



Adaptive harmonic balance method for nonlinear time-periodic flows [☆]

Raymond C. Maple ^{a,*}, Paul I. King ^a, Paul D. Orkwis ^b, J. Mitch Wolff ^c

^a *Department of Aeronautics and Astronautics, Air Force Institute of Technology, AFIT/ENY, 2950 P. Street Bldg 640, Wright-Patterson, AFB OH 45433-7765, USA*

^b *Department of Aerospace Engineering and Engineering Mechanics, University of Cincinnati, Mail Location 70, Cincinnati, OH 45221-0070, USA*

^c *Department of Mechanical and Materials Engineering, Wright State University, Dayton, OH 45435, USA*

Received 7 September 2002; received in revised form 29 July 2003; accepted 14 August 2003

Abstract

An adaptive harmonic balance approach is developed and applied to a supersonic/subsonic diverging nozzle subject to unsteady periodic outflow conditions. The adaptive method minimizes the computational work required to obtain the harmonic balance solution by augmenting the frequency content in each cell as required to capture the local flow physics. Augmentation automatically adjusts with grid density, resulting in lower frequency content on coarse grids. The new adaptive harmonic balance method produces solutions equivalent to a non-adapted harmonic balance solutions, but with up to an 86% reduction in computational effort.

© 2003 Published by Elsevier B.V.

1. 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 dynamic (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

[☆] The views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Air Force, the Department of Defense, or the US Government. This paper is a work of the US Government and is not subject to copyright protection in the United States.

* Corresponding author. Tel.: +1-937-255-6565; fax: +1-937-656-7621.

E-mail addresses: raymond.maple@afit.edu (R.C. Maple), paul.king@afit.edu (P.I. King), paul.orkwis@uc.edu (P.D. Orkwis), mitch.wolff@wright.edu (J. Mitch Wolff).

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 [1–4], the time-averaged nonlinear harmonic technique [5–7], and the harmonic balance technique [8–13].

All three techniques are closely related in that they all assume a harmonic form for the unsteady flow, which allows one to recast the unsteady problem as a steady-state problem, and thus employ steady-state convergence acceleration techniques to reduce computation time. The time-linearization technique assumes a small-disturbance (i.e., infinitesimal amplitude) form for the unsteadiness, and incorporates no unsteady nonlinearities. The time-averaged nonlinear harmonic technique also assumes a small perturbation form of the governing equations, but applies time-averaging to produce a nonlinear harmonic form that is equivalent to a first-order harmonic balance. Both approaches typically solve for one harmonic at a time, and thus do not include nonlinear harmonic coupling effects. The harmonic balance method, on the other hand, eliminates the small disturbance restriction and, by solving a set of nonlinear equations for several harmonics simultaneously, captures coupled nonlinear effects. This makes the harmonic balance technique uniquely suited for calculating nonlinear periodic flow through transonic turbomachinery. Early tests show that the harmonic balance approach can produce useful results with a small number of frequencies in an order of magnitude less time than time-accurate calculations [8,10].

The fidelity of a harmonic balance solution is dependent on grid density, on the number of harmonics retained, and on the flow being modeled. On a given grid, a flow that is smoothly unsteady (i.e., no moving discontinuities) will require fewer harmonics to achieve a desired modeling accuracy than a flow containing a moving shock. Because the computational cost increases with each harmonic included in the harmonic balance solution approach, it is preferable to use the least number of harmonics needed for a given problem to achieve a desired level of accuracy.

In a typical unsteady transonic turbomachinery problem, the nature of the flow can vary significantly throughout the domain of interest. Existing harmonic balance implementations [8–10] 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 runtime (fewer harmonics), and fidelity in the regions of strong nonlinearity. In this paper, it is shown that compromise is obviated through the use of an adaptive harmonic balance technique whereby the number of Fourier frequencies at each grid point is varied in a prescribed manner to efficiently capture the local flow. The technique is demonstrated with the solution of the quasi-one-dimensional inviscid Euler flow equations applied to a supersonic/subsonic diverging nozzle with periodic unsteady outflow conditions.

2. Split-domain harmonic balance form of the governing equations

The quasi-1D Euler equations [14] in strong conservation vector form are given by

$$\frac{\partial \mathbf{Q}}{\partial t} + \frac{\partial \mathbf{F}}{\partial x} + \mathbf{H} = 0, \quad (1)$$

where

$$\mathbf{Q} = A \begin{bmatrix} \rho \\ \rho u \\ E_t \end{bmatrix}, \quad \mathbf{F} = A \begin{bmatrix} \rho u \\ \rho u^2 + p \\ (E_t + p)u \end{bmatrix}, \quad \mathbf{H} = \begin{bmatrix} 0 \\ -p \frac{\partial A}{\partial x} \\ 0 \end{bmatrix}. \quad (2)$$

In Eq. (2), A is the cross-sectional area of the nozzle, and ρ , u , E_t , and p are density, velocity, total energy, and pressure, respectively. The perfect gas relation relates pressure to the conservative variables, and the assumption of constant ratio of specific heats closes the governing flow equations.

Eq. (1) can be transformed into a split-domain harmonic balance form [12]. To simplify the discussion, this transformation is presented for a homogeneous scalar conservation equation of the form

$$\frac{\partial \xi}{\partial t} + \frac{\partial \Phi(\xi)}{\partial x} = 0. \quad (3)$$

The harmonic balance method is based on the assumption that the dependent variable, $\xi(x, t)$, can be approximated by a truncated Fourier series:

$$\xi(x, t) \approx \sum_{n=-N}^N c_n(x) e^{in\omega t}. \quad (4)$$

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 time for one 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. (4) into the time derivative term of Eq. (3) yields

$$\frac{\partial \xi}{\partial t} \approx \sum_{n=-N}^N in\omega c_n e^{in\omega t}. \quad (5)$$

