be asymptotic with Fourier series length. The adaptive split-domain approach is applied to the 1-D and quasi-1-D Euler equations. Supersonic and subsonic Euler calculations show that the adapted and non-adapted harmonic balance solutions are equivalent. Accurate adapted quasi-1-D Euler solutions for a supersonic/subsonic diverging nozzle with periodic unsteady outflow conditions are generated in 86% less time than an equivalent non-adapted split-domain solution, demonstrating the benefit of adapting frequency content to local flow conditions.

ADAPTIVE HARMONIC BALANCE METHOD FOR UNSTEADY, NONLINEAR, ONE-DIMENSIONAL PERIODIC FLOWS

I. Introduction

High-fidelity numerical simulations of fluid flow through transonic turbomachinery are of considerable interest to designers of compressors and turbines in modern jet engines. Solutions can be obtained with conventional time-accurate computational fluid dynamics (CFD) codes, but the considerable time required to generate these solutions limits their utility to the designer. For the class of problems where the flow can be assumed to be fully developed and periodic in time, such as flow past a rotor with oscillating blades or flow through a rotor-stator, time-accurate calculations can be particularly inefficient. It is usually necessary to step through many disturbance periods before a fully developed solution is reached. To achieve shorter computation times for this class of problem, CFD techniques have been developed that take advantage of the time-periodic nature of the flow. These include the time-linearization technique (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12), the time-averaging technique (13, 14, 15), and the harmonic balance technique (16, 17, 18, 19).

Of the three CFD techniques developed specifically for time-periodic flows, the harmonic balance technique is most suitable for modeling the large disturbances and strong nonlinearities found in transonic turbomachinery. All three techniques are closely related in that they all assume a harmonic form for the unsteady flow, recast the unsteady problem as a steady-state calculation, and employ convergence acceleration techniques to reduce computation time. However, the time-linearization and time-averaging techniques have restrictions on the unsteadiness that are not present

in the harmonic balance method. The time-linearized technique models unsteadiness as a small-amplitude perturbation linearized about an nonlinear steady-state background flow. Linearization of the unsteadiness, combined with a small-disturbance assumption, limits applicability to transonic flows. In the time-averaged technique a small harmonic perturbation is time-averaged with the steady-state background equations in a process similar to Reynolds averaging. Additional stress terms result that capture some nonlinearities in the unsteady flow, but a small disturbance assumption is still made. Finally, both the time-linearized and time-averaged techniques solve for only one harmonic at a time. Multiple-harmonic solutions are built up by superposing several single-harmonic solutions, losing the effect of nonlinear harmonic coupling. The harmonic balance method, on the other hand, has no small disturbance restriction, and solves a set of nonlinear equations for several harmonics simultaneously, more accurately capturing coupled nonlinear effects. This makes it uniquely suited for calculating flow through transonic turbomachinery.

The fidelity of a harmonic balance solution is dependent on grid density, on the number of harmonics included, and on the flow being modeled. On a given grid, a flow that is smoothly unsteady (i.e., without moving discontinuities) will require fewer harmonics than a flow containing a moving shock to achieve the same level of fidelity. Because the computational cost increases with each harmonic included in the solution, it is desirable to use the minimum number of harmonics needed for a given problem.

In a typical transonic turbomachinery problem, the nature of the flow can vary significantly throughout the domain of interest. This is illustrated in Fig. 1.1 (20), which shows experimental time-pressure plots at two locations on a single inlet guide vane upstream of a transonic rotor. Existing harmonic balance implementations solve for a constant number of harmonics over the entire computational domain, so problems containing both smooth and discontinuous unsteadiness require a compromise between reduced run time (fewer harmonics), and fidelity in the regions of




Figure 1.1 Experimental Pressure Data on Inlet Guide Vane Upstream of Transonic Rotating Compressor Blade Row for 5% Span (top), and 95% Span (bottom)

strong nonlinearity. The goal of this research is to remove the need for compromise by extending the harmonic balance technique to allow a variable number of included harmonics, and to automatically identify and apply, on a cell-by-cell basis, the minimum number of harmonics required to achieve a consistent fidelity throughout the computational domain.

1.1 Overview of Harmonic Balance Method

The harmonic balance method has been used for many years as a means of analyzing the behavior of harmonic ordinary differential equations (ODEs) (21). The technique consists of assuming a solution in the form of a truncated Fourier series with a predetermined number of harmonics, substituting the assumed solution into the ODE, and algebraically manipulating the results to collect terms of like frequency. Any resulting terms with a frequency not in the Fourier series are dropped. Each harmonic is then balanced by requiring that like-frequency terms on each side of the equation satisfy equality independently. Balancing results in a system of cou-

pled algebraic equations which are solved for the Fourier coefficients of the assumed solution.

When the same technique is applied to a partial differential equation (PDE), the result is a system of ODEs or PDEs that are solved for the Fourier coefficients. For example, consider a generic scalar one-dimensional Conservation Law equation in differential form

$$\frac{\partial \xi}{\partial t} + \frac{\partial \Phi(\xi)}{\partial x} = 0. \quad (1.1)$$

In this equation, $\xi(x, t)$ is the scalar dependent variable, and Φ is a flux term that depends on ξ .

If the time response of the conserved variable ξ at an arbitrary point in space, x , is assumed to be periodic in time with radian frequency ω , then that response can be approximated by a truncated complex Fourier series

$$\xi(x, t) \approx \sum_{n=-N}^N c_n(x) e^{in\omega t}. \quad (1.2)$$

In this series, the coefficients, c_n , are functions of x only, while the exponential terms are functions of time only. To derive the harmonic balance form of the conservation equation, the approximating series is substituted into the conservation equation which, after some algebraic manipulation, is transformed into a system of ordinary differential equations dependent only on x , which are solved for the unknown Fourier coefficients. Because the variable $\xi(x, t)$ is real, the Fourier coefficients corresponding to $\pm n$ are complex conjugates of each other. Equations associated with negative frequencies can therefore be dropped, leaving $N + 1$ complex differential equations for $2N + 1$ unknowns. In vector form, these equations are given by

$$\frac{d\hat{\Phi}(\hat{\xi})}{dx} + \hat{\mathbf{S}}(\hat{\xi}) = 0 \quad (1.3)$$

where $\hat{\xi}$ is the vector of Fourier coefficients, c_n , $\hat{\Phi}$ is a vector of harmonic balance flux terms, and $\hat{\mathbf{S}}$ is a vector of source term arising from the time derivative in Eq. 1.1.

For problems of interest, Eq. 1.3 is highly nonlinear. A pseudo-time derivative of the dependent vector $\hat{\xi}$ is added (Eq. 1.4) so that a time-marching solution method can be applied to find the steady-state solution

$$\frac{\partial \hat{\xi}}{\partial \tau} + \frac{\partial \hat{\Phi}(\hat{\xi})}{\partial x} + \hat{\mathbf{S}}(\hat{\xi}) = 0. \quad (1.4)$$

When steady state is reached, the added pseudo-time derivative is equal to zero, and Eq. 1.3 is recovered.

Once a steady-state solution is obtained, useful information is recovered from the calculated Fourier coefficients. The values of the dependent variable at any point in time are easily reconstructed from the computed coefficients by performing the summation in Eq. 1.2. The time-average values of the dependent variable are also readily obtained, as they are the computed coefficients c_0 .

Previous Work in Harmonic Balance. The first use of the harmonic balance method in CFD was by Hall, Thomas, and Clark (16), who implemented a harmonic balance Reynolds-Averaged Navier-Stokes solver. They applied it to a single two-dimensional compressor blade row undergoing forced periodic vibration under transonic flow conditions. This configuration contained moderately strong shocks, but limited shock motion. Solutions containing up to 7 frequencies in the approximating series were calculated. Solutions containing 3–5 frequencies showed good agreement with standard time-accurate calculations for the time-average and first harmonic terms, but took approximately one-tenth as long to compute. Solutions with 7 frequencies failed to converge. In a follow-on effort, Thomas, Dowell, and Hall (18) coupled an inviscid harmonic balance solver with a linear structural model to study limit cycle oscillations of an aeroelastic system.

Additional work in harmonic balance CFD was conducted by McMullen, Jameson, and Alonso (17, 19), who investigated several aspects of the harmonic balance method not previously addressed. Using a somewhat different formulation than Hall *et al.*, they performed a stability analysis of the method, concluding that stability could be an issue when large numbers of frequencies are included in the solution. They applied the method to several test problems, including unsteady 1-D channel flow, 2-D oscillating flow behind a cylinder in crossflow, and a pitching airfoil. Their results show good agreement with analytical and numerical predictions for both time dependent and time-averaged data. Early cylinder crossflow calculations fixed the fundamental frequency at the theoretical frequency of vortex shedding (17). Later cylinder crossflow and pitching airfoil calculations employed a gradient-based optimization approach to dynamically determine the correct fundamental frequency (19).

1.2 Scope

The objective of this research is to extend the harmonic balance technique to accurately resolve unsteady, nonlinear, periodic flows while minimizing computational cost. To achieve this objective, the solution technique must be able to robustly solve for as many harmonics as necessary to attain the desired degree of solution fidelity. To maintain overall efficiency and minimize computational cost, higher harmonics should only be included where required by local flow conditions. The scope of this research is defined by the following thesis statement and assumptions:

1.2.1 Thesis Statement. A spatially-adaptive harmonic balance method can be implemented to accurately and efficiently compute a stationary time-periodic flow field containing regions of smooth flow and regions with strong moving discontinuities by automatically adjusting the number of frequencies in the solution, on a point-by-point basis, to match local flow conditions.

1.2.2 Assumptions. The implementation and effectiveness of the spatially-adaptive harmonic balance technique are not strongly dependent on the spatial dimension of the problem, and will be adequately demonstrated in one spatial dimension. Efficiency will be measured by comparing average frequency content and run time with a comparable non-adapted harmonic balance implementation. Accuracy will be determined through comparison with conventional time-accurate results and non-adapted harmonic balance results.

1.3 Research Approach

Several one-dimensional CFD codes were developed to investigate different aspects of the adaptive harmonic balance method. The initial work investigated harmonic balance solutions of the inviscid Burgers' equation, a scalar analog of the Euler equation, subject to a variety of periodic boundary conditions. Next, subsonic and supersonic harmonic balance solutions to Euler's equation were investigated. Finally, the adaptive split-domain harmonic balance method was demonstrated by solving the quasi-1-D Euler's equation for an unsteady diverging nozzle configuration.

Inviscid Burgers' Equation: Different implementations of the harmonic balance Burgers' equation were tested to determine which implementation provided the best performance, robustness, and accuracy. A new split-domain implementation was shown to be a superior harmonic balance implementation when large numbers of harmonics were included. The impact of shock strength, fundamental frequency, and grid density on the harmonic balance solution were investigated.

Euler's Equation: An adaptive split-domain harmonic balance Euler solver was written to investigate the frequency augmentation approach and to develop methods for properly treating transitions from one frequency to another at arbitrary points in the computational grid. A study of the impact of user-specified parameters on

the performance of the adaptive algorithm for both supersonic and subsonic flows was conducted.

Quasi-1-D Euler's Equation: The adaptive split-domain harmonic balance algorithm is demonstrated by solving unsteady diverging nozzle flow. This configuration is representative of anticipated production applications, containing regions of both low and high frequency content, with abrupt transitions between the two.

1.4 Document Organization

The remainder of this document is organized as follows:

Chapter II: Details the theory and implementation of the adaptive split-domain harmonic balance method.

Chapter III: Records the analysis and results of the harmonic balance method applied to the inviscid 1-D Burgers equation for a variety of periodic boundary conditions.

Chapter IV: Records the analysis and results of the adaptive split-domain harmonic balance method applied to the one-dimensional Euler equation for both supersonic and subsonic boundary conditions.

Chapter V: Presents a demonstration of the adaptive split-domain harmonic balance technique applied to an unsteady quasi-1-D nozzle flow.

Chapter VI: Summarizes the results and conclusions of the current research

Chapter VII: Discusses future research topics suggested by the current research.

II. Theory and Implementation of Adaptive Split-Domain Harmonic Balance

2.1 Introduction

In this chapter, a theoretical and practical foundation for the adaptive split-domain harmonic balance CFD method is established. The chapter begins by expanding on the brief overview of harmonic balance in CFD presented in Section 1.1 with a more in-depth derivation of the basic harmonic balance method and its application to the solution of conservation equations. This is followed by a brief discussion of previous harmonic balance CFD implementations, and a detailed description of the new split-domain harmonic balance approach. The focus of the discussion then turns to frequency augmentation and the adaptive harmonic balance algorithm. Finally, the chapter concludes with a discussion of the implementation of full approximation scheme (FAS) multigrid convergence acceleration with adaptive split-domain harmonic balance.

The discussion that follows is general. No specific solver implementation is described, and no assumptions are made about the specific conservation equations being solved or the discretization schemes used to obtain a numerical solution. Because of this, the concepts and algorithms described below are easily adapted to new applications. In the chapters that follow, some or all of the ideas discussed below are applied to a MacCormack discretization of the 1-D Burgers' equation (Chapter III), as well as cell-centered finite volume formulations of the 1-D Euler's equation (Chapter IV) and quasi-1-D Euler's equation (Chapter V).

2.2 Derivation of the Harmonic Balance Equations

In Section 1.1, a brief overview of the application of the harmonic balance method to a scalar conservation law was presented. The following discussion expands on that brief overview, providing a detailed derivation of the harmonic balance form

of the conservation law, and covering some aspects of numerical implementation. The derivation is carried out based on the generic scalar conservation law, Eq. 1.1, repeated here as Eq. 2.1

$$\frac{\partial \xi}{\partial t} + \frac{\partial \Phi(\xi)}{\partial x} = 0. \quad (2.1)$$

A similar derivation for a specific conservation law, the 1-D inviscid Burgers' equation (Eq. 3.1) is presented in Appendix A.

As discussed in Section 1.1, the harmonic balance method centers on the assumption that the dependent variable, ξ , can be approximated by a truncated Fourier series

$$\xi(x, t) \approx \sum_{n=-N}^N c_n(x) e^{in\omega t}. \quad (2.2)$$

In this series, N is the number of positive and negative frequencies in the truncated Fourier series, and ω is the fundamental radian frequency of the series, equal to $2\pi/P$ where P is the period of oscillation of the physical system. The complex coefficients, c_n , are functions of the spatial variable only, while the exponentials are functions of time. Substituting Eq. 2.2 into the first term on the left hand side of Eq. 2.1 and evaluating the time derivative yields

$$\frac{\partial \xi}{\partial t} \approx \sum_{n=-N}^N in\omega c_n e^{in\omega t}. \quad (2.3)$$

This expression is the origin of the harmonic balance source term.

A similar substitution is performed for the second term on the left hand side of Eq. 2.1. In this case, substitution requires the evaluation of the flux function, Φ , with the approximating series as an argument. Upon evaluation, the resulting flux is algebraically simplified and terms of like frequency are collected, resulting in an expression of the form

$$\Phi(\xi) \approx \sum_{n=-M}^M \phi_n e^{in\omega t} \quad (2.4)$$

where ϕ_n contains the sum of all coefficients of the exponential $e^{in\omega t}$. If the evaluation of Φ requires products or integer powers of the dependent variable ξ , then $M > N$. As part of the harmonic balance approximation, all exponential terms with frequency greater than $\pm N\omega$ are discarded, leaving only terms with frequencies present in the approximating series

$$\Phi(\xi) \approx \sum_{n=-N}^N \phi_n e^{in\omega t}. \quad (2.5)$$

As a simple example of the derivation of a harmonic balance flux, consider a flux function $\Phi(\xi) = \xi^2$, and an approximating Fourier series with $N = 1$. Evaluation of the flux with the approximating series as the argument yields

$$\Phi(\xi) \approx \left(\sum_{n=-1}^1 c_n e^{in\omega t} \right)^2. \quad (2.6)$$

Expanding the right-hand side of Eq. 2.6 gives

$$\left(\sum_{n=-1}^1 c_n e^{in\omega t} \right)^2 = c_{-1}^2 e^{-i2\omega t} + c_{-1}c_0 e^{-i\omega t} + c_{-1}c_1 + c_0c_{-1}e^{-i\omega t} + \quad (2.7)$$

$$c_0^2 + c_0c_1 e^{i\omega t} + c_1c_{-1} + c_1c_0 e^{i\omega t} + c_1^2 e^{i2\omega t}$$

$$= c_{-1}^2 e^{-i2\omega t} + 2c_{-1}c_0 e^{-i\omega t} + 2c_{-1}c_1 + c_0^2 +$$

$$2c_0c_1 e^{i\omega t} + c_1^2 e^{i2\omega t}. \quad (2.8)$$

Equation 2.8 contains exponential terms with frequencies $\pm 2\omega$. Since the assumed solution contains frequencies only up to $\pm\omega$, these higher frequency terms are discarded. The remaining terms in Eq. 2.8 are the harmonic balance flux terms for this simple example:

$$\phi_{-1} = 2c_{-1}c_0 \quad \phi_0 = 2c_{-1}c_1 + c_0^2 \quad \phi_1 = 2c_1c_0. \quad (2.9)$$

Substitution of Eq. 2.3 and Eq. 2.5 into the left hand side of Eq. 2.1 and collection of terms with like exponentials yields the series approximation form of the conservation equation

$$\sum_{n=-N}^N \left[in\omega c_n + \frac{d\phi_n}{dx} \right] e^{in\omega t} = 0. \quad (2.10)$$

At this point, the harmonic terms are required to balance across the equality; each frequency is required to satisfy equality independently of the other frequencies. Since the right hand side of Eq. 2.10 is identically zero, each of the terms in the summation on the left hand side must be equal to zero, i.e.,

$$\left[in\omega c_n + \frac{d\phi_n}{dx} \right] e^{in\omega t} = 0 \quad -N \leq n \leq N. \quad (2.11)$$

The terms inside the square brackets must be identically equal to zero for each of the $2N + 1$ complex equations represented by Eq. 2.11 to hold for all time. The exponential terms are dropped, leaving a system of $2N + 1$ coupled ordinary differential equations for $2N + 1$ complex coefficients. This system can be further simplified because the dependent variable, $\xi(x, t)$, is real and periodic. The negative frequency Fourier coefficients are thus the complex conjugates of the corresponding positive frequency coefficients, i.e., $c_{-n} = \tilde{c}_n$. Furthermore, since the zero-frequency term c_0 must be real, the total number of unknowns is $2N + 1$: the real and imaginary parts of N positive-frequency coefficients, plus the real zero-frequency coefficient. Since the coefficients of the negative frequencies are not independent, the corresponding equations need not be solved. Keeping just the positive frequency equations and eliminating the exponential terms, Eq. 2.11 reduces to

$$in\omega c_n + \frac{d\phi_n}{dx} = 0 \quad 0 \leq n \leq N, \quad (2.12)$$

or in vector form,

$$\frac{d\hat{\Phi}(\hat{\xi})}{dx} + \hat{\mathbf{S}}(\hat{\xi}) = 0 \quad (2.13)$$

with

$$\hat{\xi} = \begin{bmatrix} c_0 \\ c_1 \\ \vdots \\ c_N \end{bmatrix}, \quad \hat{\mathbf{S}}(\hat{\xi}) = \begin{bmatrix} 0 \\ i\omega c_1 \\ \vdots \\ iN\omega c_N \end{bmatrix}, \quad \hat{\Phi}(\hat{\xi}) = \begin{bmatrix} \phi_0 \\ \phi_1 \\ \vdots \\ \phi_N \end{bmatrix}. \quad (2.14)$$

To solve Eq. 2.13, a pseudo-time derivative of the vector of Fourier coefficients, $\hat{\xi}$, is added

$$\frac{\partial \hat{\xi}}{\partial \tau} + \frac{\partial \hat{\Phi}(\hat{\xi})}{\partial x} + \hat{\mathbf{S}}(\hat{\xi}) = 0. \quad (2.15)$$

A time-marching scheme is used to drive the solution to steady state. Once steady state is reached, the pseudo-time derivative vanishes, and Eq. 2.13 is recovered.

Eq. 2.15 is the vector harmonic balance form of the scalar one-dimensional conservation equation (Eq. 2.1). Derivation of the harmonic balance form for multi-dimensional and vector conservation equations is similar to that of the scalar equation. For a multi-dimensional problem, multiple flux derivatives must be computed, but the basic harmonic balance form remains unchanged. In the case of a vector equation (e.g. the Euler equation), each of the conserved (or primitive) variables is expanded in a separate Fourier series. The harmonic balance solution vector is a concatenation of the Fourier coefficients for all the approximated variables. An example of a vector equation derivation is found in (16).

Evaluating the harmonic balance flux vector, $\hat{\Phi}$, can be computationally expensive. For a simple flux like that of the scalar 1-D inviscid Burgers' equation (Eq. 3.1), the asymptotic complexity of the harmonic balance flux calculation is $O(N^2)$. For a more complex vector equation such as the Euler equation (Eq. 4.1), the asymptotic complexity lies between $O(N^3)$ and $O(N^4)$ (16). Considerable computational savings can be realized by taking a multi-domain approach for the flux calculation.

A multi-domain approach takes advantage of the relationship between the computed Fourier coefficients, $\hat{\xi}$, and their inverse discrete Fourier transform (IDFT) (16). Given the vector of Fourier coefficients $\hat{\xi}$, and $\mathcal{F}$, a discrete Fourier transform (DFT) operator that produces positive-frequency Fourier coefficients from a set of real numbers, define $\bar{\xi}$ such that

$$\hat{\xi} = \mathcal{F}\bar{\xi} \quad (2.16a)$$

$$\bar{\xi} = \mathcal{F}^{-1}\hat{\xi}. \quad (2.16b)$$

The vector $\bar{\xi}$ is a real vector of length $2N + 1$ that contains values of the dependent variable, $\xi(x, t)$, sampled at times $t = (0, \Delta t, 2\Delta t, \dots, 2N\Delta t)$, with $\Delta t = \frac{2\pi}{\omega(2N+1)}$, the period of oscillation divided by the number of samples. A similar relationship is assumed for the flux vector, $\hat{\Phi}$, and its IDFT (16).