A similar substitution is performed for the spatial derivative term of Eq. (3). In this case, substitution requires the evaluation of the flux, 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

$$\Phi(\xi) \approx \sum_{n=-M}^M \phi_n e^{in\omega t}, \quad (6)$$

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

$$\Phi(\xi) \approx \sum_{n=-N}^N \phi_n e^{in\omega t}. \quad (7)$$

Substitution of Eqs. (5) and (7) into Eq. (3) and collection of terms with like exponents yields the series approximation form of the conservation equation

$$\sum_{n=-N}^N \left[in\omega c_n + \frac{d\phi_n}{dx} \right] e^{in\omega t} = 0. \quad (8)$$

At this point, the harmonic terms are required to balance across the equality, i.e., each frequency is required to satisfy equality independently of the other frequencies. Since the right-hand side of Eq. (8) is identically zero, each of the terms in the summation on the left-hand side must be equal to zero, i.e.,

$$\left[in\omega c_n + \frac{d\phi_n}{dx} \right] e^{in\omega t} = 0, \quad -N \leq n \leq N. \quad (9)$$

For each of the $2N + 1$ complex equations represented by Eq. (9) to hold for all time, the terms inside the square brackets must be identically equal to zero. Dropping the exponential terms leaves 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., $c_{-n} = \overline{c_n}$. Furthermore, since the zero-frequency term c_0 must be real, the total number of unknowns is still $2N + 1$: 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. (9) reduces to

$$i\omega c_n + \frac{d\phi_n}{dx} = 0, \quad 0 \leq n \leq N \tag{10}$$

or in vector form,

$$\frac{d\hat{\Phi}(\hat{\xi})}{dx} + \hat{\mathbf{S}}(\hat{\xi}) = 0 \tag{11}$$

with

$$\hat{\xi} = \begin{bmatrix} c_0 \\ c_1 \\ \vdots \\ c_N \end{bmatrix}, \quad \hat{\mathbf{S}}(\hat{\xi}) = \begin{bmatrix} 0 \\ i\omega c_1 \\ \vdots \\ iN\omega c_N \end{bmatrix}, \quad \hat{\Phi}(\hat{\xi}) = \begin{bmatrix} \phi_0 \\ \phi_1 \\ \vdots \\ \phi_N \end{bmatrix}. \tag{12}$$

To solve Eq. (11), a pseudo-time derivative of the Fourier coefficient vector, $\hat{\xi}$, can be added:

$$\frac{\partial \hat{\xi}}{\partial \tau} + \frac{\partial \hat{\Phi}(\hat{\xi})}{\partial x} + \hat{\mathbf{S}}(\hat{\xi}) = 0. \tag{13}$$

The solution is converged to steady state with a numerical time-marching scheme, at which time the pseudo-time derivative vanishes, and Eq. (11) is recovered.

Numerical stability considerations require a reduced time step when solving Eq. (13) for large N [9,12]. To improve the stability characteristics of the numerical solution, Eq. (13) can be split into a homogeneous partial differential equation, and an ordinary differential equation

$$\frac{\partial \hat{\xi}_1}{\partial \tau} + \frac{\partial \hat{\Phi}(\hat{\xi}_1)}{\partial x} = 0, \tag{14a}$$

$$\frac{d\hat{\xi}_2}{d\tau} + \hat{\mathbf{S}}(\hat{\xi}_2) = 0. \tag{14b}$$

An approximate solution to Eq. (13) is obtained by sequentially solving Eqs. (14a) and (14b) using the solution from one as the initial condition for the other. For example, one can solve Eq. (14a) 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. (13). The approximate solution is then obtained by solving Eq. (14b) for $\hat{\xi}_2$, with the initial condition $\hat{\xi}_2(0) = \hat{\xi}_1(x, \tau)$. When a time marching numerical approach is employed, this sequential solution approach is applied every time step. Different numerical approaches can be used for each equation. The stability characteristics of the overall solution are determined by the stability of the method applied to each of Eqs. (14a) and (14b) individually [15,16]. By choosing appropriate discretizations, the large N time step restriction is greatly reduced or eliminated.

As an alternative to the simple splitting approach described above, an approximate numerical solution can also be obtained using a Strang symmetric splitting approach [15–18]. Given difference operators $\mathcal{M}_{\Delta\tau}$ and $\mathcal{L}_{\Delta\tau}$ that are second-order accurate for Eqs. (14a) and (14b), respectively, then the composite operator $\mathcal{L}_{\Delta\tau/2}\mathcal{M}_{\Delta\tau}\mathcal{L}_{\Delta\tau/2}$ is second-order accurate for Eq. (13) [15]. In contrast, the simple splitting approach results in overall first-order accuracy.

For nonlinear problems, operator splitting introduces an error that is dependent on the integration time step size [18]. In theory, the operators \mathcal{M} and \mathcal{L} could be based on either implicit or explicit temporal discretizations. For explicit schemes, the maximum stable time step is sufficiently small that splitting error is unlikely to be significant. The effect of splitting error on solutions where larger (implicit) time steps are taken has not been investigated.

Evaluating the harmonic balance flux vector, $\hat{\Phi}$, can be computationally expensive, requiring between N^3 and N^4 complex operations per evaluation. Considerable computational savings can be gained by transforming Eq. (14a) to the time domain [8,10,12]. Given the vector of Fourier coefficients $\hat{\xi}_1$, and a discrete Fourier transform (DFT) operator, \mathcal{F} , that produces positive-frequency Fourier coefficients from a set of real numbers, define $\bar{\xi}$ such that

$$\hat{\xi}_1 = \mathcal{F}\bar{\xi} \quad (15a)$$

or

$$\bar{\xi} = \mathcal{F}^{-1}\hat{\xi}_1. \quad (15b)$$

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 $\Delta t = 2\pi/(\omega(2N + 1))$, the period of oscillation divided by the number of samples. A similar relationship is assumed for the flux vector, $\hat{\Phi}$ and its inverse DFT, $\bar{\Phi}$ [8,10].

Given these relationships, Eq. (14a) can be transformed to the time domain, greatly simplifying its solution. First, the harmonic balance flux term is approximated by

$$\hat{\Phi}(\hat{\xi}_1) \approx \mathcal{F}\Phi(\mathcal{F}^{-1}\hat{\xi}_1) = \mathcal{F}\bar{\Phi}(\bar{\xi}), \quad (16)$$

where the flux function $\bar{\Phi}$ is just the original flux function Φ acting on each sampled value in $\bar{\xi}$.

Substituting Eqs. (15a) and (16) into Eq. (14a) gives

$$\frac{\partial(\mathcal{F}\bar{\xi})}{\partial\tau} + \mathcal{F}\frac{\partial(\bar{\Phi}(\bar{\xi}))}{\partial x} = 0. \quad (17)$$

Applying the inverse transform \mathcal{F}^{-1} to Eq. (17) yields the time-domain form of Eq. (14a):

$$\frac{\partial(\bar{\xi})}{\partial\tau} + \frac{\partial(\bar{\Phi}(\bar{\xi}))}{\partial x} = 0. \quad (18)$$

Eq. (18) is similar to Eq. (3), except that integration is performed in pseudo-time, and the dependent variable is a vector of time-sampled dependent variable values rather than a single scalar. Eqs. (14b) and (18) are designated the split-domain harmonic balance form of Eq. (3). Including the discrete Fourier transforms, the solution operator for these equations becomes $\mathcal{L}_{\Delta\tau/2}\mathcal{F}\bar{\mathcal{M}}_{\Delta\tau}\mathcal{F}^{-1}\mathcal{L}_{\Delta\tau/2}$, where $\bar{\mathcal{M}}_{\Delta\tau}$ is a second-order difference operator for Eq. (18).

The split-domain harmonic balance form of Eq. (1) is obtained in a similar manner, and is given by

$$\frac{\partial\bar{\mathbf{Q}}}{\partial\tau} + \frac{\partial\bar{\mathbf{F}}}{\partial x} + \bar{\mathbf{H}} = 0, \quad (19a)$$

$$\frac{\partial \hat{\mathbf{Q}}}{\partial \tau} + \hat{\mathbf{S}} = 0. \quad (19b)$$

The dependent variable in Eq. (19a) is $\overline{\mathbf{Q}}$, which consists of the $2N + 1$ time samples of the quasi-1D dependent variable \mathbf{Q} . The vectors $\overline{\mathbf{F}}$ and $\overline{\mathbf{H}}$ are the corresponding flux and quasi-1D source vectors. The dependent variable in Eq. (19b), $\hat{\mathbf{Q}}$, is formed by concatenating the DFTs of the sampled conserved variables, and the vector $\hat{\mathbf{S}}$ is the corresponding harmonic balance source term.

The split-domain harmonic balance solution approach has been successfully applied to a variety of test cases [12,13], requiring a broad range of Fourier frequencies, from less than 10 to more than 80. The number of frequencies required for a given solution fidelity depends on several factors, including the nature of the flow field (e.g., the presence or absence of moving shocks), the grid density, and the fundamental frequency of the unsteadiness. As a general rule, flows with strong moving shocks require more frequencies. For a given shock structure, flows with a lower fundamental frequency require more total frequencies than flows with a higher fundamental frequency. Finally, because high-frequency solutions require the support of a fine computational grid, solutions on coarse grids require fewer frequencies than solutions on fine grids.

3. Adaptive split-domain solver

A new CFD solver incorporating the adaptive split-domain harmonic balance approach of Maple et al. [13] was written to solve Eqs. (19a) and (19b). Eq. (19a) was discretized with a Jameson, Schmidt, Turkel (JST)-type cell-centered finite volume scheme [19], resulting in second-order spatial accuracy. The quasi-1D source term was incorporated into the flux balance. To avoid oscillations at flow discontinuities, modified JST artificial dissipation [19,20] was used in conjunction with globally applied second-order dissipation. The global dissipation term was not strictly required, but was included in some solutions to remove small oscillations near shocks that were not completely eliminated by the JST dissipation terms.

The split-domain harmonic balance solution was advanced in pseudo-time by integrating Eq. (19b) forward one-half pseudo-time step, transforming the intermediate solution to the time domain, integrating Eq. (19a) for one full pseudo-time step and finally transforming the result back to the frequency domain where Eq. (19b) was integrated another one-half pseudo-time step. Pseudo-time integration for both Eqs. (19a) and (19b) was accomplished with a three-stage Runge–Kutta scheme with good high-frequency smoothing properties and minimal storage requirements [21]. Local time stepping was employed to accelerate convergence. The time step in each cell was determined by selecting the smallest maximum step across all computed samples. The maximum stable CFL for the solver was 1.7.

Boundary conditions for the split-domain approach can be applied in either the frequency domain or the time domain. In the current implementation, standard characteristic variable steady-state boundary conditions were applied in the time domain. The subsonic nozzle outflow boundary condition was developed in [14].

3.1. Frequency augmentation

The adaptive harmonic balance method minimizes the computational cost of the basic harmonic balance solution by automatically adjusting, on a cell-by-cell basis, the number of Fourier frequencies included in the solution. This is accomplished by means of frequency augmentation. With this approach, the solution is started 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. The final solution contains the minimum number of frequencies in each cell that efficiently and robustly model the local flow.

The decision to augment frequencies at a given point is made by examining the fraction of spectral energy contained in the highest computed Fourier frequency, defined as

$$E_N = \frac{|\mathbf{C}_N|^2}{\sum_{n=0}^N |\mathbf{C}_n|^2}, \quad (20)$$

where all operations are performed element-wise. The elements of E_N are spatially smoothed (in this case with a 5-point arithmetic average) and then compared to a threshold value, E_{thresh} . When any element exceeds the threshold, additional frequencies are incorporated into the solution at that cell. Fourier coefficients for the new frequencies are initialized to zero. The threshold, which is generally greater than $5.0e-8$, is selected to achieve a desired solution fidelity.

Threshold-based augmentation was supplemented by fringe augmentation. The purpose of fringe augmentation was to increase the size of a threshold-augmented region. In cases with very rapid transitions between smooth and discontinuous flow, the location of the transitions can change as frequencies are added and the solution is refined. In those situations, it is 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.