Given these relationships, the harmonic balance flux, $\hat{\Phi}$, can be closely approximated by calculating $\bar{\xi}$ from $\hat{\xi}$, applying the time domain flux function Φ to each element of $\bar{\xi}$ to obtain a vector of time-sampled fluxes, $\bar{\Phi}$, and transforming the result back to the frequency domain (16). The asymptotic complexity of the DFT/IDFT for arbitrary numbers of terms is $O(N^2)$ (22). The complexity of the time-domain flux calculation is $O(N)$, so the complexity of the combined operation is $O(N^2)$. To put the potential computational savings in perspective, consider an Euler flux calculation with $N = 45$, and assume a computational cost proportional to N^2 for the multi-domain approach, and $N^{3.5}$ for the direct approach. Disregarding any constants of proportionality, the direct approach is over 300 times more expensive than the multi-domain approach.

In practice, the theoretical computational cost savings could be much higher. The asymptotic complexity of a modern DFT/IDFT algorithm has a lower bound of $\Omega(N \log_2 N)$ (22). The lower bound is typically obtainable only when the number of terms being transformed is a power of 2, however. Since the number of terms

transformed in the harmonic balance calculation is always odd, the ideal performance is never achieved. However, very good performance can still be achieved if the number of terms can be factored into integer powers of small primes (22), i.e., when $2N + 1 = 2^a 3^b 5^c 7^d 11^e$ for some integers a, b, c, d , and e . So the actual computational cost will lie between $N \log_2 N$ and N^2 , depending on the value of N . For the above example with $N = 45$, a direct flux calculation is between 300 and 2400 times more costly than a multi-domain calculation, disregarding constants of proportionality.

2.3 Prior Implementations of Harmonic Balance for CFD

Existing harmonic balance CFD solvers have taken two different approaches to implementing the harmonic balance equations. Both are multi-domain approaches. The first approach, developed by Hall, Thomas, and Clark (16), results in a system of equations that is integrated in the time domain. The second, developed by McMullen, Jameson, and Alonso (17), results in a system of equations integrated in the frequency domain.

Derivation of the time-domain approach begins by substituting Eq. 2.16a and the multi-domain flux calculation into Eq. 2.15 to obtain

$$\frac{\partial(\mathcal{F}\bar{\xi})}{\partial\tau} + \frac{\partial(\mathcal{F}\bar{\Phi}(\bar{\xi}))}{\partial x} + \hat{\mathbf{S}}(\mathcal{F}\bar{\xi}) = 0. \quad (2.17)$$

An IDFT applied to Eq. 2.17 yields the final form of the time-domain harmonic balance approach.

$$\frac{\partial\bar{\xi}}{\partial\tau} + \frac{\partial\bar{\Phi}(\bar{\xi})}{\partial x} + \mathcal{F}^{-1}\hat{\mathbf{S}}(\mathcal{F}\bar{\xi}) = 0. \quad (2.18)$$

Eq. 2.18 has the form of the original conservation equation, Eq. 2.1, with an added source term. The primary advantage of this formulation is that an existing steady-state solver can be easily modified to solve the harmonic balance equations.

Derivation of the frequency-domain formulation begins by rewriting Eq. 2.15 as

$$\frac{\partial \hat{\xi}}{\partial \tau} + \hat{R}(\hat{\xi}) + \hat{\mathbf{S}}(\hat{\xi}) = 0, \quad (2.19)$$

where $\hat{R}(\bar{\xi})$ is a residual containing the flux derivative term, $\frac{d\hat{\Phi}}{dx}$, (or in multiple dimensions, the sum of the flux derivatives for each dimension) as well as any terms required by a specific discretization, such as artificial dissipation terms. A frequency domain/time domain relationship is assumed for the residual, such that

$$\hat{R}(\hat{\xi}) = \mathcal{F}\overline{R}(\mathcal{F}^{-1}\hat{\xi}) = \mathcal{F}\overline{R}(\bar{\xi}) \quad (2.20)$$

where $\overline{R}(\bar{\xi})$ contains the flux derivatives and additional terms evaluated at $2N + 1$ points in time. Substitution of Eq. 2.20 into Eq. 2.19 yields the second harmonic balance form,

$$\frac{\partial \hat{\xi}}{\partial \tau} + \mathcal{F}\overline{R}(\mathcal{F}^{-1}\hat{\xi}) + \hat{\mathbf{S}}(\hat{\xi}) = 0. \quad (2.21)$$

Despite being integrated in the frequency domain, this formulation is very similar to the time-domain formulation. Due to the linearity of $\mathcal{F}$ and $\mathcal{F}^{-1}$, if the same discretization scheme is applied to both formulations, and the spatial discretization of $\overline{R}(\bar{\xi})$ contains only linear difference operators, then Eqs. 2.18 and 2.21 are mathematically equivalent.

2.4 Split-Domain Harmonic Balance

The time-domain and frequency-domain formulations of the harmonic balance equations are adequate for calculating solutions with small N , but both have aspects that reduce their utility for solutions requiring large N . The primary concern with the existing formulations is the stability of the resulting harmonic balance equations. A stability analysis of the frequency-domain approach indicates that as N becomes large, the stability limit of the approach becomes restricted (17), requiring a smaller time step and reducing efficiency. This property was confirmed experimentally for

both the time-domain and frequency domain formulations, as well as for the direct formulation (see Section 3.4.2).

The second, lesser concern with the existing formulations is the number of transforms required per point, per iteration. A multi-step time integration technique such as an explicit 3-, 4-, or 5-stage Runge-Kutta scheme requires one DFT/IDFT evaluation per point for each stage of the integrator. For small N , this is of little consequence, but for large N the cost of multiple DFT/IDFT evaluations could become significant.

A new split-domain harmonic balance formulation was developed that addresses both the stability concern and the DFT/IDFT evaluation concern. In this formulation, the inhomogeneous harmonic balance equation (Eq. 2.15) is split into a homogeneous partial differential equation and an ordinary differential equation (23).

$$\frac{\partial \hat{\xi}_1}{\partial \tau} + \frac{\partial \hat{\Phi}(\hat{\xi}_1)}{\partial x} = 0 \quad (2.22a)$$

$$\frac{d \hat{\xi}_2}{d \tau} + \hat{\mathbf{S}}(\hat{\xi}_2) = 0. \quad (2.22b)$$

An approximate solution to Eq. 2.15 is obtained by sequentially solving Eqs. 2.22a and 2.22b, using the solution from one as the initial condition for the other. For example, one could solve Eq. 2.22a for $\hat{\xi}_1(x, \tau)$ subject to the initial condition $\hat{\xi}_1(x, 0) = \hat{\xi}(x, 0)$, the initial condition for Eq. 2.15. The approximate solution is then obtained by solving Eq. 2.22b for $\hat{\xi}_2$ with the initial condition $\hat{\xi}_2(0) = \hat{\xi}_1(x, \tau)$. This split-operator approach is sometimes used to solve stiff systems of equations (i.e. systems with widely differing time scales) such as those that result when finite-rate chemistry is included in a CFD solution (24, 25).

An approximate numerical solution can be obtained using a Strang symmetric splitting approach (26, 24, 23). Given difference operators $\mathcal{M}_{\Delta\tau}$ and $\mathcal{L}_{\Delta\tau}$ that are second-order accurate for Eqs. 2.22a and 2.22b respectively, then the composed operator $\mathcal{L}_{\Delta\tau/2} \mathcal{M}_{\Delta\tau} \mathcal{L}_{\Delta\tau/2}$ is second-order accurate for Eq. 2.15 (26). Furthermore,

the stability characteristics of the composed operator are determined by the stability characteristics of the individual operators $\mathcal{M}_{\Delta\tau}$ and $\mathcal{L}_{\Delta\tau}$ (26, 23). By choosing appropriate discretizations for Eqs. 2.22a and 2.22b, the large- N time step restriction is greatly reduced or eliminated.

To take advantage of the efficiencies of the multi-domain approach, Eq. 2.22a is transformed to the time domain in a manner similar to Eq. 2.18. Together with Eq. 2.22b, this transformed equation becomes the split-domain harmonic balance form of Eq. 1.1.

$$\frac{\partial \bar{\xi}}{\partial \tau} + \frac{\partial \bar{\Phi}(\bar{\xi})}{\partial x} = 0 \quad (2.23a)$$

$$\frac{d\hat{\xi}}{d\tau} + \hat{\mathbf{S}}(\hat{\xi}) = 0. \quad (2.23b)$$

The steps required to advance the solution one iteration from time level n to time level $n + 1$ are:

step 1: Given the solution vector at time level n , $\hat{\xi}^n$, advance Eq. 2.23b one-half pseudo-time step to produce $\hat{\xi}^*$

step 2: Calculate $\bar{\xi}^* = \mathcal{F}^{-1}\hat{\xi}^*$

step 3: With $\bar{\xi}^*$, advance Eq. 2.23a one full pseudo-time step to produce $\bar{\xi}^{**}$

step 4: Calculate $\hat{\xi}^{**} = \mathcal{F}\bar{\xi}^{**}$

step 5: With $\hat{\xi}^{**}$, advance Eq. 2.23b one-half pseudo-time step to produce $\hat{\xi}^{n+1}$

The split-domain harmonic balance approach has several advantages over previous harmonic balance implementations. First and foremost, it greatly reduces the stability restriction for large N . Secondly, it requires only one DFT/IDFT calculation per point per iteration, regardless of the time integration scheme employed. Another advantage lies in the fact that the system of equations represented by the time-domain PDE, Eq. 2.23a, is uncoupled, and each equation has the same form as

the original conservation equation. Thus a solution scheme developed for the original equation can be applied in-turn, without modification, to each equation in the harmonic balance system. This makes the split-domain approach easy to implement in an existing code.

The split-domain harmonic balance formulation has one potential disadvantage. As a result of splitting the harmonic balance equation, a new discretization error is introduced. This error causes the steady-state solution to be dependent on the numerical time-integration step size, $\Delta\tau$ (25). An investigation of the impact of splitting error on split-domain harmonic balance solutions was conducted (Appendix B) and it was found that splitting error is not a significant factor for the solver implementations studied in this effort. It could become a factor, however, for an implementations that allow very large numerical time steps.

Like the time-domain PDE, the system of equations represented by the frequency domain ODE, Eq. 2.23b, is also uncoupled, with each equation having the form

$$\frac{dc_k}{d\tau} + i k \omega c_k = 0, \quad 0 \leq k \leq N. \quad (2.24)$$

Eq. 2.24 is a linear ODE with an easily obtained exact solution. Given a solution c_k^n at time level n as an initial condition, the exact solution to Eq. 2.23b at pseudo-time $\Delta\tau/2$ is given by

$$c_k^{n+\frac{1}{2}} = c_k^n e^{-ik\omega\frac{\Delta\tau}{2}}, \quad 0 \leq k \leq N. \quad (2.25)$$

Insight into the split-domain harmonic balance iterative solution process is gained by recognizing that Eq. 2.25 represents a small shift of the solution vector elements in physical time. The time-shift property for the discrete Fourier transform states that, given the time sampled sequence $\bar{\xi}$ with length $2N + 1$, and its discrete Fourier

transform $\hat{\xi}$, then (27)

$$\bar{\xi}_{k-\alpha} \Leftrightarrow \hat{\xi}_n e^{-\frac{i2\pi n\alpha}{2N+1}}, \quad (2.26)$$

where the subscripts k and n range over the elements in the vector. The right hand side of Eq. 2.25 is similar to that of Eq. 2.26, provided

$$\alpha = \frac{(2N+1)\omega\Delta\tau}{4\pi}. \quad (2.27)$$

Thus the combined result of the operations represented by the split-domain solution steps 4, 5, 1, and 2 (in that order) approximates a small physical-time shift of the elements in the vector $\bar{\xi}$ along the continuous periodic function given by the approximating Fourier series, Eq. 2.2.

With this understanding, the split-domain harmonic balance solution process can be interpreted as a sequence of forward integrations of Eq. 2.23a in pseudo-time (which, because Eq. 2.23a has the same form as Eq. 2.1, is just a scaled integration in physical time), surrounded by backward physical-time shifts of the elements in the solution vector (Fig. 2.1). The shift-integrate-shift sequence is continued until the end values are the same as the starting values, at which point the overall steady-state solution is reached.

2.5 Adaptive Split-Domain Harmonic Balance

The adaptive split-domain harmonic balance method minimizes the computational cost of the harmonic balance solution by automatically tailoring the number of Fourier frequencies included in the solution according to the flow at a given point, on a point-by-point (or for a finite volume discretization, cell-by-cell) basis. This is accomplished by means of a frequency augmentation approach. With this approach, the solution is begun with a user-specified minimum initial number of frequencies. As the solution develops, frequencies are added in fixed increments to individual grid points until a final frequency distribution and solution are obtained.




Figure 2.1 Time-sample Shifting that Occurs with Each Iteration of the Split-domain Solution Process

The frequency augmentation approach was chosen because it is simple and computationally efficient, especially for problems that are elliptic in time (e.g. subsonic flow problems). Computational efficiency stems from the fact that the low-frequency coefficients included in the early solutions tend to have low-frequency spatial variations (Fig. 2.2), and thus take the most time to converge to steady state. By beginning with a small number of Fourier frequencies, much of the work of converging the low-frequency coefficients is accomplished while solving a reduced problem. Subsequent addition of higher Fourier frequencies introduces mostly high-spatial-frequency errors which are quickly removed from the solution. In this respect, frequency augmentation is similar to a full multigrid convergence acceleration method, in which a solution is converged on a series of computational grids, beginning with a coarse grid and continuing on successively finer grids until a final solution is obtained.

To implement a frequency augmentation approach, it is necessary to identify which points in the computational grid need augmentation, and to properly handle the frequency transitions that occur when a cell and its neighbor are solved with different numbers of frequencies. Determining which cells need additional Fourier frequencies requires a means of measuring the quality of the solution, and some




Figure 2.2 Illustration of Typical Spatial Variation of a Low-frequency (top) and High-frequency (bottom) Fourier Coefficient

criteria against which that quality can be compared. It also requires a strategy that governs how often the criteria are applied. Handling frequency transitions requires the development of a robust multi-frequency approach for solving the time-domain equation, Eq. 2.23a.

2.5.1 Frequency Augmentation Criteria and Scheduling. The frequency augmentation approach requires a reliable indicator of solution fidelity at each computational cell. Since the final time-response of the flow is not known *a priori*, the indicator must rely only on the current (and possibly past) state of the solution. The indicator chosen for this research is the fraction of spectral energy contained in the highest computed Fourier Frequency, E_N , given by

$$E_N = \frac{|c_N|^2}{\sum_{n=0}^N |c_n|^2}. \quad (2.28)$$

The assumptions underlying the choice of E_N as an adaptation metric are that the majority of the spectral energy is contained in the low-frequency terms, and that the fraction of energy contained in the highest calculated frequency decreases as more terms are included in the approximating series. Thus the quality of the solution can

be gauged by the amount of energy in the highest frequency term. Physically, these assumptions require that, aside from localized discontinuities, the modeled flow field is smoothly varying throughout its period. In addition, if discontinuities do appear, they must not be impulse discontinuities, i.e., the discontinuity should be a step rather than a spike. These physical requirements are consistent with the flow fields of interest, and proved valid for all test cases examined.

The decision to augment frequencies at a point is made by comparing E_N to a threshold value, E_{thresh} . Because E_N tends to mirror the spatially oscillatory nature of the high-frequency coefficients, it is smoothed with an unweighted 5-point spatial average prior to thresholding. When the smoothed E_N exceeds the threshold, additional frequencies are incorporated into the solution at that point. Fourier coefficients for the new frequencies are initialized to zero. In the case of a vector equation, E_N is calculated for each of the variables expanded in a series, and the solution is augmented if any one of these exceeds E_{thresh} . Selection of E_{thresh} is based on experience and the desired solution fidelity. See Sections 4.4 and 5.4 for more on the impact of threshold value on adapted solutions.

Threshold-based augmentation was supplemented by two forms of non-threshold based augmentation. The first of these was fringe augmentation. The purpose of fringe augmentation was to increase the size of a threshold-augmented region. In some test cases with very rapid transitions between smooth and discontinuous flow, the location of the transitions changed as frequencies were added and the solution was refined. In those situations, it was sometimes necessary to augment a small fringe of cells adjacent to threshold-identified cells in order to allow the transition to shift in the direction of lower frequency content.

The second non-threshold-based augmentation was designed to enforce a minimum number of consecutive cells with the same frequency content. In some cases, threshold-based augmentation can result in small (1-2 cell) segments of the computational domain having a different frequency content than their neighbors. To avoid




Figure 2.3 Effect of Fringe Augmentation and Pixelation on a Threshold-based Frequency Distribution

this situation, a pixelation process was applied to each newly augmented frequency distribution. In this process, the grid was divided into contiguous, non-overlapping segments with a uniform size, or pixelation width. Within each segment, the frequency content was set to the maximum found within that segment. This guaranteed a minimum number of contiguous cells with the same frequency content, while also guaranteeing frequency content greater than or equal to that required to meet the augmentation threshold. The cumulative effect of both fringe augmentation and pixelation is illustrated in Fig. 2.3.

For the current research, both the augmentation fringe width and pixelation width were controlled by user input at run time. Each could be disabled when not needed by setting the fringe width to zero, and the pixelation width to one.

Once a cell was identified for augmentation, its frequency content was increased by a predetermined increment, chosen to minimize compute time. The majority of the run time required to solve the split-domain harmonic balance equations is composed of two components—the time associated with solving the time-domain equations (Eq. 2.23a), and the time spent performing the necessary Fourier transforms. Run

time associated with the time-domain equations increases linearly with N . There is no performance advantage favoring any particular increment, so long as the increment is relatively small. There can, however, be a significant difference in the run time of the Fourier transforms for different N , as discussed in Section 2.2. Numerical tests were conducted to determine values of N which gave the best overall performance. The results of these tests are documented in Appendix C.

Like the augmentation increment, the timing of frequency adaptation was based on run-time performance considerations. Frequency augmentation is most effective when it is based on solutions that are representative of the final solution. If adaptation is attempted too early in the solution process, unneeded frequencies could be added based on transient flow structures that are not present in the final solution. On the other hand, if the solution is allowed to develop too long before adaptation is attempted, the result could be wasted computational effort. The goal of adaptation scheduling is to identify the “right” times to adapt the solution.

A dual-trigger adaptation scheduling strategy is taken. The primary trigger is based on a modified L_2 norm of the change in the solution during one iteration. The L_2 norm, or residual, is defined as

$$R = \sqrt{\frac{\sum_{i=1}^{ni} (\Delta \hat{\xi}_i \cdot \widetilde{\Delta \hat{\xi}_i})}{ni}}. \quad (2.29)$$

The L_2 norm is a measure of the remaining error in the solution. In Eq. 2.29, ni is the number of cells in the computational grid, $\Delta \hat{\xi}$ is the change in the solution vector in one iteration, and the overset $\sim$ indicates complex conjugation. Adaptation is triggered when $\log_{10}(R)$ drops by a user-specified amount, indicating that a desired level of solution development has been reached. The secondary trigger is based on the number of iterations completed. The iteration-based scheduling serves as a backup to residual-based scheduling in the rare cases where solution convergence stalls and the specified residual drop is not achieved.

The residual drop and iteration count are measured relative to a reference residual and iteration number. The initial reference values are set after the first iteration of the solution. The reference values are then reset whenever the solution is adapted. Adaptation trigger values for the first and subsequent adaptations are different. The initial trigger values are set to allow time for the solution to develop from a poor initial guess. Subsequent trigger values are set to allow errors introduced by frequency augmentation to be removed from the solution, and to further refine the solution a small amount. The initial trigger values are typically much larger than the subsequent trigger values.

The adaptive frequency augmentation algorithm is summarized in Fig. 2.4.

2.5.2 Treatment of Frequency Transitions. Discretization of the spatial derivative in the time-domain portion of the split-domain harmonic balance equations (Eq. 2.23a) requires the addition and/or subtraction of the solution vector in a cell, $\bar{\xi}_i$, (or the corresponding fluxes, $\bar{\Phi}(\bar{\xi}_i)$) with the other solution vectors in its discretization stencil. This presents a problem when those cells have different maximum frequencies. Not only do the solution vectors have different numbers of elements, but those elements correspond to the state of the flow at different points in time; they have different sample rates. To solve this problem, the solution vectors must be resampled so that the sample rate is consistent across the discretization stencil (Fig. 2.5).

Two different resampling methods are employed in the present implementation. These methods include truncation/zero-padding in the frequency domain, and linear interpolation in the time domain. The primary means of resampling is truncation/zero-padding in the frequency domain. With this method, a vector is down-sampled by taking its Fourier transform, truncating the results, and transforming the truncated Fourier coefficient vector back to the time domain. Upsampling is achieved in a similar manner, except that instead of truncating the Fourier coefficient




Figure 2.4 Flowchart for the Adaptive Harmonic Balance Algorithm




Figure 2.5 Sample-rate Interpolation Required at a Frequency Transition from $N = 2$ to $N = 3$

vector, it is padded with new zero-valued high-frequency coefficients. This approach results in very smooth interpolation. When the results are used in a linear operation such as addition or subtraction, the approach is equivalent to performing the same linear operation directly in the frequency domain.

One drawback with frequency-domain truncation/zero padding is that the interpolated values are not bounded by the original data. When sample rates are very small, interpolation can result in non-physical values, as demonstrated in Fig. 2.6. When this occurs, linear interpolation in the time-domain is applied.

With the linear interpolation method, conservative variables are obtained at the required sample times by linearly interpolating the calculated values. Besides guaranteeing that the interpolated values are bounded by the computed values, this approach has the advantage that it can be applied entirely in the time domain, and does not require any additional Fourier transforms.

Despite the higher computational cost and potential low-sample-rate problems, frequency-domain truncation/zero-padding was selected as the primary resampling




Figure 2.6 Non-physical Interpolated Values Produced by Frequency-domain Zero-padding

method. This choice was based on the quality of the solution computed at a sample rate transition. The superiority of the method over linear interpolation is illustrated in Fig. 2.7. This figure contains plots of contours of constant Fourier coefficient magnitude for both types of interpolation at a transition from 7 frequencies (15 samples/period) to 16 frequencies (33 samples/period). Also included in the plots (dashed) are coefficient magnitude contours for a solution with a constant 16 frequencies. The frequency-domain truncation/zero padded solution is much smoother and more closely matches the constant-sample-rate results than the linearly-interpolated solution.

In the adaptive split-domain harmonic balance solver written for this research, a 3-stage Runge-Kutta time-integration scheme was used to advance the time-domain PDE, Eq. 2.23a (See Section 4.2). In order to achieve a smooth solution at frequency transitions, it was necessary to resample the transition boundaries at every integration stage. Attempts to time-lag the transitions by freezing the resampled values at the first integration stage resulted in discontinuities at the transition point, even when both sides of the transition had the same number of frequencies and no interpolation was required.




Figure 2.7 Comparison of Fourier Coefficient Magnitudes at a Transition from 7 Frequencies to 16 Frequencies, for Solutions Computed with Time-domain Linear Interpolation (left) and Frequency-domain Truncation/Zero-padding (right) Resampling Methods

As a consequence of the need to resample the transition boundaries at each integration stage, the previously uncoupled system of equations represented by Eq. 2.23a becomes effectively coupled through the resampling operation. This significantly increases the storage requirements of the solver. When the time-domain system of equations was uncoupled, the numerical integration scheme could be applied separately to each equation in the system, reducing the storage requirement for intermediate solutions to a single time sample. But when the equations are coupled, intermediate solutions for all samples must be stored. For this reason, it is important that the adaptive split-domain harmonic balance method be implemented with a numerical integration scheme with low storage requirements if that scheme is a multi-step scheme.

2.6 Multigrid and Adaptive Split-Domain Harmonic Balance

The Full Approximation Storage (FAS) multigrid method is a commonly used technique for accelerating the calculation of steady-state solutions to nonlinear conservation equations. The method works by transferring a partially converged solution to a coarser computational grid, where a coarse grid correction is calculated. This

correction is then applied to the solution on the fine grid. Because the coarse grid contains fewer grid points, calculation of the coarse grid correction is faster than calculation of the solution on the fine grid.

Implementation of FAS multigrid in an adaptive split-domain harmonic balance solver requires some special considerations. In the following discussion, the basic FAS multigrid scheme as found in (28) is presented, followed by the specifics of the adaptive split-domain implementation.

Theory of FAS Multigrid. Consider an equation of the form

$$L^h U^h = f^h, \quad (2.30)$$

defined on a grid with spacing h , where L^h is a nonlinear operator, U^h is the exact solution, and f^h is a forcing function. Let u^h be an approximate solution to Eq. 2.30, and $V^h = U^h - u^h$ be the error in the approximate solution. The goal of the FAS multigrid scheme is to quickly calculate an estimate of the solution error, V^h .

Substituting $U^h = V^h + u^h$ into Eq. 2.30 and subtracting $L^h u^h$ from both sides gives

$$L^h(V^h + u^h) - L^h u^h = f^h - L^h u^h. \quad (2.31)$$

If the terms in Eq. 2.31 are sufficiently smooth, the equation can be transferred or “restricted” onto a coarse grid with grid spacing $2h$ without much loss of accuracy. Given a coarse grid operator, L^{2h} , and a restriction operator, I_h^{2h} , the restricted equation is given by

$$L^{2h}(I_h^{2h} u^h + V^{2h}) - L^{2h}(I_h^{2h} u^h) = I_h^{2h}(f^h - L^h u^h). \quad (2.32)$$

The coarse grid operator L^{2h} may or may not be the same as the fine grid operator, L^h . V^{2h} is the approximate solution error on the coarse grid. Eq. 2.32 can be

rewritten to have the same form as Eq. 2.30, i.e.,

$$L^{2h}u^{2h} = f^{2h} \quad (2.33)$$

where

$$u^{2h} = I_h^{2h}u^h + V^{2h} \quad (2.34)$$

$$f^{2h} = I_h^{2h}(f^h - L^h u^h) + L^{2h}(I_h^{2h}u^h). \quad (2.35)$$

The FAS multigrid scheme solves for u^{2h} on the coarse grid, which, because it has half as many points as the fine grid, requires half as much work. In addition, since the coarse grid spacing is larger, in many cases a larger integration time step may be taken. Once a solution for u^{2h} is obtained, the error in the solution, V^{2h} is recovered and interpolated to the fine grid to form a coarse grid correction. The fine grid solution is updated by adding the coarse grid correction to the original approximate solution, as illustrated in Eq. 2.36.

$$(u^h)_{\text{new}} = (u^h)_{\text{old}} + I_{2h}^h(u^{2h} - I_h^{2h}u_{\text{old}}^h) \quad (2.36)$$

Because Eq. 2.33 has the same form as Eq. 2.30, the FAS multigrid algorithm is easily extended to more than two levels. Each coarse grid becomes the fine grid for a still coarser grid. In theory, this can continue until there is only one point in the interior of the coarsest grid.

A typical multigrid implementation begins by performing a small number of iterations on the finest grid to remove high spatial-frequency errors. This smoothed solution is restricted to the next coarsest grid, where the solution is again smoothed for a small number of iterations. The process of smooth and restrict is repeated until the coarsest grid is reached, at which point the solution is converged to steady state. Once steady state is reached, a coarse grid correction is made to the next finer

grid, and a small number of iterations are performed to remove errors introduced by prolongation. The process is repeated until the finest grid is reached, completing one multigrid cycle. This down-and-up pattern is referred to as a “V” cycle (Fig. 2.8). Other cycles with different restriction/prolongation patterns (e.g., the “W” cycle) are also sometimes used.




Figure 2.8 Illustration of a 4-Level Multigrid V Cycle

FAS Multigrid and Split-domain Harmonic Balance. In order to apply FAS Multigrid to the split-domain harmonic balance equations, they must be put in the form of Eq. 2.30. Let L_f be defined as the time-integration operator that advances the solution to the frequency-domain ODE (Eq. 2.23b) by one half time step. Similarly, let L_t be defined as the time-integration operator that advances the time-domain PDE by one full time step. The process of integrating the solution vector $\hat{\xi}$ from time level n to time level $n + 1$ with the split-domain scheme can then be written as

$$\hat{\xi}^{n+1} = L_f \mathcal{F} L_t \mathcal{F}^{-1} L_f \hat{\xi}^n. \quad (2.37)$$

Subtracting $\hat{\xi}^n$ from both sides and dividing by the integration time step, $\Delta\tau$, gives

$$\frac{\hat{\xi}^{n+1} - \hat{\xi}^n}{\Delta\tau} = \frac{1}{\Delta\tau} [L_f \mathcal{F} L_t \mathcal{F}^{-1} L_f - I] \hat{\xi}^n, \quad (2.38)$$

where I is the identity operator. The left-hand side of Eq. 2.38 is a first-order forward-difference approximation of the pseudo-time derivative $\frac{\partial \hat{\xi}}{\partial \tau}$ added to the steady-state harmonic balance equation to allow the use of a time-marching solution method. Removing the derivative results in an operator-notation expression for the steady-state split-domain harmonic balance equations that is in the desired form, i.e. $L\hat{\xi} = \hat{f}$, where $\hat{f} = 0$, and

$$L = \frac{1}{\Delta\tau}[L_f\mathcal{F}L_t\mathcal{F}^{-1}L_f - I]. \quad (2.39)$$

To calculate a coarse grid correction for the split-domain harmonic balance equations at time level n , the updated solution at time level $n + 1$, $\hat{\xi}^{n+1}$, must first be calculated. The change in the solution vector, $\hat{\xi}^{n+1} - \hat{\xi}^n$, is computed and divided by the pseudo-time step, $\Delta\tau$. The result, along with the solution at time level n , is restricted to the next coarsest grid and the coarse grid correction is calculated. The correction is added to $\hat{\xi}^n$. The previously calculated $\hat{\xi}^{n+1}$ is discarded.

FAS Multigrid and Frequency Augmentation. To implement adaptive frequency augmentation and FAS multigrid together, no major changes are required in either method. A strategy is required to govern what grid levels augmentation will occur on, and how frequency maps will be transferred from one grid level to another. The only additional requirement is that frequency transitions must be properly handled during restriction and prolongation.

In the current research, frequency augmentation was applied only on the finest grid. Frequency map consistency from the finest grid to the coarsest grid was maintained by setting the frequency map pixelation width to 2^{k-1} , where k is the depth of the multigrid cycle. This width guarantees that the two fine grid cells contributing to a coarse grid cell both have the same sample rate for all grid levels (Fig. 2.9). Thus every region of the grid is represented by the same frequency content on all grid levels. This approach eliminates frequency transitions during restriction, be-




Figure 2.9 Illustration of How a Pixelation Width of 2^{N-1} Preserves Spatial Frequency Content and Eliminates Frequency Transitions During Restriction

cause all frequency transitions occur between cells that contribute to different coarse grid cells. Frequency transitions are still encountered during prolongation, however. These transitions are handled by zero-padding the high-frequency Fourier coefficients of the cell with the smaller sample rate.

III. Application of Split-Domain Harmonic Balance to Burgers' Equation

3.1 Introduction

The purpose of the research documented in this chapter was to examine the behavior of the split-domain harmonic balance method when applied to problems with strong discontinuities moving over large regions of the grid. A useful model for this research was the 1-D inviscid Burgers' equation, a simplified form of Euler's equation that yields traveling discontinuities in the flow field for large amplitude periodic disturbance boundary conditions. The study included the application of the harmonic balance method to two families of unsteady boundary conditions – one based on a single-frequency sine wave and the second based on a simulated wake function. The amplitude and frequency of the boundary conditions were varied to generate test cases with a wide range of flow properties, from smooth and continuous to strongly discontinuous. The effect of number of Fourier frequencies included in the harmonic balance calculation of the solution to these different test cases was of particular interest. To provide a basis for comparison of accuracy, stability, and performance, harmonic balance solvers based on prior harmonic balance approaches were also implemented and tested.

3.2 Solver Implementation

Four harmonic balance CFD solvers were written to solve the one-dimensional inviscid Burgers equation, given by

$$\frac{\partial u}{\partial t} + \frac{1}{2} \frac{\partial u^2}{\partial x} = 0, \quad (3.1)$$

where $u(x, t)$ is the dependent variable, t is the temporal variable, and x is the spatial variable. The four solvers included a direct harmonic balance implementation

(Eq. 2.15), a time-domain implementation (Eq. 2.18), a frequency-domain implementation (Eq. 2.21), and a split-domain implementation (Eq. 2.23).

The direct, time-domain, and frequency-domain solvers were implemented with an explicit MacCormack discretization scheme, modified to incorporate source terms (24). MacCormack's scheme is a two-step solution scheme that is second order in both time and space. For a scalar conservation equation with source term, the modified MacCormack scheme is given by

$$\xi_i^* = \xi_i^n - \Delta t \left(\left[\frac{\hat{\Phi}(\xi_{i+1}^n) - \hat{\Phi}(\xi_i^n)}{2\Delta x} \right] + \mathbf{S}(\xi_i^n) \right) \quad (3.2a)$$

$$\xi_i^{n+1} = \frac{1}{2} \left(\xi_i^n + \xi_i^* - \Delta t \left(\left[\frac{\hat{\Phi}(\xi_i^*) - \hat{\Phi}(\xi_{i-1}^*)}{2\Delta x} \right] + \mathbf{S}(\xi_i^*) \right) \right). \quad (3.2b)$$

Discretizations of the direct, time-domain, and frequency-domain harmonic balance equations were obtained by substitution of the appropriate solution vectors, flux functions, and source terms into Eqs. 3.2a and 3.2b. (See Appendix A for a derivation of the harmonic balance Burgers' flux function and source term required for the direct solver.) For the frequency-domain formulation, substitution resulted in Eqs. 3.3a and 3.3b, and for the time-domain formulation, Eqs. 3.4a and 3.4b.

$$\hat{\xi}_i^* = \hat{\xi}_i^n - \Delta \tau \left(\mathcal{F} \left[\frac{(\bar{\xi}_{i+1}^n)^2 - (\bar{\xi}_i^n)^2}{4\Delta x} \right] + \hat{\mathbf{S}}(\hat{\xi}_i^n) \right) \quad (3.3a)$$

$$\hat{\xi}_i^{n+1} = \frac{1}{2} \left(\hat{\xi}_i^n + \hat{\xi}_i^* - \Delta \tau \left(\mathcal{F} \left[\frac{(\bar{\xi}_i^*)^2 - (\bar{\xi}_{i-1}^*)^2}{4\Delta x} \right] + \hat{\mathbf{S}}(\hat{\xi}_i^*) \right) \right) \quad (3.3b)$$

$$\bar{\xi}_i^* = \bar{\xi}_i^n - \Delta \tau \left(\frac{(\bar{\xi}_{i+1}^n)^2 - (\bar{\xi}_i^n)^2}{4\Delta x} + \mathcal{F}^{-1} \hat{\mathbf{S}}(\hat{\xi}_i^n) \right) \quad (3.4a)$$

$$\bar{\xi}_i^{n+1} = \frac{1}{2} \left(\bar{\xi}_i^n + \bar{\xi}_i^* - \Delta \tau \left(\frac{(\bar{\xi}_i^*)^2 - (\bar{\xi}_{i-1}^*)^2}{4\Delta x} + \mathcal{F}^{-1} \hat{\mathbf{S}}(\hat{\xi}_i^*) \right) \right) \quad (3.4b)$$

In these expressions, the squaring of a vector quantity denotes element-wise squaring. The quantities $\hat{\xi}$, $\bar{\xi}$, and $\hat{\mathbf{S}}$ are as defined in Section 2.2. The residual function in the time-domain harmonic balance equation (Eq. 2.21) has a different definition for each of the solution stages, and is defined as the expressions in square brackets in Eqs. 3.4a and 3.4b. Note that because these expressions contain only linear operations (vector subtraction and scalar division), the frequency domain solver and time-domain solver should be equivalent.

The split-domain solver was implemented with different discretization schemes for the time-domain PDE, Eq. 2.23a, and the frequency-domain ODE, Eq. 2.23b. The time-domain PDE was solved with an unmodified MacCormack discretization, given by

$$\bar{\xi}_i^* = \bar{\xi}_i^n - \frac{\Delta\tau}{4\Delta x} ((\bar{\xi}_{i+1}^n)^2 - (\bar{\xi}_i^n)^2) \quad (3.5a)$$

$$\bar{\xi}_i^{n+1} = \frac{1}{2} \left(\bar{\xi}_i^n + \bar{\xi}_i^* - \frac{\Delta\tau}{4\Delta x} \left((\bar{\xi}_i^*)^2 - (\bar{\xi}_{i-1}^*)^2 \right) \right). \quad (3.5b)$$

The frequency-domain ODE was solved with a three-stage Runge-Kutta numerical integrator

$$\begin{aligned} \hat{\xi}^* &= \hat{\xi}^n + 0.35 \frac{\Delta\tau}{2} \hat{\mathbf{S}}(\hat{\xi}^n) \\ \hat{\xi}^{**} &= \hat{\xi}^n + 0.6 \frac{\Delta\tau}{2} \hat{\mathbf{S}}(\hat{\xi}^*) \\ \hat{\xi}^{n+1/2} &= \hat{\xi}^n + \frac{\Delta\tau}{2} \hat{\mathbf{S}}(\hat{\xi}^{**}). \end{aligned} \quad (3.6)$$

Note that Eq. 3.6 advances the ODE by one-half pseudo-time step, as required by the symmetric Strang splitting approach (see Section 2.4).

All solvers required artificial dissipation to prevent oscillations near discontinuities in the solution. Dissipation was incorporated in the direct, frequency-domain, and split-domain formulations by adding second derivative smoothing (Eq. 3.7) of

the Fourier coefficients at the end of each iteration, according to

$$\hat{\xi}_{i, new} = \hat{\xi}_i + \alpha \frac{\hat{\xi}_{i-1} - 2\hat{\xi}_i + \hat{\xi}_{i+1}}{\Delta x^2}. \quad (3.7)$$

Dissipation was added to the time-domain solver by applying the same smoothing operator to the time-domain solution vector, $\bar{\xi}$. Due to the linearity of the Fourier transform and artificial dissipation operators, this added dissipation was mathematically equivalent to that of the other solvers. The dissipation parameter α was a small constant, of order 1.0e-6, that controlled the amount of applied damping.

Burgers' equation is hyperbolic in time. For positive values of the dependent variable, u , all flow information travels in the direction of increasing spatial coordinate. Boundary conditions are therefore implemented by fixing the inflow boundary values, and extrapolating the outflow boundary values from upstream. A 0th order extrapolation was implemented. For the direct, frequency-domain, and split-domain solvers, boundary conditions were applied in the frequency domain. Boundary conditions were applied in the time domain for the time-domain solver.

To accelerate convergence to steady state, local time stepping was employed. Using the definition of the Courant-Friedricks-Lewy (CFL) stability limit for the MacCormack discretization (29), the maximum time step at each point was determined according to

$$\Delta \tau_i = \text{CFL} \frac{\Delta x}{\max(\bar{\xi}_i)} \quad (3.8)$$

where CFL was a user-specified value less than or equal to 1.0, and $\max(\bar{\xi}_i)$ was the maximum element in the vector of time-sampled dependent variables at the i^{th} grid point.

3.3 Test Configuration

Solutions were obtained for two families of periodic inflow boundary conditions. The first family was constructed by varying the amplitude and frequency of a

sinusoidal oscillation about a mean value of 1.0, i.e.,

$$u(0, t) = 1.0 + a \sin(\omega t) \quad (3.9)$$

where a is the amplitude of the oscillation, and ω is the frequency of the oscillation in radians per second. The second family consisted of amplitude and frequency variations of a simulated wake function, defined by

$$u(0, t) = \begin{cases} 1.0 & 0 \leq t < \frac{3\pi}{2\omega} \\ 1.0 + a \sin^2(2\omega t) & \frac{3\pi}{2\omega} \leq t \leq \frac{2\pi}{\omega} \end{cases} \quad (3.10)$$

The basic forms of the input families are shown in Fig. 3.1.




Figure 3.1 Sinusoidal (left) and Wake Function (right) Boundary Condition Waveforms

For each of these families, combinations of three amplitudes and three frequencies were applied, for a total of nine inflow conditions. The three amplitudes were $a = 0.1, 0.3$, and 0.5 , while the frequencies, defined in terms of $f = \omega/2\pi$, were $f = 0.75, 1.5$, and 3.0 . Over a 2-unit grid, these amplitudes produce a wide range of flow behavior, from smooth and continuous to strongly shocked. The frequencies chosen are typical of the reduced frequencies found in axial compressor simulations (5, 6, 9).

Inflow boundary conditions for the time-domain solver were obtained by evaluating Eqs. 3.9 and 3.10 at $2N + 1$ points in time, given by

$$t = \frac{2\pi n}{\omega(2N + 1)} \quad 0 \leq n \leq 2N. \quad (3.11)$$

Inflow boundary conditions for the direct, frequency-domain, and split-domain solvers were obtained by computing Fourier series coefficients for N frequencies. For Eq. 3.9, the real Fourier series coefficients are given by

$$a_0 = 1.0 \quad (3.12a)$$

$$a_n = 0.0 \quad 1 \leq n \leq N \quad (3.12b)$$

$$b_n = \begin{cases} a & n = 1 \\ 0.0 & 1 < n \leq N. \end{cases} \quad (3.12c)$$

The real Fourier series coefficients for Eq. 3.10 are given by

$$a_0 = 1 + .125 a \quad (3.13a)$$

$$a_n = \begin{cases} \frac{8a}{n(n^2-16)\pi} \sin(\frac{n\pi}{2}) & 1 \leq n \leq N, n \neq 4 \\ -0.125 a & n = 4 \end{cases} \quad (3.13b)$$

$$b_n = \begin{cases} \frac{8a}{n(n^2-16)\pi} (\cos(n\pi) - \cos(\frac{n\pi}{2})) & 1 \leq n \leq N, n \neq 4 \\ 0.0 & n = 4. \end{cases} \quad (3.13c)$$

Complex Fourier series coefficients for the positive frequencies were obtained from the real Fourier series coefficient by means of the following relations

$$c_0 = \frac{1}{2} a_0 \quad (3.14a)$$

$$c_n = \frac{1}{2} (a_n - ib_n). \quad (3.14b)$$

3.4 Results

3.4.1 Accuracy. From the standpoint of accuracy, the four harmonic balance implementations were equivalent, producing solutions that were effectively identical. Representative solutions are shown for each input variation in Figs. 3.2 and 3.3. These figures compare the solution at $t = 0$ relative to the input period with an equivalent fully developed time-accurate calculation. The harmonic balance solutions were generated with 48 Fourier frequencies on a 501 point grid. The time-accurate solutions were obtained on the same grid using a validated MacCormack scheme with the same artificial dissipation used for the harmonic balance equation, Eq. 3.7.

As can be seen in Figs. 3.2 and 3.3, the input boundary conditions resulted in solutions ranging from smooth traveling waves to strong moving discontinuities. In all but two cases, the 48 frequency harmonic balance solutions were comparable to the time-accurate solutions. The two exceptions were the sine input cases with amplitudes $a = 0.5$ and 0.3 at the lowest disturbance frequency. These cases contain significant high-frequency oscillations in their solutions. It will be shown that these cases require additional terms in the approximating series.

Effect of Series Length. To determine the effect of series length on the accuracy of the harmonic balance method, solutions were generated for each input condition with series lengths ranging from 2 to 48 frequencies. A quantitative measure of the difference between each of these solutions and an equivalent time-accurate solution was obtained by calculating the difference root mean square (RMS) for 10 equally spaced temporal samples spanning one period of the disturbance. These 10 RMS differences were averaged to obtain a solution difference. The results are plotted in Figs. 3.4 and 3.5.