Once a cell was identified for augmentation, its maximum Fourier frequency was increased by a pre-determined increment, chosen to minimize computational cost. The cost of a solution is composed mainly of two components – the cost associated with solving the time-domain equations (Eq. (19a)), and the cost of performing the necessary Fourier transforms. The cost of solving the time-domain equations increases linearly with N , and there is no performance advantage favoring any particular increment, so long as the increment is relatively small. The cost of the Fourier transform varies with M , the number of terms being transformed, from order $M \log_2(M)$ for M a power of 2, to M^2 for M prime. Because the split-domain harmonic balance method is based on a set of $M = 2N + 1$ real numbers, the ideal asymptotic complexity of $M \log_2(M)$ is never attained. However, modern FFT algorithms such as those implemented in the Fastest Fourier Transform in the West (FFTW) library used in this effort yield nearly as good performance if the number of terms can be factored by powers of small prime numbers, i.e., $2N + 1 = 2^a 3^b 5^c 7^d 11^e$ for some positive integers a , b , c , d , and e [22]. Candidate frequency augmentation increments were selected by testing the solver (which incorporated FFTW library calls) for various N and identifying those N that minimized overall compute time.

Like the augmentation increment, the timing of frequency adaptation was based on run-time performance considerations. A frequency augmentation approach 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, unnecessary frequencies can 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, compute times can increase. The goal of adaptation scheduling is to adapt the solution at times that yield the best performance.

In the current work, a dual-trigger adaptation scheduling approach was employed. The primary trigger was based on a measure of the flow development, as given by a modified L_2 norm of the change in the solution during one iteration. The L_2 norm, or residual, was defined as

$$R = \sqrt{\frac{\sum_{i=1}^{ni} (\Delta \hat{\mathbf{Q}}_i \cdot \widetilde{\Delta \hat{\mathbf{Q}}_i})}{ni}}, \quad (21)$$

where, ni is the number of cells in the computational grid, $\Delta \hat{\mathbf{Q}}$ is the change in the solution vector in one iteration, and the over-tilde ($\widetilde{}$) indicates complex conjugation. Adaptation was triggered when $\log_{10}(R)$

dropped by a user-specified amount. The secondary trigger was based on the number of iterations completed. Iteration-based scheduling served as a backup to residual-based scheduling in the rare cases where convergence stalled and the specified residual drop was not achieved.

Both 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. Typical values for the initial residual-based trigger are 0.25 for supersonic flows and 2.25 for subsonic flows. Subsequent residual trigger values of 0.0 (supersonic) and 0.2 (subsonic) remove errors introduced by frequency augmentation and, in the subsonic case, further refine the solution a small amount between adaptations. Values for the iteration-based triggers are set such that the desired residual drop could reasonably be achieved if the solution converged normally. For a supersonic problem, values of order 100 (initial) and 25 (subsequent) are typical, while for a subsonic problem, values of order 10,000 and 2000 are typical.

3.2. Sample interpolation

Discretization of Eq. (19a) requires the addition and subtraction of the vectors $\bar{\mathbf{Q}}$ (and $\bar{\mathbf{F}}$) within a discretization stencil. This presents a problem when cells in the discretization stencil have different sample rates. Not only do the vectors have different numbers of elements, but those elements correspond to the state of the flow at different points in time. To resolve this problem, cell quantities must be resampled so that the sample rate is consistent across the discretization stencil.

Two different resampling methods were implemented: truncation/zero-padding in the frequency domain and linear interpolation in the time domain. With the first method, a time-domain 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 the Fourier coefficient vector is padded with new zero-valued high-frequency coefficients rather than being truncated. 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 sometimes results in non-physical values.

When non-physical interpolation occurred, linear interpolation in the time-domain was applied. With the linear interpolation method, conservative variables were obtained at the required sample times by linearly interpolating the calculated values. Besides guaranteeing that the interpolated values were bounded by the computed values, this approach had the advantage that it was applied entirely in the time domain, and did not require any additional Fourier transforms.

Despite the higher computational cost and potential non-physical interpolations, frequency-domain truncation/zero-padding was selected as the primary resampling method. This choice was based on the quality of the solution computed at sample rate transitions. The superiority of the frequency-domain method over linear interpolation is illustrated in Fig. 1. 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 appearing in the plots as dashed lines 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 order to achieve a smooth solution at frequency transitions, it is necessary to resample the transition boundaries at every Runge–Kutta (RK) integration stage. Freezing the resampled values at the first

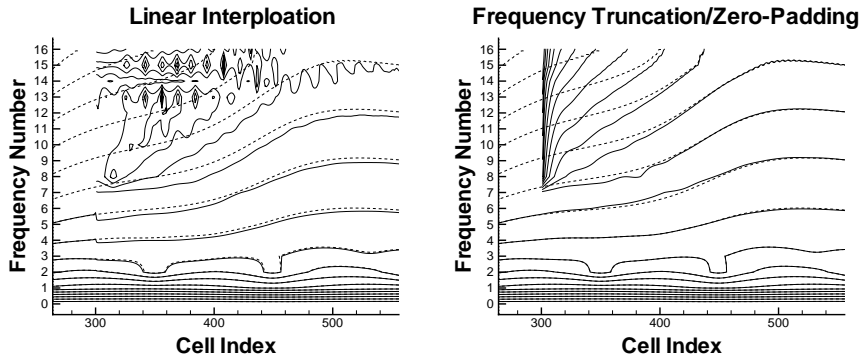


Fig. 1. Comparison of Fourier coefficient magnitude contours (momentum term) for different interpolation schemes at a transition from 7 to 16 frequencies with contours from a constant 16 frequency solution.

integration stage results in discontinuities at the transition point, even when both sides of the transition have the same number of frequencies and no interpolation is required.

As a consequence of the need to resample the transition boundaries at each integration stage, the previously uncoupled system of equations represented by Eq. (19a) became effectively coupled through the resampling operation. This made it necessary to store intermediate solutions for all of the equations in the system, significantly increasing the storage requirements of the solver. For this reason, it was important that the adaptive split-domain harmonic balance method be implemented with a numerical integration scheme with low storage requirements.