The results show that the differences between the harmonic balance and time-accurate solutions did not go to zero, but were asymptotic with respect to approximating series length. In each case where a good solution was obtained, there was




Figure 3.2 Comparison of Time-accurate (Solid Line) and 48-Frequency Harmonic Balance (Symbols) Solutions for the Sine Input for $t = 0$. Inset plots Show Dependent Variable Magnitude vs. Nondimensional Distance




Figure 3.3 Comparison of Time-accurate (Solid line) and 48-Frequency Harmonic Balance(Symbols) Solutions for the Wake Function Input for $t = 0$. Inset Plots Show Dependent Variable Magnitude vs. Nondimensional Distance




Figure 3.4 RMS Error Between Time-accurate and Harmonic Balance Solutions—Sine Input




Figure 3.5 RMS Error Between Time-accurate and Harmonic Balance Solutions—Wake Function Input




Figure 3.6 Comparison of 48-Frequency and 97-Frequency Solutions for $a = 0.5$, $f = 0.75$, on a 501 Point Grid

a series length corresponding to some Fourier frequency beyond which no improvement occurred. This frequency is hereafter referred to as the asymptotic frequency, and the associated solution is called the asymptotic solution. Solutions based on fewer frequencies are called sub-asymptotic, while those with more frequencies are called super-asymptotic. It is clear from Fig. 3.4 that the two cases with oscillatory solutions are sub-asymptotic, and thus required additional frequencies to minimize error. This was confirmed by generating a 97 frequency solution to the $a = 0.5$, $f = 0.75$ case, which is compared with the 48 frequency solution in Fig. 3.6. In the higher-frequency solution, most of the high frequency oscillations have been removed.

The fact that the differences between the harmonic balance and time-accurate solutions were asymptotic with respect to series length does not mean that the harmonic balance solutions did not continue to converge. The truncation error in the harmonic balance solution simply became insignificant compared to differences caused by other factors such as a slight difference in shock location.

Both disturbance amplitude and disturbance frequency influenced the asymptotic frequency. As amplitude increased, the asymptotic frequency also increased. This was attributed to the presence of stronger discontinuities in the larger amplitude solutions. In contrast, as the disturbance frequency increased, the asymptotic frequency usually decreased. This behavior is explained qualitatively.

If the shape and period of the time response of Burgers' equation at every point in the grid is fixed, the response then has a fixed frequency content with significant energy up to some frequency ω_{max} . In the harmonic balance solution, ω_{max} is expressed as some multiple of the fundamental frequency, say $n_{max}\omega$. Then $n_{max} = \omega_{max}/\omega$ is inversely proportional to ω , increasing as ω decreases.

True inverse proportionality was observed in some test cases (e.g., the sine input with frequency change from $f = 1.5$ to $f = 3.0$, all amplitudes), but for the majority of the cases the asymptotic frequency was less than that predicted by the simple model. This is consistent with the fact that a discontinuity in the time response is generally sharper for higher disturbance frequencies, and thus ω_{max} is not constant.

Effect of Grid Density. All of the harmonic balance solutions became dissipative to some degree as the computational grid was coarsened. The impact on solution quality depended on the nature of the flow field. For solutions with strong discontinuities the effect was relatively minor, and sometimes beneficial, while for smooth, small-amplitude solutions the effect introduced severe damping.

One example of a beneficial dissipative effect is shown in Fig. 3.7. This figure shows that while the coarse grid solution experiences some smearing of the shock, there is almost complete elimination of the non-physical oscillations present in the 49-frequency fine grid solution. Eliminating these oscillations on a fine grid would require the use of a much longer approximating Fourier series.




Figure 3.7 Smoothing Effect of Coarse Grid on Harmonic Balance Solution

An example of unfavorable effect of a coarse grid is illustrated in Fig. 3.8. This figure compares the 501-point time-accurate solution, the 101-point time-accurate solution, and the 101 point harmonic balance solution. The coarse-grid, time-accurate solution shows some degradation, primarily in the form of a phase-lag in the peaks of the solution. In contrast, the coarse grid damping effect caused considerable degradation in the harmonic balance solution. The harmonic balance method was more sensitive to grid density than the time-accurate method.

The effect of grid density in both cases was traced to the Fourier coefficients corresponding to the higher computed frequencies. Figure 3.9(a) shows the variation in magnitude of one high frequency ($n = 47$) coefficient for the $a = 0.5$, $f = 0.75$ sine input on the 501 point grid. The computed coefficient shows rapid oscillation in the spatial dimension. Figure 3.9(b) shows the same coefficient calculated on the 101 point grid. In this case, the coarse grid did not contain sufficient spatial resolution to capture the oscillations in magnitude, and the magnitude of the coefficient was under-predicted. The impact of poorly resolved high-frequency coefficients depended on the relative importance of those frequencies in the harmonic balance solution.




Figure 3.8 Severe Damping Effect of Coarse Grid on Harmonic Balance Solution

In the case shown in Fig. 3.7, most of the energy in the solution was contained in relatively low frequencies, and coarse-grid damping of the highest frequencies resulted in beneficial smoothing. For the case shown in Fig. 3.8, however, the damped frequencies comprised a significant part of the solution, and the overall accuracy was degraded. The generalized loss of higher-frequency information caused the harmonic balance method to be more sensitive to grid density than the time-accurate method.

The results suggest that while grid density is important for harmonic balance solutions, grid density is even more important for smooth solutions that require fewer terms in the approximating series. Problems that require many terms may require less grid resolution, partially offsetting the cost of the additional terms.

3.4.2 Stability. One of the primary motivations for the split-domain approach was improved high-frequency stability. The superiority of the split-domain approach was clearly evident in the test results. For low-frequency ($N < 10$) solutions, all of the harmonic balance solvers exhibited good stability, successfully computing solutions with $\text{CFL} > 0.95$. As the number of frequencies increased, the




Figure 3.9 Comparison of Spatial Variation of High-frequency Coefficient on (a) 501 Point Grid and (b) 101 Point Grid

maximum stable CFL of the direct, time-domain, and frequency-domain solvers was significantly reduced, while the maximum stable CFL of the split-domain solver remained unchanged. For example, for the $a = 0.3$, $f = 1.5$ test case on a 501 point grid, the split-domain solver was able to compute a 48-frequency solution with a CFL of 1.0, the maximum CFL for the MacCormack discretization scheme. For the same configuration, the direct, time-domain, and frequency-domain solvers became unstable when the CFL exceeded 0.42.

For super-asymptotic solutions, increased series length sometimes had a detrimental effect on numerical stability, even for the split-domain solver. All sub-asymptotic and asymptotic solutions, however, were successfully converged to a residual (L_2 norm) of $5.0\text{E-}8$ with a CFL of 0.95, independent of the number of frequencies used.

The need for a reduced CFL was not consistently related to the number of super-asymptotic frequencies included in the solution. Only test cases with a disturbance frequency of 3.0 required a reduced CFL to obtain 48-frequency solutions. In these cases, reductions of 40% to 60% in CFL were required. All other test cases converged with an un-reduced CFL, despite the fact that some of those cases had smaller asymptotic frequencies than some of the $f = 3.0$ cases.

3.4.3 Performance. The best overall run-time performance was obtained from the split-domain solver, especially for large N solutions. For a typical 49-frequency solution, the split-domain solver ran approximately 3 times faster (based on CPU time) than the frequency-domain and time-domain solvers, and 7 times faster than the direct solver. Factors contributing to the faster run time were the larger time step allowed by the higher CFL, and the fact that the split-domain scheme calculates fluxes in the time domain with only a single Fourier transform pair.

A third factor that heavily influenced the performance of all the multi-domain solvers, including the split-domain solver, was the choice of $N = 49$. As discussed in Section 2.2, the computational cost of an FFT, and thus of a multi-domain flux calculation, is highly dependent on the number of terms being transformed. This is dramatically illustrated by examining run times for a typical 48-frequency solution.

At 48 frequencies, the split-domain solver is only about 1.8 times faster than the direct solver, while the time-domain and frequency domain solvers are about 2.4 times slower than the direct solver. The drop in performance is explained by looking at the number of terms being transformed by the FFT for each N . For $N = 49$, the number of terms transformed is $2N + 1 = 99$, which is easily factored into $3^2 11^1$, and thus the proportional cost of the FFT is close to the best-case $N \log_2 N$. But when $N = 48$, the number of terms transformed is 97, which is prime and cannot be factored into products of small primes. In this case, the cost of the FFT is proportional to the worst-case N^2 . As a result, all of the multi-domain solvers take appreciably longer to run when N is reduced from 49 to 48. The split-domain solver manages to outperform the direct solver, but only because it is running at a higher CFL and computes a single transform pair per point per iteration.

3.5 Summary

Large amplitude, time-periodic solutions to Burgers' equation were computed with a split-domain harmonic balance solver, and compared to solutions computed

using prior harmonic balance approaches. The split-domain method produced solutions comparable to those of the prior methods, while successfully eliminating a stability restriction experienced by those methods when a large number of Fourier frequencies are included in the solution. This, combined with a reduction in the number of FFTs required to implement the split-domain method, resulted in significantly reduced run times.

Solutions for boundary conditions containing moving waves ranging from smooth disturbances to strong discontinuities were successfully computed. Comparison with conventional time-accurate calculations showed that the error in the harmonic balance solutions was asymptotic with respect to the number of frequencies included in the approximating solution. When the number of frequencies was equal to or greater than the asymptotic frequency, the harmonic balance solutions were comparable to the time-accurate solutions. Several factors were found to influence the asymptotic frequency, including disturbance frequency, the strength of the moving wave, and the computational grid density.

IV. Application of Adaptive Split-Domain Harmonic Balance to Euler's Equation

4.1 Introduction

This chapter documents an investigation of the adaptive harmonic balance method for CFD. The primary goal of the investigation was to establish that the energy-based frequency augmentation approach described in Section 2.5 reliably and effectively matches harmonic balance frequency content to local flow conditions, producing accurate solutions in less time than a non-adapted harmonic balance approach. The effects of grid density, augmentation threshold, pixelation width, and adaptation scheduling on the quality of adapted solutions and performance of the adaptation algorithm were examined. Finally, the compatibility of the adaptive split-domain harmonic balance approach with FAS multigrid convergence acceleration was confirmed.

To accomplish these objectives, adapted harmonic balance solutions were computed for a variety of supersonic and subsonic inviscid 1-D periodic flows governed by the 1-D Euler's equation. Results were compared with non-adapted harmonic balance solutions, and with solutions obtained using conventional time-accurate CFD techniques.

4.2 Solver Implementation

The one-dimensional Euler equations in strong conservation form are given by

$$\frac{\partial \mathbf{Q}}{\partial t} + \frac{\partial \mathbf{F}}{\partial x} = 0 \quad (4.1)$$

where

$$\mathbf{Q} = \begin{bmatrix} \rho \\ \rho u \\ E_t \end{bmatrix}, \quad \mathbf{F} = \begin{bmatrix} \rho u \\ \rho u^2 + p \\ (E_t + p)u \end{bmatrix} \quad (4.2)$$

and ρ , u , and p are the density, velocity, and pressure, respectively. Total energy, E_t , is defined as $\rho(e + \frac{1}{2}u^2)$, where e is specific internal energy. All quantities are nondimensional. Together with the perfect gas relation and an assumption of a constant ratio of specific heats, $\gamma = 1.4$, Eq. 4.1 represents a closed system of equations.

Equation 4.1 provides the form of the time-domain PDE in a split-domain harmonic balance implementation of the 1-D Euler equation. In the harmonic balance implementation, the dependent vector $\mathbf{Q}$ is replaced by a new vector $\hat{\mathbf{Q}}$ which is composed of the $2N + 1$ instances of the vector $\mathbf{Q}$ sampled at times $t = (0, \Delta t, 2\Delta t, \dots, 2N\Delta t)$, with $\Delta t = \frac{2\pi}{\omega(2N+1)}$, the period of oscillation divided by the number of samples, $2N + 1$. However, because the equations corresponding to each sample are independent of the other samples, they can be treated independently. For this reason, the following discussion is based on a single sample.

A cell-centered finite-volume solver (30) was written to solve the time-domain PDE. The finite-volume solver is based on a discretization of the integral form of Eq. 4.1

$$\frac{\partial}{\partial \tau} \int_V \mathbf{Q} dV + \int_S \mathbf{F} \cdot \hat{\mathbf{n}} dS = 0. \quad (4.3)$$

In Eq. 4.3, dV is a differential volume element, $\hat{\mathbf{n}}$ is a unit vector normal to the control volume surface, and dS is a differential surface element. Integration is performed over the interior (first term) and surface (second term) of the control volume. The physical time variable t has been replaced by a pseudo-time variable τ . If the control volume is taken to be a single grid cell, then for a one-dimensional computational

grid, Eq. 4.3 reduces to

$$V \frac{\partial \mathbf{Q}_i}{\partial \tau} + \mathbf{F}_i^r - \mathbf{F}_i^l = 0 \quad (4.4)$$

for each cell. The one-dimensional cell volume, V , is just the cell size Δx , and $\mathbf{Q}_i$ is the average, or cell center, value of the conserved variables inside the cell. The fluxes at the left and right faces, $\mathbf{F}^l$ and $\mathbf{F}^r$, are constructed by averaging the fluxes evaluated at the cell centers on either side of the face. For a grid with uniform cell size, this scheme is second-order accurate in space and is equivalent to a second-order central difference scheme.

To avoid oscillations at discontinuities, second and fourth order modified Jameson-Schmidt-Turkel (JST) artificial dissipation (30, 31) was implemented. JST artificial dissipation adds an additional term, $\mathbf{D}$, of the form

$$\mathbf{D} = (\mathbf{D}^2 - \mathbf{D}^4)\mathbf{Q} \quad (4.5)$$

to the left hand side of Eq. 4.4. $\mathbf{D}^2$ and $\mathbf{D}^4$ are second and fourth order difference operators defined as

$$\mathbf{D}^2 \mathbf{Q}_i \equiv \bar{\nabla}(\lambda_{i+1/2} \varepsilon_{i+1/2}^{(2)}) \underline{\Delta} \mathbf{Q}_i \quad (4.6)$$

$$\mathbf{D}^4 \mathbf{Q}_i \equiv (\bar{\nabla} \underline{\Delta})(\lambda_i \varepsilon_i^{(4)} \bar{\nabla} \underline{\Delta}) \mathbf{Q}_i \quad (4.7)$$

where

$$\varepsilon_{i+1/2}^{(2)} = K^{(2)} \max(\nu_i, \nu_{i+1}) \quad (4.8)$$

$$\varepsilon_i^{(4)} = \max(0, K^{(2)} - K^{(4)} \nu_i) \quad (4.9)$$

$$\nu_i = \left| \frac{p_{i+1} - 2p_i + p_{i-1}}{p_{i+1} + 2p_i + p_{i-1}} \right|. \quad (4.10)$$

In the above equations, $\underline{\Delta}$ and $\bar{\nabla}$ are first-order forward and backward difference operators, $\lambda_i = |u_i| + a_i$ is the maximum eigenvalue of the flux Jacobian matrix

$\frac{\partial \mathbf{F}_i}{\partial \mathbf{Q}_i}$, $\lambda_{i+1/2} = 1/2(\lambda_i + \lambda_{i+1})$, and a_i is the local speed of sound. The parameters $K^{(2)}$ and $K^{(4)}$ control the amount of dissipation applied, and are set by the user at run time. For the test problems described below, typical values of $K^{(2)}$ ranged from 0.4 to 0.55, while values of $K^{(4)}$ ranged from 0.004 to 0.04.

Experience with the code showed that solution quality and convergence properties were sometimes improved by including additional second order dissipation. The additional dissipation was included by means of a modified $\mathbf{D}^2$ operator, given by

$$\mathbf{D}^2 \mathbf{Q}_i \equiv \overline{\nabla}(K^{(2g)} + \lambda_{i+1/2} \varepsilon_{i+1/2}^{(2)}) \underline{\Delta} \mathbf{Q}_i. \quad (4.11)$$

The new parameter $K^{(2g)}$ controls the amount of additional second order dissipation. For positive $K^{(2g)}$, second order dissipation in the form of second difference smoothing is applied throughout the grid. If $K^{(2g)}$ is zero, the original modified JST dissipation is recovered. Typical values of $K^{(2g)}$ used in this study ranged from 0.0 for supersonic flows to 0.001 for subsonic flows.

Including the artificial dissipation term and replacing the cell volume with the cell size, Δx , the semi-discretized one-dimensional finite-volume Euler's equation becomes

$$\frac{\partial \mathbf{Q}_i}{\partial \tau} = -\frac{1}{\Delta x}(\mathbf{F}_i^r - \mathbf{F}_i^l + \mathbf{D}_i). \quad (4.12)$$

A three-stage Runge-Kutta (RK) scheme (Eq. 4.13) with good high-frequency smoothing properties and minimal storage requirements (32) was used to advance the solution in pseudo-time.

$$\begin{aligned} \mathbf{Q}^* &= \mathbf{Q}^n + 0.35 \Delta \tau R(\mathbf{Q}^n) \\ \mathbf{Q}^{**} &= \mathbf{Q}^n + 0.6 \Delta \tau R(\mathbf{Q}^*) \\ \mathbf{Q}^{n+1} &= \mathbf{Q}^n + \Delta \tau R(\mathbf{Q}^{**}) \end{aligned} \quad (4.13)$$

In Eq. 4.13, $R(\mathbf{Q})$ is the right hand side of Eq. 4.12 evaluated at the indicated time levels. Artificial dissipation terms were evaluated once during the first RK stage, and frozen for the remaining stages.

Implementation of the split-domain harmonic balance approach also required the solution of the frequency-domain ODE, Eq. 2.23b. Two solution approaches were tested—an exact integration approach, given by Eq. 2.25, and a numerical approach based on the Runge-Kutta integration scheme applied to the time-domain PDE. It was found that the exact approach provided no significant improvement in solution quality or stability compared to numerical integration, and sometimes slowed or stalled convergence to steady state. For this reason, the numerical approach was employed for the results documented below.

Boundary conditions were enforced by setting a single ghost cell outside the boundaries of the computational domain (see Appendix D). Because only one ghost cell was maintained, the fourth-order dissipation term, $\mathbf{D}^4 \mathbf{Q}$, was set to zero in the cells adjacent to the boundaries. All boundary conditions were calculated and applied in the time-domain.

Local time stepping was employed to accelerate convergence to steady state. The maximum stable time step allowed in the i^{th} cell was determined as follows:

$$\Delta\tau_{max} = \text{CFL} \min_{samples} \left(\frac{\Delta x_i}{\lambda_i} \right). \quad (4.14)$$

Local time steps were calculated every iteration in conjunction with the artificial dissipation term, and used immediately. This meant that the time step used for the first frequency-domain ODE integration (step one of the split-domain iteration, Section 2.4) was different from that used for steps three and five, the time-domain and second frequency-domain integration. This did not affect the steady-state solution, however, because the change in time step from one iteration to another became negligible as the solution approached steady state.

The maximum stable CFL for the solver was 1.7. In a number of test cases with large sample rates, a reduced CFL was required. A simple variable CFL function was implemented that linearly varied the CFL from a user-specified maximum at the lowest sample rate to a user-specified minimum at the highest sample rate. During the early stages of the solution when sample rates were low, a large CFL was applied. As the solution progressed and frequencies were added to the solution, the CFL was reduced to maintain a stability solution. The linear scaling approach provided a conservative but robust scaling of CFL with sample rate.

In addition to local time stepping, full approximation storage (FAS) multigrid convergence acceleration (28, 33) was implemented as described in Section 2.6. Two different restriction operators were used to transfer the fine-grid solution to the next coarsest grid (33). Split-domain residuals (Eq. 2.39) were transferred using a conservative transfer operator defined by

$$I_{k+1}^k R_{k+1} = \frac{1}{V_k} \sum (V_{k+1} R_{k+1}), \quad (4.15)$$

where the subscripts k and $k + 1$ identify coarse and fine grid values respectively, and V is cell volume. Summation occurs over the two fine grid cells making up each coarse grid cell. A volume-weighted transfer operator, defined by

$$I_{k+1}^k \hat{\mathbf{Q}}_{k+1} = \frac{\sum (V_{k+1} \hat{\mathbf{Q}}_{k+1})}{\sum (V_{k+1})} \quad (4.16)$$

was used to transfer the solution values themselves. Immediately after restriction, initial local time steps for the coarse grid were calculated according to Eq. 4.14. Once a coarse grid correction was computed, it was transferred to the next finest grid using linear interpolation as shown in Fig. 4.1.




Figure 4.1 Linear Interpolation from Coarse Grid to Fine Grid

4.3 Test Configurations

Adaptive frequency augmentation was applied to several one-dimensional test cases. The first set of cases consisted of a family of flow fields with sinusoidally varying supersonic inflow. For these cases, the inflow Mach number was varied according to

$$M = 2.0 + a \sin(\omega t) \quad (4.17)$$

while maintaining a constant nondimensional density and static pressure of 1.0 and 0.17857, respectively. The disturbance frequency, ω was chosen so that approximately two complete disturbance cycles would occur each flow-through period on a unit grid. The magnitude of the sinusoidal disturbance was varied to achieve different flow characteristics.

Pressure distributions at several snapshots in time for a flow field with $a = 0.25$ are shown in Fig. 4.2. These solutions were obtained from a fully-developed time-accurate Roe solver calculation. Because the supersonic flow was hyperbolic in time with all flow information traveling in the downstream direction, the fully-developed solution was achieved very quickly. The snapshots plotted in Fig. 4.2 were taken after $5 \frac{1}{2}$ flow-through periods.




Figure 4.2 Static Pressure for Supersonic Test Configuration SS3, $M = 2.0 \pm 0.25$, Generated with Time-accurate Roe Solver on 1000 Cell Grid




Figure 4.3 Flow Interactions for Simulated Oscillating Piston in an Open Tube

The last test case consisted of a purely subsonic flow field loosely approximating the flow in an open tube with an oscillating piston at one end. The piston action was simulated by imposing a sinusoidal velocity $u(t) = a \sin(\omega t)$ at a fixed boundary while maintaining a zero pressure gradient and constant total enthalpy. The open end of the tube was modeled with a characteristic-variable far-field boundary condition (34). Details of the boundary condition implementation are given in Appendix D.