Because the adaptive solution approach presented above is primarily concerned with individual cell properties, extension to multiple dimensions is straightforward. In addition, the adaptive solution approach presented here is fully compatible with Full Approximation Storage (FAS) multigrid convergence acceleration [23], though it is not employed in this effort. Multigrid efficiency comparable to that of a non-harmonic-balance steady-state problem with similar geometry and flow features can be achieved.

4. Test configurations

The quasi-1D 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. 2).




Fig. 2. Unsteady diverging nozzle configuration.

A common steady-state quasi-1D nozzle test case [24] 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:

$$\rho = 0.003954 + 0.0001 \sin(\omega t) \text{ slugs/ft}^3. \quad (22)$$

For computation, all properties were non-dimensionalized by the inlet values.

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 [24]. With the addition of the specified unsteadiness at the exit, the flow behind the shock becomes continuously 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. A reduced disturbance frequency $\omega = 0.125$ (based on the inlet velocity and inlet nozzle height) was selected to maximize the extent of shock motion and produce multiple frequencies in the subsonic flow downstream of 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 subsonic flow with lower frequency content.

Flow solutions were obtained for three computational grids with densities of 256, 512, and 1024 cells, respectively. 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 (zero-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 for a given cell. The available candidates are listed in Table 1. Since the flow in much of the grid was steady-state, each of the solutions was begun with zero 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. All of these candidates locally minimize run time as discussed above.

The initial adaptation triggers were set to 2.6 (residual drop) and 7000 (iterations). After the initial adaptation, trigger values of 0.1 and 1200 were used. Fringe augmentation widths of 4, 8, and 25 were required for the 256-, 512-, and 1024-cell grids, respectively.

All solutions except one were computed with a fixed CFL of 1.7. The one exception was the solution on the 1024-cell grid with augmentation threshold of $5.0e-4$. For this one case, a transient in the shock region made it necessary to decrease the CFL in cells with high frequency content in order to achieve a converged

Table 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

steady-state solution. All harmonic balance solutions were converged to a residual below $1.0e-6$, a drop of approximately 4.6–4.9 orders of magnitude.

5. Results

5.1. Comparison with exact solutions

The quasi-1D harmonic balance solver was validated by calculating zero-frequency (steady-state) solutions for each grid density and comparing results 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 [24,25].

All solutions showed good agreement with the theoretical solution, as illustrated for the 1024-cell grid in Fig. 3. The root-mean-square (RMS) percent error was 4.9%, 3.5%, and 2.5% for the 256-, 512-, and 1024-cell solutions, respectively. Most of the observed error was attributed to smearing of the normal shock across several cells. 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%, respectively.

Validation of the unsteady harmonic balance solutions was primarily achieved through comparison with conventional time-accurate calculations. A comparison of time-average static pressure for the adapted harmonic balance and time-accurate calculations on the 1024-cell grid is shown in Fig. 3. Other comparisons are discussed below. One quantity, the time-average mass flux, was readily validated against an analytic solution. Conservation of mass requires that

$$\frac{(\overline{\rho u})_x}{(\overline{\rho u})_{\text{inlet}}} = \frac{A_{\text{inlet}}}{A_x}, \quad (23)$$

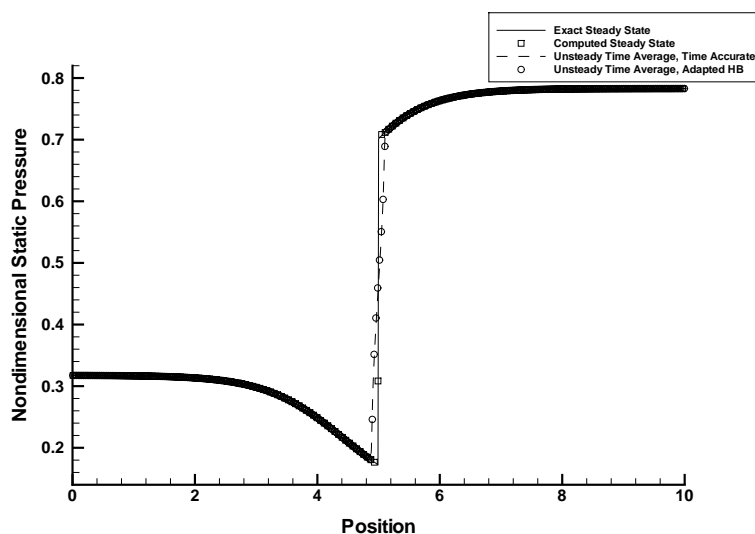


Fig. 3. Calculated zero-frequency (steady-state) and theoretical exact pressures for 1024-cell grid diverging nozzle, with the computed time-averaged pressure for time-accurate and harmonic balance calculations.

where $(\overline{\rho u})_x$ is the time-average mass flux at location x . In Fig. 4, the adaptive harmonic balance time-average mass flux (the zero-frequency Fourier coefficient of the ρu term) for the 256- and 1024-cell grids is plotted 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. In all three cases, the error peaks at the edges of the




Fig. 4. Comparison of adapted harmonic balance time-average mass flux with theoretical solution.




Fig. 5. Error distribution for adapted harmonic balance time-average mass flux.

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 these magnitudes are expected for a diverging nozzle solution (cf. [24]).

5.2. Unsteady nozzle solutions

A typical growth history of an adapted frequency distribution is shown in Fig. 6. The figure shows the adaptation history for a solution on the finest grid with an augmentation threshold of $5.0e-7$. Each frame shows the frequency distribution after frequency augmentation at the indicated iteration, from the first adaptation at iteration 4725, to the final adaptation at iteration 16,990. As can be seen in the figure, the frequency distribution had stabilized outside the shock region by the fourth adaptation. Subsequent adaptations added frequencies only to the region of shock motion.

In Fig. 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 distribution envelope. This figure illustrates how the number of frequencies included in the adapted solution varies with local flow behavior. Fig. 7 also illustrates the ability of the adapted harmonic balance solution to sharply capture the oscillating normal shock and agree with the time-accurate solution at all points in time. The spatially varying frequency content produced no visible artifacts in the computed solution.