While these boundary conditions do not exactly model a physical system, they do result in a complex, periodic, subsonic flow field containing alternating right-running shock and expansion waves interacting with reflections from the open end of the tube (Fig. 4.3). Initial flow conditions of $u = 0$, $\rho = 1$, and $p = 0.7142857$ were assumed. By varying the length of the tube and the magnitude and frequency of the imposed sinusoidal piston velocity, the flow field was tuned so that a stationary periodic flow field with moderately strong features was produced. Tuning was accomplished by computing a large number of piston cycles with a validated time-accurate CFD solver, starting from a zero-velocity initial solution. The time history of the mass in the tube was used to determine when a stationary solution was achieved. The chosen configuration achieved fully developed flow in approximately 150 piston cycles, as shown in Fig. 4.4.

Figure 4.5 contains pressure plots for the fully developed subsonic test flow field at several points in time, computed during the 199th piston cycle of the time-accurate calculation.




Figure 4.4 Time-accurate Development of the Stationary Flow Field for the Piston-In-Tube Configuration

All harmonic balance solutions for both the supersonic and subsonic test cases were converged to an overall residual (Eq. 2.29), of $1.0\text{e-}9$, or approximately 6.5 orders of magnitude. For comparison purposes, time-accurate solutions such as those discussed above were also computed. A summary of case names and defining parameters of all the test cases is contained in Table 4.1.

Table 4.1 Summary of Test Configuration Parameters

Config	Unsteady Amp	ω	Grid Length
SS1	0.05	4π	2.0
SS2	0.1	4π	1.0
SS3	0.25	4π	1.0
SS4	0.5	4π	1.0
SU1	0.1π	5π	2.0




Figure 4.5 Static Pressure for Subsonic Virtual Piston Test Configuration SU1, Generated with Time-accurate Roe Solver on 1000 Cell Grid

4.4 *Tuning the Adaptive Solver*

The following paragraphs discuss several user-definable parameters that control frequency augmentation and their impact on solution accuracy and solver performance. For a detailed discussion of the adaptive algorithm and these parameters, see Section 2.5.

4.4.1 Selection of Augmentation Threshold. The key parameter of the adaptive harmonic balance method is the augmentation threshold, E_{thresh} . This threshold value indirectly controls the sample rate distribution, and thus the run time of a calculation and the accuracy of the resulting solution.

To study the impact of E_{thresh} on frequency content, accuracy, and run time, solutions were generated for both supersonic and subsonic configurations with thresholds ranging from $5.0\text{e-}4$ to $1.0\text{e-}10$. Examples of the frequency content and solution quality resulting from different augmentation thresholds are shown in Figs. 4.6–4.9. These figures show frequency and pressure distributions for configurations SS3 and SU1, respectively, for $E_{thresh} = 1.0\text{e-}4$, $1.0\text{e-}7$, and $1.0\text{e-}10$. The pressure distributions were reconstructed from the computed Fourier coefficients at time $t = 0$ relative to the disturbance period. It is important to note that the reconstruction could just as easily have been generated for any time with equal fidelity.

Figure 4.6 shows adapted frequency distributions for configuration SS3. From this plot, one can determine the number of Fourier frequencies included at each cell in the computational grid for each of the three augmentation thresholds. As seen in the figure, there was a significant increase and refinement of the frequency distribution with the change from $1.0\text{e-}4$ to $1.0\text{e-}7$. The increased frequency content was accompanied by a significant improvement in solution quality as shown in Fig. 4.7. However, further decreasing the threshold to $1.0\text{e-}10$ resulted in no perceptible change to the reconstructed pressures, despite another large increase in frequency content.




Figure 4.6 Adapted Frequency Distributions for Three Different Augmentation Thresholds, Case SS3, 500 cells




Figure 4.7 Nondimensional Static Pressure for Adapted Solutions at Three Different Augmentation Thresholds, Case SS3, 500 Cells




Figure 4.8 Adapted Frequency Distributions for Three Different Augmentation Thresholds, Case SU1, 500 Cells




Figure 4.9 Nondimensional Static Pressure for Adapted Solutions at Three Different Augmentation Thresholds, Case SU1, 500 Cells

The adapted frequency content of the subsonic case also increased significantly with each decrease in E_{thresh} (Fig. 4.8). Unlike the supersonic case, however, the subsonic pressure distributions for thresholds of 1.0e-4 and 1.0e-7 are nearly identical (Fig. 4.9). The increased frequency content resulting from the decrease in threshold did not improve the solution. A further threshold decrease to 1.0e-10 resulted in a reduction in solution quality, as high-frequency oscillations appeared upstream of the strongest shocks.

The results for both the subsonic and supersonic cases are consistent with the Burgers' equation results presented in Section 3.4. There it was found that including more frequencies in the harmonic balance solution did not always improve the solution; there was an asymptotic frequency beyond which no improvement occurred. The same should be true for augmentation threshold, since it indirectly controls frequency content. For the current test configurations, solution quality did not improve as E_{thresh} was decreased below 5.0e-8 (supersonic) and 1.0e-4 (subsonic). In Section 3.4.2, it was also found that in some cases, a reduced CFL was required to obtain a stable solution incorporating frequencies above the asymptotic frequency. This was the case for the subsonic configuration with thresholds below 1.0e-6 (Table 4.2). Evidence suggests that the low CFL required to obtain a solution with $E_{thresh} = 1.0e-10$ contributed to the appearance of high-frequency oscillations in the subsonic solution (see Appendix B).

Table 4.2 Maximum Fourier Frequency and Stable CFL for Configuration SU1, 500 Cells, for Decreases in Augmentation Threshold

Threshold	Max N	Max CFL
1.0e-6	40	1.7
1.0e-7	45	1.5
1.0e-8	58	1.1
1.0e-9	67 ^a	1.0
1.0e-10	82 ^a	0.7

^a Maximum allowed for run

The effect of E_{thresh} on problem size (given by the average frequency content of the solution), and run time is shown in Figs. 4.10 and 4.11. Figure 4.10 shows the average number of Fourier frequencies included in each adapted solution. For both the supersonic and subsonic configurations, the average frequency content was inversely proportional to the $\log_{10}(E_{thresh})$. The size of the supersonic adapted harmonic balance problem grew approximately 8 times from the largest threshold to the smallest, while the size of the subsonic problem grew approximately 3.5 times.

The growth in run time with decreasing E_{thresh} is shown in Fig. 4.11. Run time for the supersonic case grew at approximately the same rate as the problem size, increasing 6.7 times from the largest to smallest thresholds. Down to a threshold of $5e-8$, run time for the subsonic case also grew at the same rate as problem size. For smaller thresholds, however, the subsonic run time grew at approximately twice the rate of problem size. This was due to the reduced CFL required at these low thresholds.




Figure 4.10 Change in Average Frequency Content with Changing Augmentation Threshold for Configurations SS3 and SU1 on 500 Cell Grid




Figure 4.11 Change in Relative Compute Time with Changing Augmentation Threshold for Configurations SS3 and SU1 on 500 Cell Grid

These results suggest that a reasonable lower limit for the augmentation threshold would be $5.0\text{e-}8$. Adapting to a smaller threshold increased run times, but resulted in no improvement in solution quality. Determining a suitable upper limit for the augmentation threshold is less clear. Based on these results, a conservative upper bound for a high-fidelity solution would be on the order of $5.0\text{e-}7$. A larger threshold may be acceptable for some problems, however, as was seen with configuration SU1, where a threshold of $1.0\text{e-}4$ produced comparable results (c.f., Fig. 4.9). It is likely that a maximum acceptable value will vary widely depending on the flow being modeled, the grid used, and the desired solution quality. If a lower-fidelity solution is acceptable (e.g., only the time-average solution is of interest) then a much higher threshold would suffice.

4.4.2 Fringe Width. None of the test cases required fringe augmentation. Therefore the fringe augmentation width was set to zero.

4.4.3 Pixelation Width. Solutions for both subsonic and supersonic test cases were insensitive to the adaptation pixelation width. Widths as small as 1 (no pixelation) and as large as 25 were applied, with no significant change in the solution. For the subsonic case, solutions with larger pixelation widths converged slightly faster than those with smaller widths (4% faster for a width of 25, vice a width of 1); but for the supersonic case no consistent advantage was observed.

Based on these results, there is no requirement for frequency pixelation. A slight performance improvement might be realized by including pixelation, but any advantage will be case dependent. Other factors, such as multigrid implementation, may increase the importance of pixelation, however (see Section 2.6).

4.4.4 Adaptation Trigger. Solutions were computed for various values of the primary residual-based adaptation triggers, labeled κ_1 (initial) and κ_2 (subsequent). To ensure that augmentation was based on these triggers, the iteration-based triggers were disabled.

The adaptation triggers were found to have little impact on the final solution. They did, however, have a significant impact on run times. The behavior for a given set of triggers was quite different for supersonic and subsonic solutions. For the supersonic configurations, tests showed that it was best to adapt after only a small reduction in residual. The hyperbolic nature of the supersonic flow fields meant that any disturbance introduced in the flow had to propagate out of the grid before the solution converged. Thus, there was no advantage in converging a solution prior to adding new frequencies. The optimum trigger values for the supersonic cases were found to be $\kappa_1 = 0.25$ and $\kappa_2 = 0$. In comparison, a solution with $\kappa_1 = 2.0$ and $\kappa_2 = 0.1$ took 50% longer to compute.

The opposite behavior was observed for the subsonic case. Here, the elliptic nature of the problem meant that converging the low-frequency solutions before adapting provided a distinct performance advantage. The best run time for the

subsonic case were achieved with $\kappa_1 = 2.5$ and $\kappa_2 = 0.1$. The longest run time (approximately 55% longer) was achieved with $\kappa_1 = 0.5$ and $\kappa_2 = 0.1$. Attempts to run with $\kappa_1 = 0.25$ resulted in unstable solutions.

One subsonic test, conducted with $\kappa_1 = 2.75$, never adapted. The solution residual at the initial frequency content ($N = 3$) plateaued after a drop of approximately 2.55 orders of magnitude, and failed to achieve the specified drop. Had the iteration-based triggers been enabled, they would have ensured that the solution was eventually adapted.

4.5 Results

4.5.1 Accuracy. Figures 4.12–4.15 compare results for both adapted and non-adapted harmonic balance solutions and a conventional time-accurate solution. Figs. 4.12 and 4.13 contain pressure and coefficient magnitudes for configuration SS4, while Figs. 4.14 and 4.15 contain similar information for configuration SU1.

The pressure distributions shown for configuration SS4 (Fig. 4.12) are typical of all the supersonic configurations. The reconstructed pressures for the adapted and non-adapted solutions are equivalent, and compare favorably with the time-accurate calculation. Both include slight second order Gibbs effects at the shocks that would be removed with an increase in artificial dissipation. Examining the magnitudes of the momentum term Fourier coefficients (Fig. 4.13) confirms that, except for some small differences in the highest frequencies and near frequency transitions, the adapted and non-adapted results are essentially identical. This suggests that the omitted frequencies in the augmented solution have no significant impact on accuracy.

The reconstructed pressure distributions for the adapted and non-adapted harmonic balance solutions for configuration SU1 (Figs. 4.14) are comparable. There are slight differences in the two solutions, however, as can be seen by examining the coefficient magnitude contours in Fig. 4.15. Unlike the supersonic case, there are small




Figure 4.12 Computed Nondimensional Static Pressure for Adapted (Augmentation Threshold $7.0\text{e-}8$) and Non-adapted, 49-Frequency Harmonic Balance Solutions, Configuration SS4, 601 Cells, Compared with 1000 Cell Time-accurate Solution




Figure 4.13 Fourier Coefficient Magnitudes (Momentum Component) for Adapted (Augmentation Threshold $7.0\text{e-}8$) and Non-adapted, 49-Frequency Harmonic Balance Solutions, Configuration SS4, 601 Cells




Figure 4.14 Computed Nondimensional Static Pressure for Adapted (Augmentation Threshold $7.0\text{e-}8$, Variable CFL) and Non-adapted, 58-Frequency Harmonic Balance Solutions, Configuration SU1, 601 Cells, Compared with 1000 Cell Time-accurate Solution




Figure 4.15 Fourier Coefficient Magnitudes (Momentum Component) for Adapted (Augmentation Threshold $7.0\text{e-}8$, Variable CFL) and Non-adapted, 58-Frequency Harmonic Balance Solutions, Configuration SU1, 601 Grid Cells

but noticeable differences between the adapted and non-adapted harmonic balance Fourier coefficients. These differences result from the use of a variable CFL in the adapted solution. The steady-state solution computed by the split-domain harmonic balance method has a small dependence on time step size (see Appendix B). Application of CFL scaling resulted in larger time steps in some cells, and thus slightly changed the steady-state solution. The two solutions would be made nearly identical by freezing the CFL of the adapted solution at the same value used for the non-adapted solution, but this would increase run time and result in only minute changes to the reconstructed pressures.

Both adapted and non-adapted harmonic balance solutions for configuration SU1 show some discrepancies when compared to the time-accurate solution, particularly in the trough regions. These discrepancies were traced to the boundary condition at the open (right) end of the grid. A comparison of the computed pressure at that boundary for the harmonic balance and time accurate solutions is shown in Fig. 4.16. The pressure recovery immediately following the period of negative velocity (inflow) is much sharper for the time-accurate solution than for the harmonic balance solution. The steady-state solutions calculated for time samples in the recovery region have a shock wave located at the exit boundary (Fig. 4.17). The far field boundary condition applied at that boundary was not designed for such an extreme gradient; the pressures at the boundary were over-predicted. The same boundary condition was used in the time-accurate code, but in that case, the shock at the boundary was moving, and so its effect on the boundary condition was less severe. The end result was a delay and attenuation of the left-running expansion wave reflected from the right boundary.

4.5.2 Grid Density. The maximum significant Fourier frequency for any harmonic balance solution is highly dependent on grid density. High frequency coefficients vary rapidly in the spatial dimensions, and require a fine grid for accurate resolution. For a grid too coarse to support a given Fourier frequency, the coeffi-




Figure 4.16 Computed Pressure for Time-accurate and Harmonic Balance Solutions at the Open End of the Tube




Figure 4.17 Steady-state Harmonic Balance Pressure Distribution at Time Sample Just after Period of Negative Velocity

cient magnitudes are damped, and less energy is contained in that frequency. This establishes a natural ceiling for frequency augmentation, which leads to an observed benefit of the adaptive harmonic balance approach – it automatically matches the included frequency content to the computational grid. The variation in adapted frequency distribution with grid density is illustrated in Figs. 4.18 and 4.19, which show the final frequency distributions at several grid densities for configurations SS3 and SU1, respectively.




Figure 4.18 Variation in Adapted Frequency Distributions with Changing Grid Density for Test Configuration SS3, Augmentation Threshold $5.0e-8$




Figure 4.19 Variation in Adapted Frequency Distributions With Changing Grid Density for Test Configuration SU1, Augmentation Threshold $5.0e-8$

4.5.3 Performance. To gauge the performance benefit of the adaptive harmonic balance approach, two performance metrics were examined. The first metric measured problem size reduction, as measured by the difference between the maximum and average frequency content of an adapted solution. The second metric was the run time required to converge a solution to steady state.

For each of the test configurations SS2, SS3, SS4, and SU1, an adapted solution was calculated with an augmentation threshold of $7.0\text{e-}8$, optimal triggers, and a pixelation width of 7. No multigrid acceleration was used. A non-adapted solution was then calculated, again with no multigrid, based on the highest frequency content in the adapted solution. Both adapted and non-adapted solutions for configurations SS2, SS3, and SS4 were computed with a fixed CFL of 1.7. For case SU1, the adapted solution was computed with a CFL that scaled from 1.7 at the lowest frequency to 1.35 at the highest frequency, while the non-adapted solution required a fixed CFL of 1.3.

The reduction of frequency content and run time for the adapted solution is shown in Fig. 4.20. In all cases, adaptation resulted in reduced run times, despite relatively small reductions in frequency content. The subsonic case showed the most improvement, with more than a 50% reduction in run time compared to the non-adapted case. Much of this reduction was due to the variable CFL. The higher average CFL, coupled with the efficiency gained by solving the large-scale flow structures with fewer frequencies, resulted in a run time reduction 20% larger than the problem size reduction.

Since the adapted and non-adapted supersonic cases were computed with the same CFL, the adapted solutions had no time-step advantage. Due to the overhead of the adaptive approach, the number of iterations required to converge the adapted solutions increased relative to the non-adapted solutions by as much as 15%. Despite this increase in iterations, the reduced average frequency content still resulted in a small but significant reduction in run time.

4.6 *Multigrid*

To demonstrate that the adaptive harmonic balance technique is compatible with multigrid acceleration, FAS multigrid was applied to one test case. A low-disturbance-magnitude supersonic case (configuration SS1) was solved on a 1025




Figure 4.20 Reduction in Average Frequency Content and Run Time for Adapted vs Non-adapted Harmonic Balance Solutions

point grid using a 5-level FAS Multigrid V cycle. This configuration resulted in a solution that is relatively smooth over most of the grid. The adapted frequency distribution contained just 2 frequencies at the inlet, and gradually increased in 1-frequency increments to 7 frequencies at the exit. A representative pressure plot for this configuration computed with and without multigrid convergence acceleration is shown in Fig. 4.21.

One minor modification to the augmentation algorithm was made to simplify implementation of the multigrid scheme. The primary residual-based adaptation trigger was replaced with a trigger based on the number of multigrid cycles completed. For this configuration, adapting every 3 multigrid cycles was found to be effective.

Run times for both adapted and non-adapted solutions, with and without multigrid, are given in Table 4.3. In both the adapted and non-adapted cases, run times with multigrid acceleration are approximately half those without. The non-adapted solution benefits slightly more than the adapted solution.




Figure 4.21 Nondimensional Static Pressure for Adapted (Augmentation Threshold $7.0\text{e-}8$) Harmonic Balance Solutions, Configuration SS1, 1025 Cells, With and Without Multigrid Acceleration

Table 4.3 Run-time Performance Comparison for Adapted vs Non-adapted Harmonic Balance Solutions, With and Without Multigrid Acceleration, Case SS1, 1025 Cells, in Seconds

Non-Multigrid		Multigrid	
7 Freq	Adapted	7 Freq	Adapted
102.3	88.4	50.9	45.7

4.7 Summary

The adaptive split-domain harmonic balance method was successfully applied to a variety of supersonic and subsonic one-dimensional flow fields containing strong moving shocks. The energy-based augmentation approach reliably identified cells where additional frequency content was needed and could be supported by the computational grid. The resulting adapted harmonic balance solutions were equivalent to non-adapted harmonic balance solutions, and compared well with conventional time-accurate solutions.

Adaptation scheduling was found to have a significant impact on the run-time performance of the adaptive solver. In some cases, applying the wrong scheduling approach more than doubled solution compute time. For supersonic flows, rapid adaptation with minimal flow development time produced the best performance. For subsonic flows, a scheduling approach that allowed more time for flow development between adaptations was best.

The compatibility of the adaptive split-domain harmonic balance approach and FAS multigrid acceleration was demonstrated. Solutions were computed with and without frequency augmentation, and with and without multigrid acceleration. The relative performance benefits of each approach remained consistent.

For all test configurations, adapted harmonic balance solutions took less time to compute than equivalent non-adapted solutions. Reductions from 25% to greater than 50% were observed. These reductions were obtained despite the fact that most of the test configurations contained strong discontinuities throughout the solution domain, and thus had high average adapted frequency content.

V. Application of Adaptive Split-Domain Harmonic Balance to an Unsteady Diverging Nozzle

5.1 Introduction

The test problems solved in the previous chapter using the frequency augmentation approach were useful for developing the adaptive harmonic balance approach, but they were not ideal for demonstrating performance improvements. In those test problems, there were moving shocks throughout the computational domain, which led to a high average frequency content. In a transonic turbomachinery application, one would expect large variations in flow time response at different points in the computational domain, from near steady-state flow far upstream of a blade row, to a region of strong moving shocks near the leading edge of the blades.

The purpose of the analysis documented in this chapter is to demonstrate the adaptive split-domain harmonic balance approach on a problem representative of a transonic turbomachinery problem. The chosen test case consists of a diverging nozzle with supersonic inflow and subsonic outflow. Unsteadiness is introduced by varying the subsonic outflow conditions. The resulting flow field contains steady-state flow over most of the first half of the nozzle, a narrow region near the center of the nozzle with an oscillating normal shock, and smoothly varying flow over the latter part of the nozzle. These flow features make a good test case for the frequency augmentation algorithm. The case contains mostly smooth flow with low (or no) frequency content. It abruptly transitions to a small region where the time response of the flow is essentially a square wave, and thus has high frequency content. It then abruptly transitions back to low frequency content. The robustness of the algorithm is demonstrated by handling the transitions, and the effectiveness is demonstrated by a low average frequency content and reduced computation time.

5.2 Solver Implementation

The one-dimensional split-domain harmonic balance Euler solver developed in Chapter IV was modified to solve the quasi-1-D Euler equation. The quasi-1-D Euler equation (35) in finite volume form is given by

$$\frac{\partial}{\partial \tau} \int_V \mathbf{Q}' dV + \int_S \mathbf{F}' \cdot \hat{\mathbf{n}} dS + \int_V \mathbf{H} dV = 0 \quad (5.1)$$

where

$$\mathbf{Q}' = A\mathbf{Q} = A \begin{bmatrix} \rho \\ \rho u \\ E_t \end{bmatrix}, \quad (5.2)$$

$$\mathbf{F}' = A\mathbf{F} = A \begin{bmatrix} \rho u \\ \rho u^2 + p \\ (E_t + p)u \end{bmatrix}, \quad (5.3)$$

and

$$\mathbf{H} = \begin{bmatrix} 0 \\ -p \frac{dA}{dx} \\ 0 \end{bmatrix} \quad (5.4)$$

where A is the cross sectional area of the nozzle and the other quantities are as defined for the 1-D Euler equation in Section 4.2. Substituting the approximation $\frac{\partial A}{\partial x} = \frac{\Delta A}{\Delta x}$ and assuming the computational grid does not change with time, Eq. 5.1 reduces to

$$\bar{A}_i \Delta x_i \frac{\partial \mathbf{Q}_i}{\partial \tau} + A_i^r \mathbf{F}_i^r - A_i^l \mathbf{F}_i^l - p_i (A_i^r - A_i^l) = 0 \quad (5.5)$$

for the i th grid cell, where $\bar{A}$ is the average area over the cell, and A^l and A^r are the cross sectional areas at the left and right cell interfaces respectively. Adding

artificial dissipation as discussed in Section 4.2, and recognizing that $\bar{A}_i \Delta x_i = V_i$, the approximate volume of a vertical slice through the nozzle, the semi-discretized quasi-1-D Euler equation becomes

$$\frac{\partial \mathbf{Q}_i}{\partial \tau} = -\frac{1}{V_i} (\mathbf{F}_i^r A_i^r - \mathbf{F}_i^l A_i^l - p_i(A_i^r - A_i^l) + \mathbf{D}_i). \quad (5.6)$$

Modifying the 1-D Euler code to solve the Quasi-1-D Euler equations thus required changing the definition of the cell volume, multiplication of the left and right fluxes by the appropriate cross sectional areas, and inclusion of the pressure term. A characteristic-based, specified-density outflow boundary condition (35) designed for the diverging nozzle was also incorporated (See Appendix D). Other aspects of the solver remain as described in Section 4.2.

5.3 Test Configuration

The quasi-1-D adaptive split-domain harmonic balance solver was applied to the problem of flow through a diverging nozzle with constant supersonic inflow and unsteady subsonic outflow (Fig. 5.1).




Figure 5.1 Unsteady Diverging Nozzle Configuration

A common steady-state quasi-1-D nozzle test case (36) was modified to generate an unsteady flowfield. Assumed conditions at the inlet were: Mach number = 1.5, density = 0.002241 slugs/ft³, pressure = 2000 lb_f/ft², and temperature = 520R.

Unsteadiness was introduced by adding a sinusoidal variation to the exit density.

$$\rho = 0.003954 + 0.0001 \sin(0.1\pi t) \text{ slugs/ft}^3 \quad (5.7)$$

Density was varied as a matter of convenience—the desired unsteady behavior was produced, and a specified-density nozzle boundary condition was available. For computation, all values were nondimensionalized by the inlet velocity and density.

Without the unsteady term in the exit density, these boundary conditions result in a flow field with a normal shock near the midpoint of the nozzle and subsonic exit flow (36). With the addition of the specified unsteadiness at the exit, the flow behind the shock becomes smoothly unsteady, but remains subsonic. The location of the shock oscillates about its steady-state location, while the supersonic flow ahead of the upstream limit of shock motion is not affected and remains steady.

The unsteady amplitude was selected by first finding the range of constant densities for which a steady-state solution could be obtained using the initial conditions described below. The amplitude was set to the largest value that produced upper and lower extremes within that range. The frequency $\omega = 0.1\pi$ was selected to maximize the extent of shock motion and still produce an interesting subsonic flow behind the moving shock. Higher disturbance frequencies produced a more complex subsonic flow, but resulted in a restricted range of shock motion. Lower disturbance frequencies resulted in essentially the same range of shock motion, but produced less complex flow behind the shock.

Flow solutions were obtained for three computational grids with densities of 256, 512, and 1024 cells. A uniform cell size was maintained in each grid. Solutions were initialized to inlet conditions over the first 28% of the grids. Over the remainder of the grids, density and total energy were initialized to inlet conditions, while velocity was initialized to 34% of inlet velocity. Adapted harmonic balance solutions to the unsteady problem were computed for six augmentation thresholds from 5.0e-2 to

5.0e-7 in order-of-magnitude increments. In addition, one steady-state (0-frequency) solution was computed on each grid to validate the basic solver.

Augmentation increments available to the adaptation algorithm were specified in terms of candidate numbers of frequencies that could be included in the solution in a given cell. The available candidates are listed in Table 5.1. Since the flow in much of the grid was steady-state, each of the solutions was begun with 0 frequencies. The unsteady flow downstream of the shock was smoothly varying, and thus had low frequency content, so all of the lower-frequency candidates were included. Finally, higher frequency candidates in evenly spaced increments were included to allow refinement of the moving shock region.

Table 5.1 Candidate Numbers of Frequencies Available to the Adaptive Harmonic Balance Solver for the Unsteady Nozzle Test Case

0	1	2	3	4	5
6	7	10	13	16	19
22	24	27	32	37	42

A subsonic adaptation scheduling strategy was adopted (See Section 4.4). The primary adaptation triggers for the initial and following adaptations were set to 2.6 and 0.1, respectively, while the secondary, iteration based triggers were set to 7000 and 1200. In addition, fringe augmentation widths of 4, 8, and 25 were required for the 256, 512, and 1024 cell grids, respectively.

All solutions were initiated with a CFL of 1.7. One test case (1024-cell grid with augmentation threshold of 5.0e-4) required a scaled CFL to maintain stability. For this case, the CFL was scaled down to 1.25 at the highest frequency content. All harmonic balance solutions were converged to a residual below 1.0e-6, a drop of approximately 4.6 to 4.9 orders of magnitude.

For comparison purposes, conventional time-accurate unsteady solutions were also computed on each grid. Each solution was begun from a steady-state calculation based on the mean outflow properties, and run time-accurately for 20 outflow cycles.




Figure 5.2 Time-accurate Development of the Fully-developed Unsteady Nozzle Flow Field

The mass in the nozzle was tracked as the time-accurate solution developed to verify that a fully-developed solution was reached. Based on this measure, fully developed flow was reached at about the 12th cycle (Fig. 5.2).

5.4 Results

5.4.1 Comparison with Exact Solutions. To validate the basic steady-state quasi-1-D harmonic balance solver, 0-frequency (steady-state) solutions were calculated for each grid density and compared with an exact solution. The exact solution for a diverging nozzle with constant outflow properties is easily obtained from the isentropic flow equations and Rankine-Hugoniot normal shock relations (37, 36).

All solutions showed good agreement with the theoretical solution, as illustrated in Fig. 5.3. The root-mean-square (RMS) percent error was 4.9%, 3.5%, and 2.5% for the 256 cell, 512 cell and 1024 cell solutions, respectively. Most of the observed error was attributed to smearing of the normal shock across several cells.




Figure 5.3 Comparison of 0-Frequency (Steady State) and Theoretical Pressures for 1024 Dell Grid Diverging Nozzle

Excluding the error in the shock region reduced the RMS percent error for the coarse, medium, and fine grids to 0.31%, 0.16%, and 0.08%.

The validity of the unsteady harmonic balance solutions was confirmed by examining the time-average mass flux throughout the grid. With constant supersonic flow at the inlet boundary, conservation of mass requires that the time-averaged mass flux at any given point in the nozzle be inversely proportional to the ratio of nozzle cross-sectional area at that point to the area at the inlet. In Fig. 5.4, the computed time-average mass flux (the zero-frequency Fourier coefficient of ρu) for the 256 and 1024 cell grids is plotted along with the theoretical value. The harmonic balance solutions show good agreement with the theoretical values everywhere except in the region of shock motion, where the effect of shock smearing introduces errors. The error distribution of the time-average momentum for the 256, 512, and 1024 cell harmonic balance calculations is shown in Fig. 5.5. In all three cases, the error peaks at the edges of the region of shock motion, with the maximum error ranging

from 3.5% for the 256 cell solution to 1.5% for the 1024 cell solution. Errors of this magnitude are acceptable for an unsteady analysis of this type.

5.4.2 Unsteady Nozzle Solutions. The unsteady-outflow nozzle configuration proved to be a challenging test of the adaptive harmonic balance solver. The transition from steady-state flow to high-frequency content flow presented difficulties during the early transient stages of solution development. As the solution was refined, the upstream limit of shock motion shifted upstream slightly. Higher frequencies were needed in the transition cells to allow the disturbance to shift upstream, but none were available to the solution. Because the upstream flow was supersonic and steady-state, the transition cells had zero energy in the first Fourier frequency, and so no frequencies were added by previous threshold-based augmentation. In many cases, the solution in the frequency-deficient cells broke down before an adaptation was triggered and more frequencies were added.

This situation was successfully resolved by including fringe augmentation (Section 2.5). This effectively pre-augmented a number of cells ahead of the developing unsteadiness and gave the developing unsteadiness room to propagate upstream. (Two alternative solutions to the problem—starting the solution with a larger number of frequencies, and/or decreasing the amount of flow development between adaptations by lowering the repeat adaptation triggers—were rejected as having too negative an impact on run-time performance.) Fringe augmentation provided a reliable solution to the problem with minimum increase in run time.

A typical growth history of an adapted frequency distribution is shown in Fig. 5.6. The figure shows the complete adaptation history for a solution on the finest grid with an augmentation threshold of $5.0\text{e-}7$. Each frame shows the frequency distribution after frequency augmentation at the indicated iteration, from the first adaptation at iteration 4,725, to the final adaptation at iteration 16,990. As can be seen in Fig. 5.6, the frequency distribution had stabilized outside the shock region by




Figure 5.4 Comparison of Time-average Adaptive Harmonic Balance Mass Flux with Theoretical Solution




Figure 5.5 Error Distribution for Time-average Adaptive Harmonic Balance Mass Flux




Figure 5.6 Evolution of the Adapted Frequency Distribution for Each Adaptation of the Solution, 1024 Cells, Augmentation Threshold $5.0e-7$

the fourth adaptation. Subsequent adaptations added frequencies only to the region of shock motion.

In Fig. 5.7, pressure distributions for this same adapted solution are shown at 10 snapshots in time spanning one period of the flow oscillation. Also shown are snapshots from a time-accurate calculation on the same computational grid, along with the final frequency envelope. This figure illustrates how the adapted frequency distribution relates to the computed flow solution. Figure 5.7 also illustrates how the adapted harmonic balance solution sharply captured the oscillating normal shock and agreed with the time-accurate solution at all points in time. The spatially varying frequency content produced no visible artifacts in the reconstructed harmonic balance solution.

A quantitative assessment of the adapted harmonic balance solution was obtained by calculating the RMS of the percentage differences between the adapted solution and the time-accurate solution for each of the 10 samples in Fig. 5.7, and averaging the results. The average RMS difference was 0.14%. (The maximum RMS difference over the 10 samples was 0.27%.) A similar calculation was performed for solutions obtained on all three grids with all six augmentation thresholds. The results are plotted in Fig. 5.8. While there was some difference in the average RMS values for large augmentation thresholds, the solutions on all three grids converged to provide nearly identical agreement at a threshold of $5.0\text{e-}7$.

Nearly all of the difference between the time-accurate solutions and adapted harmonic balance solutions with high augmentation thresholds was in the region of the moving shock. Figure 5.9 illustrates how this region developed with decreases in augmentation threshold. This figure shows solutions in the region of shock motion for the 256 cell and 1024 cell grids at a time corresponding to $1/4$ of the oscillation period.

For a given augmentation threshold, the solutions on both grids had similar properties. At a threshold of $5.0\text{e-}2$, the shock was poorly defined, and there was




Figure 5.7 Comparison of Adaptive Harmonic Balance and Time-accurate Solutions at 10 Points in Time Spanning One Period of Oscillation




Figure 5.8 Root-Mean-Square of Difference Between Adaptive Harmonic Balance and Time-accurate Solutions with Changing Adaptation Threshold




Figure 5.9 Time Snapshot of Adaptive Harmonic Balance Solution Illustrating Refinement of Oscillating Normal Shock with Decreased Augmentation Threshold

significant overshoot behind the shock. At a threshold of $5.0\text{e-}4$, the shock was well defined, but there was ringing before and after the shock. With a threshold of $5.0\text{e-}6$, the shock was crisply defined on both grids, and most of the ringing around the shock was removed. While the defects at each threshold level were more pronounced on the finer grid due to higher grid resolution and frequency content, the basic form of the defects was the same on both grids.

The frequency distributions for the solutions shown in Fig. 5.9 are shown in Fig. 5.10. These plots illustrate the relative number of frequencies required to achieve each level of solution fidelity. They also illustrate how the adapted frequency distribution automatically adjusted for grid density. Though a given threshold achieved a qualitatively similar solution on each grid, the number of frequencies included in the shock region on the fine grid was significantly higher than the number included on the coarse grid.

Figure 5.10 also includes the final frequency distribution for solutions with a threshold of $5.0\text{e-}7$. As can be seen, a large increase in frequency content was needed to achieve a small improvement in the solution as indicated in Fig. 5.8. The variation in maximum frequency content with grid density and augmentation threshold is shown for all grids and thresholds in Fig. 5.11.

While it was the maximum frequency content that determined the fidelity of each solution, it was the average frequency content that determined the computational cost of each solution. The average frequency content for each adapted solution is shown in Fig. 5.12. Since the discontinuous flow was limited to a small portion of the computational domain, the overall average frequency content remained small for all augmentation thresholds. On the finest grid, average frequency content at the smallest threshold grew by approximately 4.5 frequencies, while the maximum frequency grew by 35 frequencies.

Low average frequency content translated directly into reduced compute times, as shown in Fig. 5.13. As can be seen in this figure, the compute time grew more




Figure 5.10 Adapted Frequency Distributions for Different Augmentation Thresholds on Coarse Grid (top) and Fine Grid (bottom)




Figure 5.11 Change in Maximum Frequency Content of Adapted Harmonic Balance Solutions with Changes in Augmentation Threshold

slowly than the average frequency content. This was due to the inherent efficiency of resolving the flow features that are slowest to converge during the early stages of the adaptive solution when the frequency content was low. On the coarsest grid, the compute time for the lowest augmentation threshold is less than twice that of the highest threshold. On the finest grid, the compute time grew only 2.3 times from the highest threshold to the lowest.

To determine the efficiency of the adaptive harmonic balance approach relative to non-adapted harmonic balance, each of the harmonic balance solutions was re-computed with a fixed number of frequencies. The maximum frequency content for each of the adapted solutions was identified and used to generate the fixed-frequency solutions, ensuring that both solutions maintained the same fidelity. The efficiency of the adaptive method was examined by comparing problem sizes and compute times.

The problem size for the fixed frequency calculation was determined by the number of frequencies included in the solution. The equivalent problem size for the adapted calculation is given by the average frequency content. The difference be-




Figure 5.12 Change in Average Frequency Content of Adapted Harmonic Balance Solutions with Changes in Augmentation Threshold




Figure 5.13 Increase in Relative Compute Time with Decreasing Augmentation Threshold

tween these two quantities is shown in Figure 5.14 in terms of percent reduction. For the largest augmentation threshold, the reduction in problem size was approximately 45% for all grid densities. This represents a significant reduction, considering that the average number of frequencies in each solution was approximately one, the minimum possible number of frequencies for a frequency augmentation approach. Any improvement would require a frequency decimation approach where frequencies are identified as unnecessary and removed from the solution. Such an approach was not implemented in the current research. The problem size reduction was more substantial for the small thresholds, approaching 90% for the smallest threshold on the finest grid.




Figure 5.14 Reduction in Average Frequency Content for Adapted vs Non-adapted Harmonic Balance Solutions

The difference between compute times for the adapted and fixed-frequency solutions is shown in Fig. 5.15, again in terms of percent reduction. It is clear from these results that a given reduction in problem size did not necessarily result in an equivalent reduction in compute time, especially for higher augmentation thresholds and lower grid densities. Defining an efficiency of the adaptive solver as the ratio of the compute time reduction to the problem size reduction, the average efficiency




Figure 5.15 Reduction in Compute Time for Adapted vs Non-adapted Harmonic Balance Solutions

for a threshold of 5.0×10^{-2} was only 47%. For lower thresholds the efficiency was much higher, averaging 91% for a threshold of 5.0×10^{-6} , and 95% for a threshold of 5.0×10^{-7} . The adaptive solution on the finest grid with an augmentation threshold of 5.0×10^{-7} achieved a 99% adaptive efficiency with an 86% reduction in compute time. The adaptive efficiencies for all solutions are shown in Fig. 5.16.

5.5 Summary

The adaptive split-domain harmonic balance method developed and analyzed in previous chapters was successfully applied to solve the quasi-1-D inviscid flow in a supersonic-subsonic diverging nozzle with unsteady outflow properties. This problem is representative of the types of problems the adaptive harmonic balance approach was designed for: a stationary unsteady flow field containing mostly smooth, low-frequency content flow, but with isolated regions of highly nonlinear high-frequency content flow.




Figure 5.16 Efficiency of Adaptive Harmonic Balance Solutions (Run Time Reduction/Average Frequency Reduction)

Adapted harmonic balance solutions were generated for a variety of grid densities and augmentation thresholds. The energy-based frequency augmentation approach proved effective in matching frequency content to underlying flow. Regions of the computational domain requiring higher frequency content were identified, and frequencies were added until augmentation threshold levels were achieved. Given thresholds produced qualitatively uniform solutions across tested grid densities. At the lowest augmentation thresholds, the adapted harmonic balance solutions showed good agreement with both theoretical and time-accurate numerical solutions.

The combination of an energy-based frequency augmentation approach, fringe augmentation, and the split-domain solver proved to be robust and stable. The adaptation algorithm and solution scheme successfully handled a rapid transition from steady-state supersonic flow to the highly nonlinear, unsteady, mixed supersonic/subsonic flow in the region of the moving shock. In all but one case, the maximum stable CFL of the solver was able to be used for the entire solution.

The adaptive harmonic balance approach effectively reduced the time needed to obtain a high-fidelity solution with the harmonic balance method. The compute time for solutions with thresholds of $5.0\text{e-}6$ and $5.0\text{e-}7$ averaged greater than 80% reductions relative to equivalent non-adapted solutions on all grids. With these reductions, the difference in compute time between low-fidelity and high-fidelity solutions was also reduced so that the highest fidelity solutions took only 2–2.5 times longer than the lowest-fidelity solutions.

VI. Summary and Conclusions

6.1 Summary of Research

A new adaptive split-domain harmonic balance CFD method was developed and applied to a variety of one-dimensional problems. The new method employed a unique multi-domain split-operator (split-domain) solution approach to efficiently solve harmonic balance equations with large numbers of Fourier frequencies. The split-domain approach successfully removed a numerical stability restriction present in previous harmonic balance CFD implementations. It also reduced the computational work required to calculate a solution by reducing the number of FFTs required to just one per point, per iteration. Improved stability and reduced computational work both resulted in reduced run time.

To further improve performance, the split-domain solution method was combined with a frequency-adaptation algorithm that minimized the size of the harmonic balance problem by varying the number of frequencies included in the harmonic balance solution. A frequency-augmentation adaptation approach was employed. With this approach, harmonic balance solutions were begun with a minimum number of frequencies, and periodically examined to determine where additional frequencies were needed to capture the local flow. The decision to add frequencies was based on an examination of the fraction of spectral energy contained in the largest included Fourier frequency. Frequencies were added to cells that contained more energy in the highest frequency than a user-specified maximum threshold. By changing the threshold, the number of frequencies in the solution, and thus the solution accuracy and run time, were controlled.

The non-adaptive split-domain harmonic balance CFD method was applied to the 1-D inviscid Burgers' equation. Large amplitude, time-periodic solutions were computed and compared to solutions computed with prior harmonic balance approaches. The split-domain method produced solutions comparable to those of

the prior methods, while successfully eliminating a stability restriction experienced by those methods when a large number of Fourier frequencies are included in the solution. It was found that the difference between harmonic balance solutions and conventional time-accurate solutions was asymptotic with respect to the number of Fourier frequencies included in the harmonic balance solution. When the number of frequencies was equal to or greater than some asymptotic frequency, the harmonic balance solutions were comparable to the time-accurate solutions. Several factors were found to influence the asymptotic frequency, including disturbance frequency, the strength of the moving discontinuity, and the computational grid density.

The full adaptive split-domain harmonic balance CFD method was applied to the 1-D Euler equation. The method was successfully employed to solve a variety of supersonic and subsonic one-dimensional flow fields containing strong moving shocks. The energy-based augmentation approach reliably identified cells where additional frequency content was needed and could be supported by the computational grid. The resulting adapted harmonic balance solutions were equivalent to non-adapted harmonic balance solutions, and compared well with conventional time-accurate solutions.

Adaptation scheduling was found to have a significant impact on the run-time performance of the adaptive solver. In some cases, applying the wrong scheduling approach more than doubled solution compute time. For supersonic flows, rapid adaptation with minimal flow development time produced the best performance. For subsonic flows, a scheduling approach that allowed more time for flow development between adaptations was best.

The compatibility of the adaptive split-domain harmonic balance approach and FAS multigrid acceleration was demonstrated. Solutions were computed with and without frequency augmentation, and with and without multigrid acceleration. The relative performance benefits of each approach remained consistent.

Finally, the adaptive split-domain approach was demonstrated by solving for quasi-1-D flow in a supersonic-subsonic diverging nozzle with unsteady outflow properties. This problem contained a wide range of flow regimes, including supersonic steady-state flow, continuous, unsteady subsonic flow, and an oscillating normal shock. Adapted harmonic balance solutions were generated for a variety of grid densities and augmentation thresholds. The energy-based frequency augmentation approach proved effective in matching frequency content to the underlying flow. Regions of the computational domain requiring higher frequency content were identified, and frequencies were added until augmentation threshold levels were achieved. However, it was necessary to supplement threshold-based augmentation with fringe augmentation in order to successfully handle shifting frequency transitions during the early stages of solution development. Given thresholds produced qualitatively uniform solutions across tested grid densities. At the lowest augmentation thresholds, the adapted harmonic balance solutions showed good agreement with both theoretical and time-accurate numerical solutions.

The adaptive harmonic balance approach effectively reduced the time needed to obtain an accurate solution with the harmonic balance method. The compute time for solutions with thresholds of $5.0\text{e-}6$ and $5.0\text{e-}7$ averaged greater than 80% reductions relative to equivalent non-adapted split-domain solutions on all grids. With these reductions, the difference in compute time between low-fidelity and high-fidelity solutions was also reduced, to the point where the highest fidelity solutions took only 2–2.5 times longer than the lowest-fidelity solutions.