A quantitative assessment of the adapted harmonic balance solution was obtained by averaging the RMS of the percent differences between the adapted solution and a time-accurate solution for each of the 10 samples in Fig. 7. The average RMS difference was found to be 0.14%. A similar calculation was performed for solutions obtained on all three grids with all six augmentation thresholds. The results are plotted in Fig. 8. While there were some differences across grids 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 $5.0e-7$.

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. Fig. 9 illustrates the development of this region 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 $5.0e-2$, the shock was poorly defined, and there was significant overshoot behind the shock. At a threshold of $5.0e-4$, the shock was well defined, but there was ringing before and after the shock. For a threshold of $5.0e-6$, the shock was well 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 in Fig. 9 are shown in Fig. 10. These plots illustrate the relative number of frequencies required to achieve each level of solution fidelity. They also illustrate the tendency of the adapted frequency distribution to automatically adjust 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.

Fig. 10 also includes the final frequency distribution for solutions with a threshold of $5.0e-7$. As can be seen, a large increase in frequency content was needed to achieve a small improvement in the solution as measured by the RMS difference with the time-accurate solution (Fig. 8). The variation in maximum frequency content with grid density and augmentation threshold is shown for all grids and thresholds in Fig. 11.

While the maximum frequency content determines the fidelity of each solution, the average frequency content determines the cost of each solution. The average frequency content for each adapted solution is shown in Fig. 12.

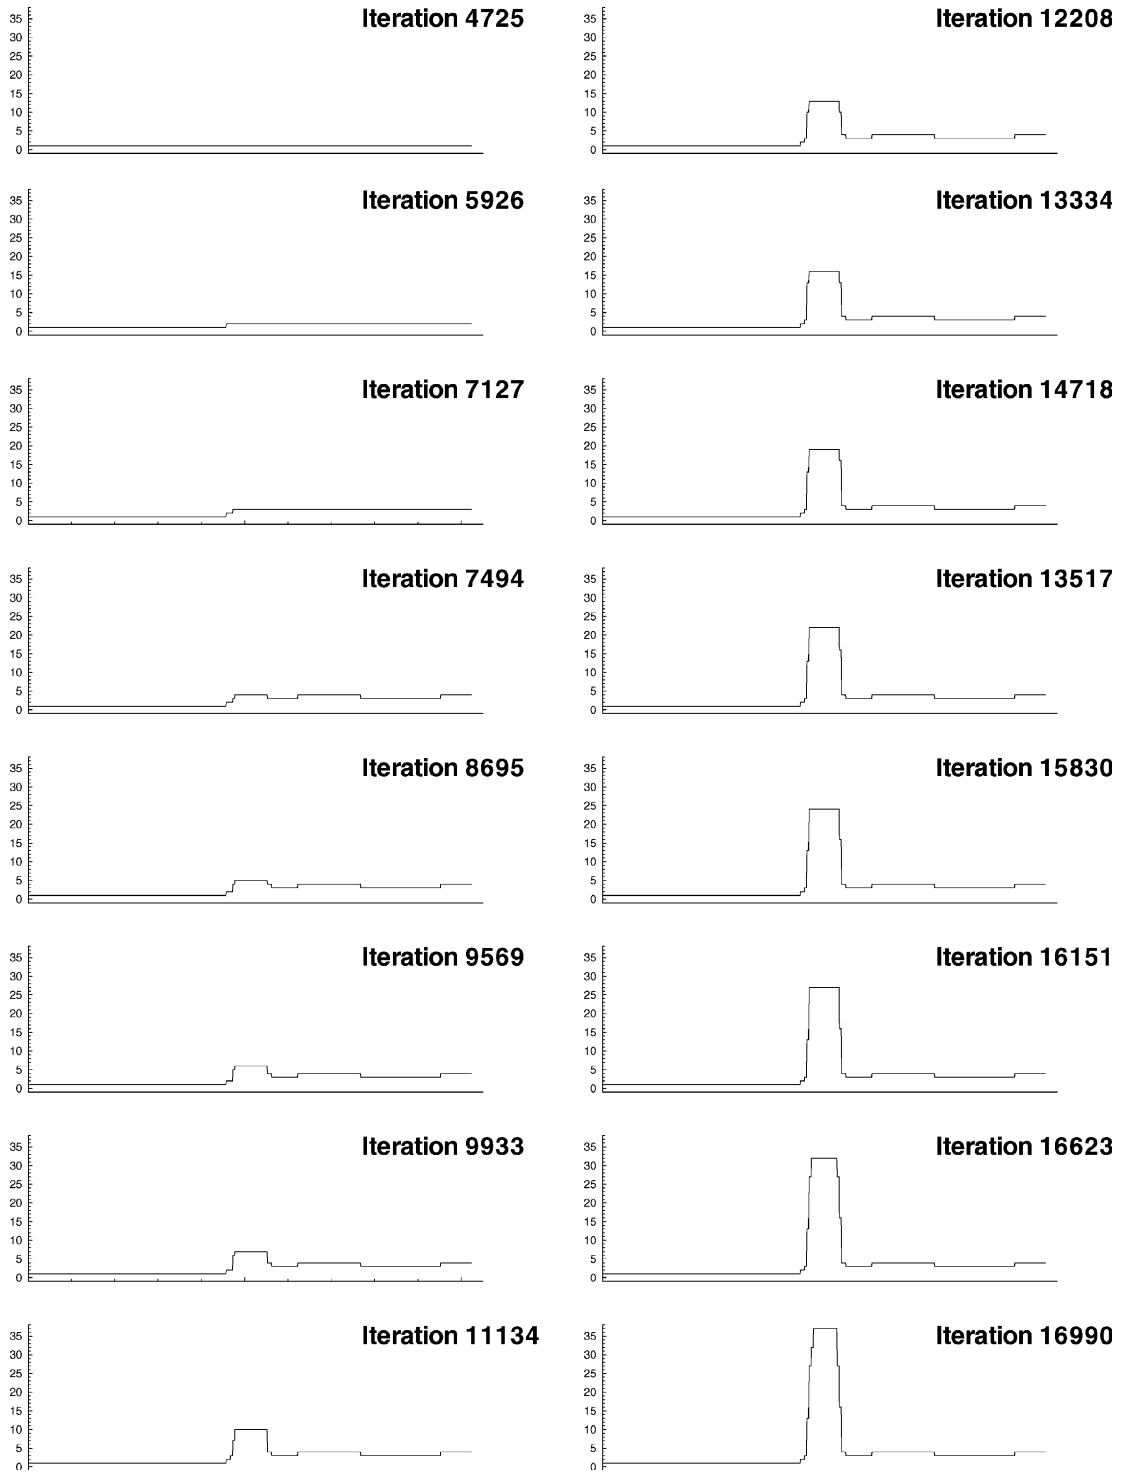


Fig. 6. Evolution of the adapted frequency distribution, 1024 cells, augmentation threshold $5.0e-7$.