6.2 Conclusions

As a result of the research conducted and documented in this dissertation, the following major conclusions were reached:

1. *The harmonic balance method produces accurate solutions to strongly nonlinear time-periodic flow problems, provided that a sufficient number of Fourier fre-*

quencies are included. There is a practical asymptotic limit that is dependent on the flow field and on the density of the computational grid. Thus, adding frequencies does not automatically improve a harmonic balance solution.

2. *A split-operator, multi-domain (split-domain) solution approach can be employed to efficiently solve a harmonic-balance CFD problem with a sufficient number of frequencies to accurately model a strongly nonlinear periodic flow.* The split-domain approach effectively reduces or removes stability restrictions associated with high-frequency harmonic balance solutions. It also reduces the required number of Fourier transform pairs to one, regardless of the time integration scheme implemented. Finally, the approach is fully compatible with FAS multigrid convergence acceleration.
3. *The number of frequencies included in a split-domain harmonic balance solution can be varied from cell to cell without negatively impacting solution quality.* With proper implementation, transitions between regions with different frequency content are transparent and do not affect a reconstructed solution. Frequency transitions cannot be treated as time-lagged boundaries; both sides of the transition must be synchronously integrated in time. Resampling *via* truncation or zero-padding of Fourier coefficients was an effective way to compute differences across frequency transitions.
4. *For flows with time-responses that are mostly continuous with a finite number of discontinuities and no impulses, the fraction of spectral energy contained in the highest computed Fourier frequency of a harmonic balance solution (E_N) is an accurate indicator of how completely the solution has modeled local flow behavior.* Solutions containing an insufficient number of Fourier frequencies will have a higher E_N than those that contain a sufficient number. When a computational grid is too coarse to support high frequencies, the E_N is reduced. A given minimum solution fidelity can be maintained by selecting an upper

bound on E_N and including sufficient frequencies to ensure E_N is below this bound.

5. *A spatially-adaptive harmonic balance method can be implemented to accurately and efficiently compute a stationary time-periodic flow field containing regions of smooth flow and regions with strong moving discontinuities by automatically adjusting the number of frequencies in the solution, on a point-by-point (cell-by-cell) basis, to match local flow conditions.* A restatement of the research thesis. An adaptive split-domain harmonic balance CFD solver employing an energy-based frequency augmentation adaptation approach was demonstrated and shown to significantly reduce the time required to compute an accurate harmonic balance solution.

VII. Recommendations for Future Research

The next logical step in the development of the adaptive split-domain harmonic balance approach is the extension of the method to higher dimensions and application to turbomachinery problems. Extension of the basic approach should be straightforward, requiring only minor modifications to the current implementation. Part of this effort should be the development of a boundary condition to model a rotor-stator interface. Additional areas that merit further research are discussed below.

Combined Grid/Frequency Adaptation

Grid adaptation provides a means of efficiently resolving flow features when their precise location is not known *a priori*. Applying grid adaptation to conventional unsteady CFD problems is difficult, however, because flow features are in constant motion. Since the harmonic balance method solves a steady-state problem, implementation of grid adaptation should be greatly simplified. One of the primary goals of research into combined grid and frequency adaptation would be to determine if frequency content information could be used along with flow feature information to drive adjustments to the computational grid.

Improved Multigrid for Harmonic Balance CFD

One of the primary incentives for developing harmonic balance CFD approaches is that a steady-state problem is solved, which allows convergence acceleration techniques such as multigrid to be employed. The current research demonstrated the compatibility of FAS multigrid and adaptive split-domain harmonic balance, but identifying the best multigrid approach was beyond scope. Identifying an optimal multigrid approach for harmonic balance problems would be of great value.

One property of a harmonic balance solution that limits the effectiveness of multigrid is that the real and imaginary parts of the high frequency Fourier coeffi-

cients tend to be highly oscillatory in the spatial dimensions. Conventional multigrid techniques are most effective when the spatial behavior of the solution is smooth, so that little information is lost when the solution is transferred to a coarser grid. When multigrid is applied to oscillatory problems, information is lost during grid transfer operations, and convergence acceleration is reduced. Some attempts have been made to develop multigrid techniques for oscillatory elliptic PDEs (e.g. (38)). One avenue of research would be to adapt these techniques to the harmonic balance equations.

Another approach suggested by the current research would be to transfer solutions between grids using the magnitude-phase form of the Fourier coefficients, rather than the real and imaginary form. It was observed that the magnitude of the complex Fourier coefficients was considerably less oscillatory than the real and imaginary components. The phase, when “unwrapped” so that it was continuous and had values greater than 2π , was nearly linear with respect to the spatial dimension. The smoother properties of a magnitude/phase should allow more accurate transfer between grids. The major challenge facing such an approach is developing a reliable means of unwrapping phase.

Magnitude/Phase Harmonic Balance Solver

In order to accurately capture the oscillatory behavior of the real and imaginary coefficients of higher-frequency Fourier coefficients, a fine computational grid is necessary. As noted above, spatial variation of the magnitude and unwrapped phase of the coefficients is much smoother, and thus can be captured with a much coarser grid (39). Development of a harmonic balance method based on a magnitude/phase representation of the complex coefficients, could significantly reduce the number of grid cells required to compute an accurate harmonic balance solution, and thus reduce the memory and run-time requirements as well. A magnitude/phase solver

would also improve multigrid efficiency, increasing the accuracy of the corrections computed on coarse grids.

Many aspects of a magnitude/phase solver would be similar to the current split-domain solver. Boundary condition and flux calculations could still be done in time domain. However, face fluxes would likely be computed from averaged cell magnitude and phase, instead of averaged cell fluxes as in the current implementation. As with magnitude/phase multigrid transfers, a major research challenge is determining how to correctly implement continuous phase.

Efficient Single-domain Harmonic Balance

The suggested research areas discussed thus far have all involved the extension or modification of the adaptive split-domain harmonic balance method. An investigation of an efficient single-domain harmonic balance approach, while suggested by the split-domain approach, is a research area that could result in a significant departure from the current method.

The basis of the proposed research is the physical interpretation of the split-domain solution process discussed in Section 2.4. It was shown that the split-domain solution approach is equivalent to a sequence of successive integrate-shift operations in which Eq. 2.23a, repeated here as Eq. 7.1, is integrated forward in pseudo time, and the elements of the solution vector $\bar{\xi}$ are shifted backward in physical time.

$$\frac{\partial \bar{\xi}}{\partial \tau} + \frac{\partial \Phi(\bar{\xi})}{\partial x} = 0 \quad (7.1)$$

A steady-state solution is achieved when the change in $\bar{\xi}$ due to the pseudo-time integration step is the negative of the change due to the physical time shift.

The combined size of the physical-time shifts was shown to be equal to $\frac{(2N+1)\omega\Delta\tau}{2\pi}\Delta t$, where Δt is the physical time sample increment. If $\Delta\tau = \Delta t = \frac{2\pi}{(2N+1)\omega}$, then the size of the physical time shift is equal to the integration step size, Δt . In

this case, the distinction between pseudo-time and physical time disappears, and the change in $\bar{\xi}$ due to the shift operation is simply the difference between consecutive elements of $\bar{\xi}$. For this integration step size, there is no need to transform to the frequency domain to compute the shift.

Instead of simply fixing the time integration step size and replacing the frequency domain operations, an alternative form of the harmonic balance equations could be employed. If pseudo-time is replaced by physical time and Eq. 7.1 is integrated from an initial condition at time t to time $t + \Delta t$, then the harmonic balance problem reduces to

$$\mathbf{A} \left[\bar{\xi} + \frac{\partial \Phi(\bar{\xi})}{\partial x} \right] = 0, \quad (7.2)$$

where $\mathbf{A}$ is a $2N + 1 \times 2N + 1$ matrix with the form

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & \dots & 0 & -1 \\ -1 & 1 & 0 & & 0 \\ 0 & -1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix}. \quad (7.3)$$

This approach would be more flexible, because no specific integration time step is required.

Both single-domain approaches eliminate the need for Fourier transforms during the solution process, and thus might reduce the time required to compute a solution. The second approach has the additional advantage that the frequency of oscillation, ω , does not appear in the equations, and thus would not need to be known *a priori*; it would enter the solution only through the boundary conditions.

References

1. Verdon, J. M. and Caspar, J. R. "Subsonic Flow past an Oscillating Cascade with Finite Mean Flow Deflection," *AIAA Journal*, (18):540 1980.
2. Verdon, J. M. and Caspar, J. R. "Development of a Linear Unsteady Aerodynamic Analysis for Finite-Deflection Subsonic Cascades," *AIAA Journal*, (20):1259 1982.
3. Verdon, Joseph M. and Caspar, Joseph R. "A Linearized Unsteady Aerodynamic Analysis for Transonic Cascades," *Journal of Fluid Mechanics*, 149:403–429 December 1984.
4. Verdon, Joseph M. "Linearized Unsteady Aerodynamic Theory." *AGARD Manual on Aeroelasticity in Axial-Flow Turbomachines Volume 1, Unsteady Turbomachinery Aerodynamics* Number AGARD-AG-298 in AGARDograph, edited by Max F. Platzer and Carta O. Franklin, AGARD, 1987.
5. Hall, Kenneth C. and Crawley, Edward F. "Calculation of Unsteady Flows in Turbomachinery Using the Linearized Euler Equations," *AIAA Journal*, 27(6):777–787 June 1989.
6. Hall, Kenneth C. and Clark, William S. "Linearized Euler Predictions of Unsteady Aerodynamic Loads in Cascades," *AIAA Journal*, 31(3):540–550 March 1993.
7. Verdon, Joseph M., Barnett, Mark and Ayer, Timothy C. *Unsteady Aerodynamic Models for Turbomachinery Aeroelastic and Aeroacoustic Applications*. NASA Contractor Report NASA/CR 4698, NASA, November 1995.
8. Cizmas, P. G. A. and Hall, K. C., "A Viscous-Inviscid Model of Unsteady Small-Disturbance Flows in Cascades." AIAA 95-2655, 31st AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, San Diego CA, July 10–12 1995.
9. Verdon, Joseph M., Montgomery, Matthew D. and Chuang, Andrew H. *Development of a Linearized Unsteady Euler Analysis With Application to Wake/Blade-Row Interactions*. NASA Contractor Report NASA/CR 1999-208879, NASA, November 1999.
10. Hoyniak, Daniel and Clark, William S., "Aerodynamic Damping Predictions Using a Linearized Navier-Stokes Analysis." ASME 99-GT-207, International Gas Turbine & Aeroengine Congress & Exhibit, Indianapolis IN, June 7–10 1999.
11. Clark, William S. and Hall, Kenneth C., "A Time-Linearized Navier-Stokes Analysis of Stall Flutter." ASME 99-GT-383, International Gas Turbine & Aeroengine Congress & Exhibit, Indianapolis IN, June 7–10 1999.

12. Florea, R. and Hall, K. C., "Sensitivity Analysis of Unsteady Inviscid Flow through Turbomachinery Cascades." AIAA 2000-0130, 38th AIAA Aerospace Sciences Meeting & Exhibit, Reno NV, January 10–13 2000.
13. Ning, W. and He, L. "Computation of Unsteady Flows Around Oscillating Blades Using Linear and Nonlinear Harmonic Euler Methods," *Journal of Turbomachinery*, 120:508–514 1998.
14. He, L and Ning, W. "Efficient Approach for Analysis of Unsteady Viscous Flows in Turbomachines," *AIAA Journal*, 17(11):2005–2012 November 1998.
15. Chen, T, Vasanthakumar, P. and He, L. "Analysis of Unsteady Blade Row Interaction Using Nonlinear Harmonic Approach," *AIAA Journal of Propulsion and Power*, 17(3):651–658 May 2001.
16. Hall, K. C., Thomas, J. P. and Clark, W. S. *Computation of Unsteady Nonlinear Flows in Cascades Using a Harmonic Balance Technique*. Technical Report, Lyon, France: 9th International Symposium on Unsteady Aerodynamics, Aeroacoustics, and Aeroelasticity of Turbomachines, September 2000.
17. McMullen, Matthew, Jameson, Antony and Alonso, Juan J., "Acceleration of Convergence to a Periodic Steady State in Turbomachinery Flows." AIAA 2001-0152, 39th AIAA Aerospace Sciences Meeting & Exhibit, Reno NV, January 8–11 2001.
18. Thomas, Jeffrey P., Dowell, Earl H. and Hall, Kenneth C., "Nonlinear Inviscid Aerodynamic Effects on Transonic Divergence, Flutter, and Limit Cycle Oscillations." AIAA 2001-1209, 42nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference and Exhibit, Seattle WA, April 16–19 2001.
19. McMullen, Matthew, Jameson, Antony and Alonso, Juan J., "Application of a Non-Linear Frequency Domain Solver to the Euler and Navier-Stokes Equations." AIAA 2002-0120, 40th AIAA Aerospace Sciences Meeting & Exhibit, Reno NV, January 14–17 2002.
20. Hutton, T. M., Leger, T. J., Johnston, D. A. and Wolff, J. M., "High Spatial Resolution IGV Forced Response in a Transonic Compressor obtained with a Flex Sensor Pressure Array." AIAA 02-4222, 38th AIAA/ASME/SAE/ASEE Joint Propulsion Meeting, Indianapolis, IN, July 7-10 2002.
21. Hayashi, Chihiro. *Nonlinear Oscillations in Physical Systems*, 28–31. Princeton, NJ: Princeton University Press, 1964.
22. Frigo, Matteo. "A Fast Fourier Transform Compiler." *Proceedings of the 1999 ACM SIGPLAN Conference on Programming Language Design and Implementation*. May 1999.

23. Toro, Eleuterio F. *Fiemann Solvers and Numerical Methods for Fluid Dynamics*, chapter 15. New York: Springer-Verlag, 1997.
24. LeVeque, R. J. and Yee, H. C. "A Study of Numerical Methods for Hyperbolic Conservation Laws with Stiff Source Terms," *Journal of Computational Physics*, 86:187–210 1990.
25. Schwer, Douglas A. and Green, Jr., William H., "Split-Operator Methods for Computing Steady-State Reacting Flow Fields." AIAA 2001-2635, 15th AIAA Computational Fluid Dynamics Conference, Anaheim, CA, June 11–14 2001.
26. Strang, G. "On the Construction and Comparison of Difference Schemes," *SIAM Journal on Numerical Analysis*, 5(3):506–517 1968.
27. Brigham, E. Oran. *The Fast Fourier Transform*, 124. Englewood Cliffs, NJ: Prentice-Hall, 1974.
28. Cannizzaro, F. E., Elmiligi, A., Mulson, D. N. and von Lavante, E., "A Multi-block Multigrid Method for the Solution of the Three-dimensional Euler Equations." AIAA 90-0105, 28th Aerospace Sciences Meeting, Reno NV, January 8-11 1990.
29. Tannehill, John C., Anderson, Dale A. and Pletcher, Richard H. *Computational Fluid Mechanics and Heat Transfer* (Second Edition), 187–189. Washington D.C.: Taylor & Francis, 1997.
30. Jameson, A., Schmidt, W. and Turkel, E., "Numerical Solutions of the Euler Equations by Finite Volume Methods Using Runge-Kutta Time-Stepping Schemes." AIAA 81-1259, 14th AIAA Fluid and Plasma Dynamics Conference, Palo Alto CA, June 23–25 1981.
31. Swanson, R.C. and Turkel, E., "Artificial Dissipation and Central Difference Schemes for the Euler and Navier-Stokes Equations." AIAA 87-1107, 8th Computational Fluid Dynamics Conference, Palo Alto CA, June 23–25 1987.
32. Enander, Rickard. "Implicit Explicit Residual Smoothing for the Multidimensional Euler and Navier-Stokes Equations," *SIAM Journal of Scientific Computing*, 18(5):1243–1254 September 1997.
33. Swanson, R. C. and Turkel, Eli. *Multistage Schemes with Multigrid for Euler and Navier-Stokes Equations*. NASA Technical Paper 3631, NASA, August 1997.
34. Janus, Jonathan M. *The Development of a Three-Dimensional Split Flux Vector Euler Solver with Dynamic Grid Applications*. MS thesis, Mississippi State University, August 1984.
35. Shubin, G. R., Stephens, A. B. and Glaz, H. M. "Steady Shock Tracking and Newton's Method Applied to One-Dimensional Duct Flow," *Journal of Computational Physics*, 39:364–374 1981.

36. Hoffman, Klaus A. and Chiang, Steve T. *Computational Fluid Dynamics* (Fourth Edition), *Volume II*, 133–151. Wichita: Engineering Education System, August 2000.
37. Zucrow, Maurice J. and Hoffman, Joe D. *Gas Dynamics, Volume I*, 162–178, 335–351. New York: John Wiley & Sons, 1976.
38. Wan, W. L., Chan, Tony F. and Smith, Barry. “An Energy-Minimizing Interpolation for Robust Multigrid Methods,” *SIAM Journal of Scientific Computing*, 21(4):1632–1649 2000.
39. Orkwis, Paul D. Private conversation, 2002.
40. Frigo, M and Johnson, S. G. “FFTW: An Adaptive Software Architecture for the FFT.” *International Conference on Acoustics, Speech, and Signal Processing 1998 Volume 4*. 1381–1384.
41. Bertin, John J. and Smith, Michael L. *Aerodynamics for Engineers* (Second Edition), 308–309. Prentice Hall, 1989.

Appendix A. Derivation of Harmonic Balance Burgers' Equation for a Complex Fourier Series

In this appendix, a direct-substitution implementation of the one-dimensional harmonic balance Burgers' equation is derived.

The one-dimensional inviscid Burgers equation is given by

$$\frac{\partial u}{\partial t} + \frac{1}{2} \frac{\partial u^2}{\partial x} = 0 \quad (\text{A.1})$$

where $u(x, t)$ is the dependent variable, t is the temporal variable, and x is the spatial variable. Let $u(x, t)$ be approximated by a truncated complex Fourier series

$$u(x, t) \approx \sum_{n=-N}^N c_n(x) e^{in\omega t} \quad (\text{A.2})$$

where N is the number of positive frequencies in the truncated series, and ω is the fundamental frequency of the series. In this series, the complex coefficients, a_n , are functions of the spatial variable only. Substituting Eq. A.2 into Eq. A.1 and evaluating the time derivative, one obtains

$$\sum_{n=-N}^N in\omega c_n e^{in\omega t} + \frac{1}{2} \frac{\partial}{\partial x} \left(\sum_{m=-N}^N \sum_{n=-N}^N c_m c_n e^{i(m+n)\omega t} \right) = 0. \quad (\text{A.3})$$

As a result of squaring the approximating series in the flux term, Eq. A.3 contains high frequency terms with $|m + n| > N$. As part of the harmonic balance approximation, these terms are discarded, leaving only terms with frequencies included in the approximating Fourier series. Gathering the remaining terms of like frequency gives

$$\sum_{k=0}^N \left[ik\omega c_k + \frac{1}{2} \frac{\partial}{\partial x} \left(\sum_{m=-N}^N \sum_{n=-N}^N c_m c_n \delta_{(m+n),k} \right) \right] e^{ik\omega t} = 0. \quad (\text{A.4})$$

In these equations, $\delta_{(m+n),k}$ is the Kronecker delta function, defined as 1 if $(m+n) = k$, and 0 if $(m+n) \neq k$, thus the summation inside the spatial derivative contains only products of c_m and c_n such that $m+n = k$.

The frequencies in Eq. A.4 are now “balanced”. Since the right hand side is identically zero, this simply means that each of the terms in the outer summation is required to satisfy equality individually. For the resulting $2N+1$ complex equations to hold for all time, the terms inside the square brackets must be identically equal to zero. The exponential terms can be dropped, leaving a system of $2N+1$ coupled ordinary differential equations for $2N+1$ complex coefficients. Since the dependent variable, $u(x, t)$, is real and periodic, the negative frequency Fourier coefficients are just the complex conjugates of the corresponding positive frequency coefficients, i.e. $c_{-n} = \tilde{c}_n$. Furthermore, the steady-state coefficient c_0 must be real, bringing the total number of unknowns to $2N+1$: the real and imaginary parts of N positive-frequency coefficient, plus the steady-state coefficient.

Keeping just the positive frequency equations and eliminating the exponential terms, Eq. A.4 reduces to

$$ik\omega c_k + \frac{1}{2} \frac{d}{dx} \left(\sum_{m=-N}^N \sum_{n=-N}^N c_m c_n \delta_{(m+n),k} \right) = 0 \quad 0 \leq k \leq N \quad (\text{A.5})$$

or in vector form,

$$\hat{\mathbf{S}}(\hat{\mathbf{U}}) + \frac{d\hat{\mathbf{F}}(\hat{\mathbf{U}})}{dx} = 0 \quad (\text{A.6})$$

where

$$\hat{\mathbf{U}} = \begin{bmatrix} c_0 \\ c_1 \\ \vdots \\ c_N \end{bmatrix} \quad \hat{\mathbf{S}}(\hat{\mathbf{U}}) = \begin{bmatrix} 0 \\ i\omega c_1 \\ \vdots \\ iN\omega c_N \end{bmatrix} \quad \hat{\mathbf{F}}(\hat{\mathbf{U}}) = \frac{1}{2} \begin{bmatrix} \sum_{m=-N}^N \sum_{n=-N}^N c'_m c'_n \delta_{(m+n),0} \\ \sum_{m=-N}^N \sum_{n=-N}^N c'_m c'_n \delta_{(m+n),1} \\ \vdots \\ \sum_{m=-N}^N \sum_{n=-N}^N c'_m c'_n \delta_{(m+n),N} \end{bmatrix} \quad (\text{A.7})$$

and c'_n is defined as c_n if $n \geq 0$, and $\tilde{c}_n$ if $n < 0$. An example of the harmonic balance Burgers' equation flux vector is presented here for $N = 3$.

$$\hat{\mathbf{F}}(\hat{\mathbf{U}})_{N=3} = \begin{bmatrix} \frac{1}{2}(c_0 + |c_1|^2 + |c_2|^2 + |c_3|^2) \\ c_0 c_1 + \tilde{c}_1 c_2 + \tilde{c}_2 c_3 \\ c_0 c_2 + \tilde{c}_1 c_3 \\ c_0 c_3 \end{bmatrix} \quad (\text{A.8})$$

Appendix B. Impact of Operator Splitting Error in the Split-Domain Harmonic Balance Solution

Introduction

One potential problem with a split-operator solution technique such as the split-domain harmonic balance formulation, is that for nonlinear problems, the operator splitting is not exact (23). In addition, discretization of a split equation typically results in the introduction of an error that is dependent on the size of the numerical time step taken to solve the problem. Therefore, it is necessary to understand the nature of the splitting error for a specific application before the approach is used. This appendix examines the error introduced by the split-domain harmonic balance approach both analytically and experimentally.