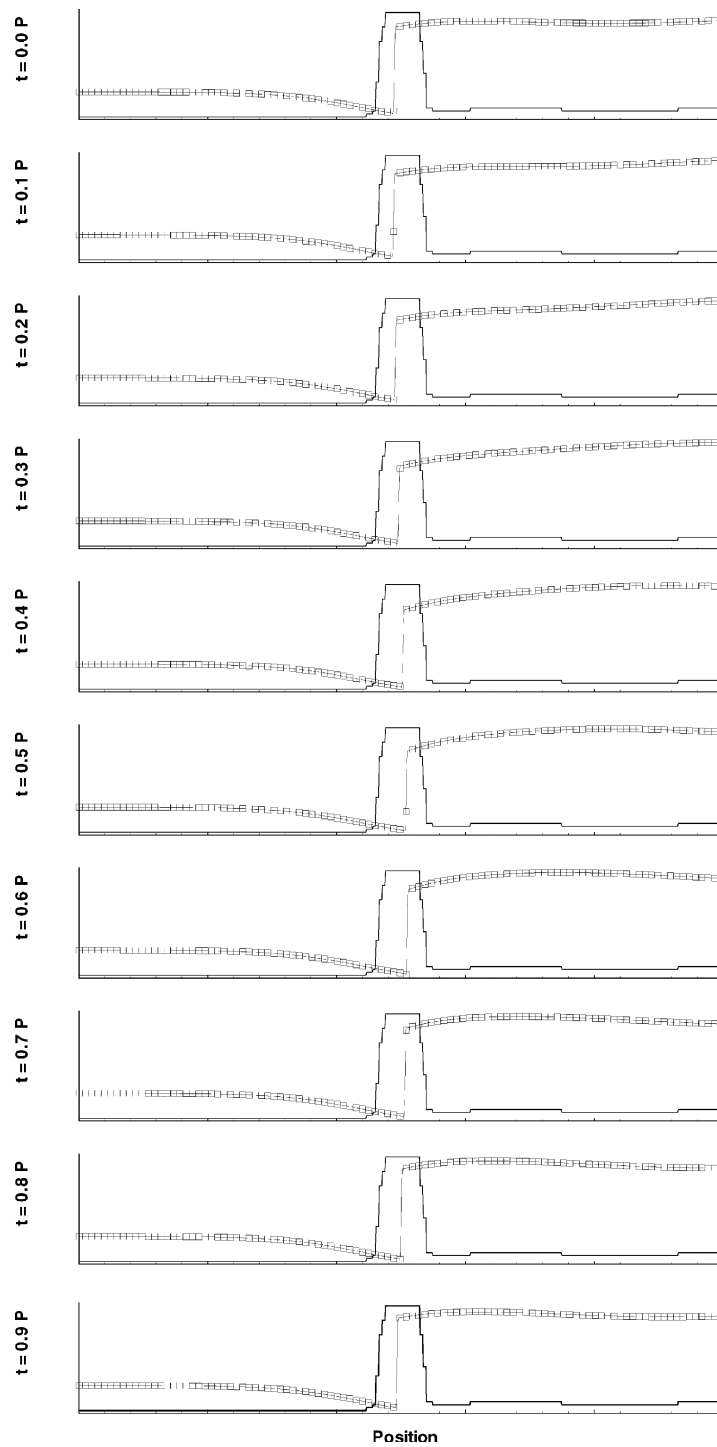


Fig. 7. Comparison of adaptive harmonic balance (symbols) and time-accurate solutions (solid) at 10 points in time spanning one period of oscillation, with adapted frequency distribution (heavy solid).




Fig. 8. Root-mean-square of difference between adaptive harmonic balance and time-accurate solutions with changing adaptation threshold.

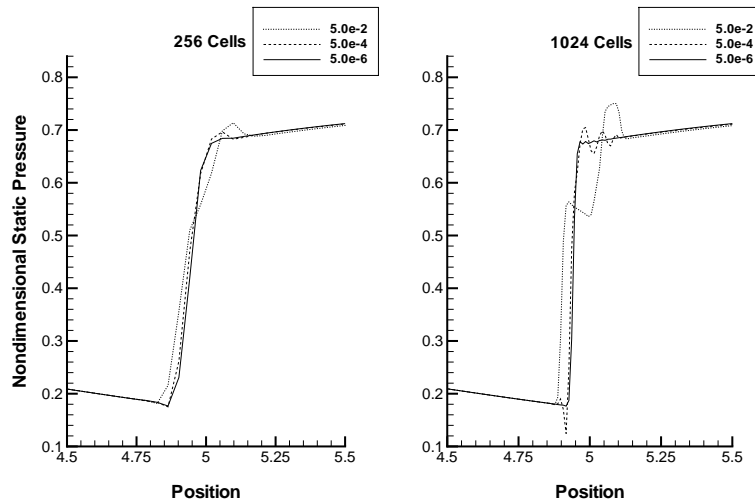


Fig. 9. Time snapshot of adaptive harmonic balance solution illustrating refinement of oscillating normal shock with decreased augmentation threshold.