Linear analysis of Splitting-induced Error

Consider the linear scalar equation

$$\frac{\partial \xi}{\partial t} + a\xi_x - b\xi = 0 \tag{B.1}$$

where ξ and b are complex scalars, and ξ_x is the spatial derivative of ξ . Splitting into two homogeneous equations gives

$$\frac{\partial \xi}{\partial t} + a\xi_x = 0 \tag{B.2}$$

$$\frac{d\xi}{dt} - b\xi = 0. \tag{B.3}$$

For simplicity, assume a first-order forward-Euler discretization in time. The effective finite difference equation resulting from a symmetric Strang splitting which evaluates the ODE twice is obtained by successively applying the time discretization.

Integrate the ODE 1/2 time step:

$$\xi^* = \xi + \frac{\Delta t}{2} b\xi \quad (\text{B.4})$$

Integrate the PDE a full time step using the results of the previous integration:

$$\begin{aligned} \xi^{**} &= \xi^* - \Delta t (a(\xi^*)_x) \\ &= \xi + \frac{\Delta t}{2} b\xi - \Delta t a\left(\xi + \frac{\Delta t}{2} b\xi\right)_x \\ &= \xi + \Delta t \left(\frac{b}{2}\xi - a\xi_x\right) - \frac{\Delta t^2}{2} ab\xi_x \end{aligned} \quad (\text{B.5})$$

Finally, integrate the ODE an additional 1/2 time step, again using the results of the previous integration step:

$$\begin{aligned} \xi^{n+1} &= \xi^{**} + \frac{\Delta t}{2} b\xi^{**} \\ &= \xi + \Delta t \left(\frac{b}{2}\xi - a\xi_x\right) - \frac{\Delta t^2}{2} ab\xi_x + \frac{\Delta t}{2} b\xi + \\ &\quad \Delta t^2 \left(\frac{b^2}{4}\xi - \frac{ab}{2}\xi_x\right) - \frac{\Delta t^3}{4} ab^2\xi_x \\ &= \xi + \Delta t (b\xi - a\xi_x) + \Delta t^2 \left(\frac{b^2}{4}\xi - ab\xi_x\right) - \frac{\Delta t^3}{4} ab^2\xi_x \end{aligned} \quad (\text{B.6})$$

Subtracting $\xi + \Delta t (b\xi - a\xi_x)$ from both sides of Eq. B.6 and dividing by Δt gives

$$\frac{\xi^{n+1} - \xi}{\Delta t} + a\xi_x - b\xi = \Delta t \left(\frac{b^2}{4}\xi - ab\xi_x\right) - \frac{\Delta t^2}{4} ab^2\xi_x \quad (\text{B.7})$$

At steady-state, the first term on the left hand side is equal to zero. The semi-discretized equation satisfied at steady-state is therefore given by

$$a\xi_x - b\xi = \Delta t \left(\frac{b^2}{4}\xi - ab\xi_x\right) + \text{H.O.T.} \quad (\text{B.8})$$

The first term on the right hand side of Eq. B.8 represents the time-step-dependent splitting error. More insight into the effect of this error is gained by discarding the higher order terms and rewriting the equation.

$$(b\Delta t + 1)a\xi_x - \left(b + \frac{b^2\Delta t}{4}\right)\xi = 0 \quad (\text{B.9})$$

This shows that the splitting error affects both the wave speed and the strength of the source term. Since the current research involved only explicit solvers, $b\Delta t$ is assumed to be much less than a , and the error in wave speed is disregarded. The remaining discussion will concentrate on the source term error component.

The harmonic balance source term is given by

$$\hat{\mathbf{S}}(\hat{\xi}) = \begin{bmatrix} 0 \\ i\omega c_1 \\ \vdots \\ iN\omega c_N \end{bmatrix}. \quad (\text{B.10})$$

The n^{th} harmonic balance source term has the same form as the source in the model equation, with $b = i\omega$. Assuming the results of the simplified analysis can be applied to each term independently, the effective source term for the split-domain harmonic balance equations becomes

$$\hat{\mathbf{S}}(\hat{\xi}) = \begin{bmatrix} 0 \\ \left(-\frac{\omega^2\Delta t}{4} + i\omega\right)c_1 \\ \vdots \\ \left(-\frac{\omega^2N^2\Delta t}{4} + iN\omega\right)c_N \end{bmatrix}. \quad (\text{B.11})$$

The splitting error adds a negative real component to the harmonic balance source coefficient. Several conclusions can be drawn from this result. First, it is clear that the steady-state (zero-frequency) equation is unaffected by splitting error. The

error should remain small for low frequencies, but will increase rapidly in the high frequencies due to the N^2 term. Since the error term contains the square of the fundamental frequency, the splitting error will be somewhat more prominent for high-frequency unsteady-flows.

Experimental Analysis of Splitting-induced Error

To assess the actual effect of splitting error on a harmonic balance solution, supersonic and subsonic solutions (configurations SS3 and SU1 from Chapter IV) were computed at a variety of CFLs. Each of the solutions was computed on a 601 point grid. The supersonic solutions were computed with 32 frequencies, while the subsonic solutions were computed with 45 frequencies.

Figures B.1–B.4 show a comparison of the real part of the 5th, 15th, 25th, and 32nd Fourier coefficients of momentum, respectively, for the supersonic case at CFLs of 0.5, 1.1, and 1.7. As predicted by the simple linear analysis, there was little difference between the solutions for the low frequency coefficient (Fig. B.1). For higher Fourier frequencies, the effect of splitting error became more pronounced (Figs. B.2, B.3, and B.4). As the splitting error increased with larger time steps (CFLs), the solution became more dissipative. The frequency of the spatial oscillations in each coefficient remained essentially unchanged. Results for the imaginary part of the Fourier coefficients were similar.

Figures B.5–B.9 show similar comparisons of the real part of the 5th, 15th, 25th, 35nd and 45th Fourier coefficients of momentum for the subsonic case at CFLs of 0.6, 1.0, and 1.4. For low to mid frequencies (Figs. B.5, B.6, and B.7), the splitting error exhibits the same increasingly dissipative effect as in the supersonic case. However, at higher frequencies, the dissipative effect of the splitting error seems to decrease again.

The discussion so far has focused on the impact of splitting error on individual Fourier coefficients. Figures B.10 and B.11 illustrate how the splitting error affects the reconstructed solutions.

As can be seen in Fig. B.10, splitting error had little effect on the supersonic solution. Solutions at the lowest and highest CFL were qualitatively similar except for the presence of increased oscillations behind the moving shocks with smaller time steps. The subsonic solution also experiences an increase in oscillations behind the moving shocks with a smaller time step (Fig. B.11). However, in the trough regions, the low-CFL solution shows a slight improvement over the high-CFL solution, partially overcoming the error introduced at the exit boundary (See Section 4.5).

For both the supersonic and subsonic cases, the oscillations that occurred behind the shocks for low-CFL solutions were not readily removed by increasing the amount of artificial dissipation in the solution. Any dissipation strong enough to affect these oscillations also damped low-frequency modes that were not significantly affected by changes in time step size, resulting in an overall degradation of the solution.

Conclusions

A simplified analysis of splitting-induced error in the split-domain harmonic balance method suggested that any time-step size dependency in a converged solution would be most evident in the higher-frequency terms. Just such an error, in the form of damping of high-frequency Fourier coefficients, was observed experimentally for both supersonic and subsonic flow fields. The observed error had little effect on reconstructed flow properties such as pressure. If anything, the increased splitting error resulting from a larger time step had a positive rather than negative effect on the solution, eliminating high-frequency oscillations that were otherwise difficult to remove.




Figure B.1 Effect of Splitting Error on 5th Fourier Coefficient with Changing CFL for Supersonic Flow




Figure B.2 Effect of Splitting Error on 15th Fourier Coefficient with Changing CFL for Supersonic Flow




Figure B.3 Effect of Splitting Error on 25th Fourier Coefficient with Changing CFL for Supersonic Flow




Figure B.4 Effect of Splitting Error on 32nd Fourier Coefficient with Changing CFL for Supersonic Flow




Figure B.5 Effect of Splitting Error on 5th Fourier Coefficient with Changing CFL for Subsonic Flow




Figure B.6 Effect of Splitting Error on 15th Fourier Coefficient with Changing CFL for Subsonic Flow




Figure B.7 Effect of Splitting Error on 25th Fourier Coefficient with Changing CFL for Subsonic Flow




Figure B.8 Effect of Splitting Error on 35th Fourier Coefficient with Changing CFL for Subsonic Flow




Figure B.9 Effect of Splitting Error on 45th Fourier Coefficient with Changing CFL for Subsonic Flow




Figure B.10 Effect of Splitting Error on Pressure with Changing CFL for Supersonic Flow




Figure B.11 Effect of Splitting Error on Pressure with Changing CFL for Subsonic Flow

Appendix C. Determination of Optimal Frequency Augmentation Increments

The number of Fourier frequencies N included in an adapted split-domain harmonic balance solution can have a large influence on the time needed to compute the solution. To find the values of N that give the best overall performance, a numerical experiment was conducted. In this experiment, a one-dimensional split-domain Euler solver (see Section 4.2) was run a fixed number of iterations for every number of frequencies N from 3 to 100. The execution time for each run was recorded, normalized by the execution time for $N = 3$, and plotted in Fig. C.1.




Figure C.1 Effect of Series Length on Split Domain Solver Run Time

Values of N that result in locally minimized execution time were identified as candidates for frequency augmentation. These are listed in Table C.1. Note that

Table C.1 Candidate Numbers of Fourier Frequencies, N , for an Efficient Adaptive Harmonic Balance Solution

0	1	2	3	4	5	6
7	10	12	13	16	17	19
22	24	27	31	32	37	38
40	45	49	52	58	62	67
73	82	87	94	97		

frequency counts 0, 1, and 2 are also included in this table, though they were not included in the experiment.

The only computations in the split-domain harmonic balance method that are not linearly dependent on N are the FFT and inverse FFT used to transform the solution vector between the frequency domain and the time domain. The values of N identified in Table C.1 are therefore those that minimized the cost of the FFT calculations. In the tested solver implementation, FFT calculations were performed via “Fastest Fourier Transform in the West” (FFTW) library calls (40). The locally minimizing values of N could change if a different FFT library was employed.

While any adaptive harmonic balance computation could include every value of N listed in Table C.1, some values could be omitted if information is known about the flow being calculated. If the flow field is known to contain strong moving discontinuities throughout the entire computational domain, then it will have a high frequency content everywhere as well. In such a case, it is unnecessary to include every possible low-frequency candidate. For such a calculation, the following values might be appropriate: 6, 12, 17, 24, 32, 37, 40, 45, 49, ... In this set of N , some lower values have been omitted, because it is known that the lower N will not appear in the final frequency distribution. The low N that are included simply provide “stepping stones” to the final solution. All of the higher N have been included, however, so that the final frequency distribution can be fit to the flow as closely as possible. If, on the other hand, the flow field being modeled contains regions of smoothly unsteady (or even steady-state) flow, it would be appropriate to include all

of the candidate N . In any case, it will not affect the final solution if all candidate frequencies are included in a calculation. It simply may take longer to compute than would otherwise be necessary.

Appendix D. Implementation of Boundary Conditions for the 1-D and Quasi-1-D Euler Equation

Boundary conditions for the adaptive split-domain 1-D Euler and quasi-1-D Euler solvers were enforced by calculating appropriate values for a single ghost cell outside the boundary of the grid. Ghost cell values were computed after the first integration of the frequency-domain ODE and subsequent transformation to the time domain, but before integration of the time domain equation (between steps 2 and 3 of the split-domain update process as defined in Section 2.4, page 2-10.) This approach allowed the use of conventional steady-state CFD boundary conditions.

Supersonic Boundary Conditions

A characteristic analysis of 1-D supersonic flow shows that all flow information propagates in the downstream direction (34). Therefore, at a supersonic inflow boundary, all flow information comes from outside the computational domain. Conversely, at a supersonic outflow boundary all flow information comes from inside the computational domain.

Supersonic Inflow. At a supersonic inflow boundary, all ghost cell values must be specified. At the inlet, Mach number, density and pressure were given. Inlet velocity was obtained using the definition of Mach number and the relation for the speed of sound in a perfect gas:

$$u = M \sqrt{\frac{\gamma p}{\rho}} \quad (\text{D.1})$$

Unsteady supersonic boundary conditions were obtained by varying Mach number while holding pressure and density constant. The unsteady Mach number was given by

$$M(t) = 2.0 + a \sin(\omega t) \quad (\text{D.2})$$

where a and ω were the amplitude and frequency of the unsteady variation. For harmonic balance solutions, samples of the unsteady Mach number were obtained at $2N + 1$ points spanning one period of oscillation, where N was the number of frequencies included in the harmonic balance solution in the interior cell adjacent to the boundary.

$$M(n) = 2.0 + a \sin\left(\frac{2\pi n}{2N + 1}\right) \quad n = 0, 1, \dots, 2N \quad (\text{D.3})$$

Conservative flow variables were constructed from the specified primitive variables.

$$\mathbf{Q}_g = \begin{bmatrix} \rho \\ \rho u \\ \rho \left(\frac{p}{(\gamma-1)\rho} + \frac{1}{2}u^2 \right) \end{bmatrix} \quad (\text{D.4})$$

Supersonic Outflow. At a supersonic outflow boundary, ghost cell values were obtained by means of a zeroth-order extrapolation from the interior cell adjacent to the boundary.

$$\mathbf{Q}_g = \mathbf{Q}_{\text{adj}} \quad (\text{D.5})$$

Piston Boundary Condition

The action of an oscillating piston was simulated by specifying a sinusoidal velocity in the ghost point, while maintaining a zero pressure gradient and constant stagnation enthalpy. A zero-pressure gradient was maintained by setting the pressure in the ghost cell equal to the pressure in the interior cell adjacent to the boundary, i.e., $p = p_{\text{adj}}$. For inviscid flow of a perfect gas, stagnation enthalpy, given by $H = e + p/\rho + \frac{1}{2}u^2$, is constant (41). Therefore, it could be computed from the initial conditions in the tube. Using the perfect gas relation $p = (\gamma - 1)\rho e$, and the fact

that the initial velocity was zero, the stagnation enthalpy was given by

$$H = \frac{p_{\text{init}}}{(\gamma - 1)\rho_{\text{init}}}. \quad (\text{D.6})$$

For time-accurate solutions, the velocity in the ghost cell was given by

$$u(t) = a \sin(\omega t). \quad (\text{D.7})$$

For harmonic balance solutions, samples of the sinusoidal velocity were obtained at $2N + 1$ points spanning one period of oscillation, where N was the number of frequencies included in the harmonic balance solution in the interior cell adjacent to the boundary.

$$u(n) = a \sin\left(\frac{2\pi n}{2N + 1}\right) \quad n = 0, 1, \dots, 2N \quad (\text{D.8})$$

Given u , p , and H , the conservative variables in the ghost cell were computed.

$$\mathbf{Q}_g = \begin{bmatrix} \frac{\gamma p}{(\gamma-1)(H-\frac{1}{2}u^2)} \\ \rho u \\ \rho \left(\frac{p}{(\gamma-1)\rho} + \frac{1}{2}u^2 \right) \end{bmatrix} \quad (\text{D.9})$$

Subsonic Characteristic Farfield Boundary Condition

A characteristic analysis of 1-D subsonic flow reveals that there are two characteristics traveling downstream, and one characteristic traveling upstream (34). Thus at a subsonic inflow boundary, two properties must be specified and one must be computed from the interior flow field. Conversely, at a subsonic outflow boundary, only one property can be specified, and two properties must be computed from the interior flow field. The boundary conditions implemented in this research are developed in (34). Only the implementation is presented here.

Inflow. Given specified farfield properties ρ_f , u_f , and p_f , and properties in the first interior cell, ρ_{adj} , u_{adj} , and p_{adj} , define

$$\rho_0 = \frac{1}{2}(\rho_f + \rho_{\text{adj}}) \quad (\text{D.10})$$

$$p_0 = \frac{1}{2}(p_f + p_{\text{adj}}) \quad (\text{D.11})$$

$$a_0 = \sqrt{\frac{\gamma p_0}{\rho_0}} \quad (\text{D.12})$$

The primitive variables at the boundary are given by

$$p_b = p_f + p_{\text{adj}} + \rho_0 a_0 (u_f - u_{\text{adj}}) \quad (\text{D.13})$$

$$\rho_b = \rho_f + \frac{p_f - p_b}{a_0^2} \quad (\text{D.14})$$

$$u_b = u_f + \frac{p_f - p_b}{\rho_0 a_0} \quad (\text{D.15})$$

The primitive variables in the ghost cells were obtained through a first-order extrapolation of the boundary values and the values in the first cell in the grid.

$$\rho_g = 2\rho_b - \rho_{\text{adj}} \quad (\text{D.16})$$

$$u_g = 2u_b - u_{\text{adj}} \quad (\text{D.17})$$

$$p_g = 2p_b - p_{\text{adj}} \quad (\text{D.18})$$

These primitive variables were used to compute conservative variables in the ghost point using Eq. D.4.

Outflow. For the subsonic outflow boundary condition, pressure was specified, and density and velocity were computed. The primitive variables at a

subsonic outflow boundary were given by

$$p_b = p_f \quad (\text{D.19})$$

$$\rho_b = \rho_{\text{adj}} + \frac{p_b - p_{\text{adj}}}{a_0^2} \quad (\text{D.20})$$

$$u_b = u_{\text{adj}} + \frac{p_{\text{adj}} - p_b}{\rho_0 a_0}. \quad (\text{D.21})$$

These values were extrapolated to the ghost point and conservative variables were computed using Eq. D.4.

Subsonic Nozzle Outflow Boundary Condition

A variation of the characteristic subsonic outflow BC was incorporated into the quasi-1-D Euler solver. This boundary condition, developed in (35), takes into account the varying geometry of the nozzle. In addition, the specified property at the boundary is density rather than pressure.

At the exit boundary, the time-accurate unsteady density was specified according to

$$\rho(t) = \rho_{\text{avg}} + a \sin(\omega t). \quad (\text{D.22})$$

For harmonic balance solutions, samples of the unsteady density were obtained at $2N + 1$ points spanning one period of oscillation, where N was the number of frequencies included in the harmonic balance solution in the interior cell adjacent to the boundary.

$$\rho(n) = \rho_{\text{avg}} + a \sin\left(\frac{2\pi n}{2N + 1}\right) \quad n = 0, 1, \dots, 2N \quad (\text{D.23})$$

The remaining primitive variables at the exit were given by the solution to two differential equations (35)

$$-p \frac{\partial \rho}{\partial x} + \rho^2 \frac{\partial e}{\partial x} = 0 \quad (\text{D.24a})$$

$$(a + u) \left(\frac{\partial p}{\partial x} + \rho a \frac{\partial u}{\partial x} \right) + \frac{1}{A} \frac{dA}{dx} u \rho a^2 = 0 \quad (\text{D.24b})$$

In Eq. D.24b, A is the cross sectional area of the nozzle at the exit, and a is the local speed of sound.

Equation D.24a was solved for e_b by approximating the spatial derivatives with first-order backward differences, and substituting the known exit density.

$$e_b = \frac{e_{\text{adj}}}{1 - (\gamma - 1)(1 - \rho_{\text{adj}}/\rho)} \quad (\text{D.25})$$

The exit velocity, u , was obtained by solving Eq. D.24b. The spatial derivatives were replaced with first-order backward differences and terms were rearranged to form a quadratic expression for u . Of the two possible solutions, one produced a non-physical answer when the nozzle cross sectional area was assumed constant and the pressure gradient $\frac{\partial p}{\partial x}$ was negative. This solution was discarded, leaving the following equation for u :

$$u = \frac{-\beta + \sqrt{\beta^2 + 4\rho a^2(p_{\text{adj}} - p + \rho a u_{\text{adj}})}}{2\rho a} \quad (\text{D.26})$$

where

$$\beta = \rho a^2 \frac{\Delta A}{A_{\text{exit}}} + (p - p_{\text{adj}}) + \rho a(a - u_{\text{adj}}) \quad (\text{D.27})$$

and the quantity ΔA was the change in nozzle cross sectional area across the cell adjacent to the boundary.

The quantities r , u , and e , were used to compute conservative variables in the ghost point using Eq. D.4.

Vita

Lieutenant Colonel Raymond C. Maple graduated from Union Grove High School in Gladewater, TX in May 1981. He was awarded a 4-year Air Force Reserve Officer Training Corps scholarship and accepted into the Sibley School of Mechanical Engineering at Cornell University. In June 1985, he received a Bachelors Degree in Mechanical Engineering and was commissioned a Second Lieutenant in the U. S. Air Force.

In October 1985, he entered active duty with the 3925th Intercontinental Ballistic Missile Facilities Engineering Squadron, Offutt AFB NE, where he worked as hardness and survivability engineer. In 1988, he was assigned to the Air Force Institute of Technology (AFIT), where he earned an M. S. in Aeronautical Engineering in December 1989. He spent the next four years with the Munitions Directorate of Wright Laboratories, Eglin AFB, where he developed computational fluid dynamics software to analyze aircraft store separation. In 1994, he reported to United States Strategic Command, Offutt AFB NE, where he worked as a program manager in the Strategic War Planning System program office. In 1997, he joined the Air Force Materiel Command Inspector General team, where he served as a Command Science and Technology Management Inspector.

In August 1999, he returned to AFIT to pursue a Doctorate of Philosophy in Aerospace Engineering. His current research activities include Computational Fluid Dynamics algorithm development and parallel computing. Upon graduation, he will join the AFIT faculty as an Assistant Professor of Aerospace Engineering.