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, the maximum frequency content grew by 35 frequencies at the smallest threshold, while the average frequency content grew by approximately 4.5 frequencies

Low average frequency content translated into reduced compute times, as shown in Fig. 13. A comparison of Figs. 12 and 13 shows that the compute time grew more 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,

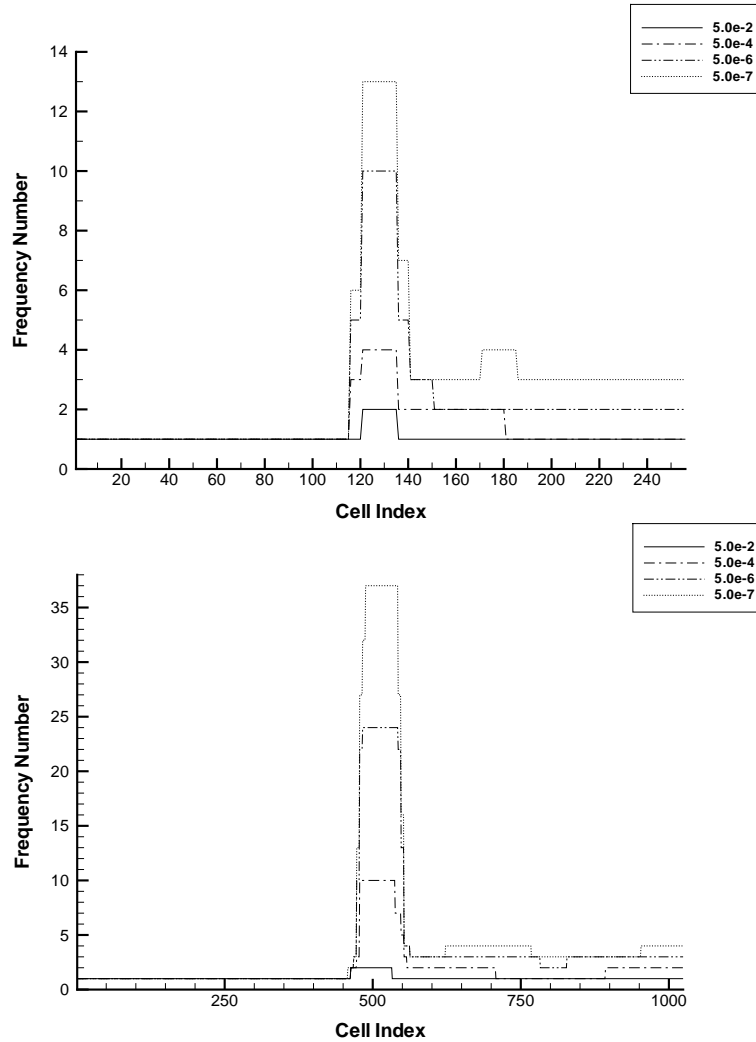


Fig. 10. Adapted frequency distributions for different augmentation thresholds on coarse grid (top) and fine grid (bottom).

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 (fixed-frequency) harmonic balance, each of the harmonic balance solutions was recomputed with a fixed number of frequencies. The maximum frequency content from each of the adapted solutions was 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 is the number of frequencies included in the solution. The equivalent problem size for the adapted calculation is the average frequency content. The difference between these two quantities is shown in Fig. 14 as a percent reduction.

For the largest augmentation threshold, the reduction in problem size was approximately 45% for all grid densities, which represents a significant reduction, considering that the average number of frequencies

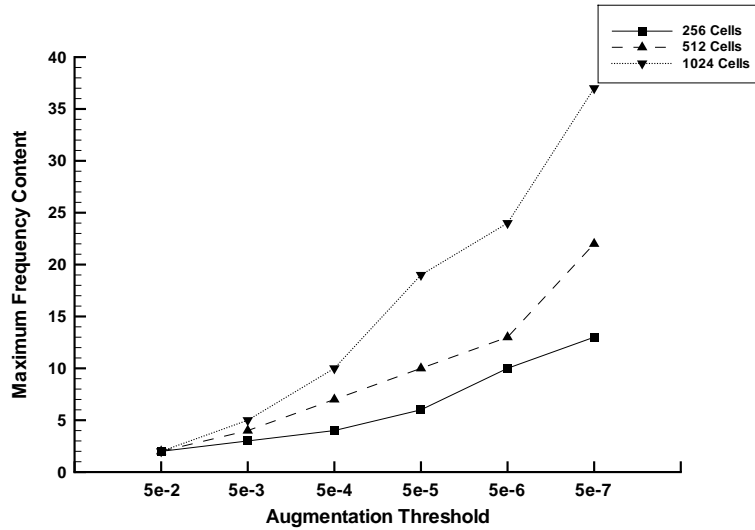


Fig. 11. Change in maximum frequency content of adapted harmonic balance solutions with changes in augmentation threshold.




Fig. 12. Change in average frequency content of adapted harmonic balance solutions with changes in augmentation threshold.

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. The problem size reduction was more substantial for the small thresholds, approaching 90% for the smallest threshold on the finest grid.

The difference between compute times for the adapted and fixed-frequency solutions is shown in Fig. 15, again in terms of percent reduction. To determine how problem size reduction (Fig. 14) translated into run time reduction (Fig. 15) with changes in augmentation threshold and grid density, the ratio of compute time to problem size was calculated.




Fig. 13. Increase in compute time with decreasing augmentation threshold.




Fig. 14. Reduction in average frequency content for adapted vs non-adapted harmonic balance solutions.

At a threshold of $5.0e-2$ the overhead of the adaptation algorithm was significant relative to the maximum absolute run-time reduction on all grids, resulting in an average ratio (converted to percent) of 47%. For lower thresholds the overhead was less significant, and ratios were higher, averaging 91% for a threshold of $5.0e-6$, and 95% for a threshold of $5.0e-7$. The adaptive solution on the finest grid with an augmentation threshold of $5.0e-7$ achieved a ratio of 99% with an 86% reduction in compute time. The adaptive ratios for all solutions are shown in Fig. 16.

To investigate the potential benefit of applying the adaptive technique to typical problems currently addressed with non-adaptive harmonic balance approaches, the time-variation of integrated pressure loads for adapted and non-adapted solutions were compared. Fig. 17 shows loads for non-adapted solutions with one, three, and five frequencies (typical of current harmonic balance applications) compared to the load for

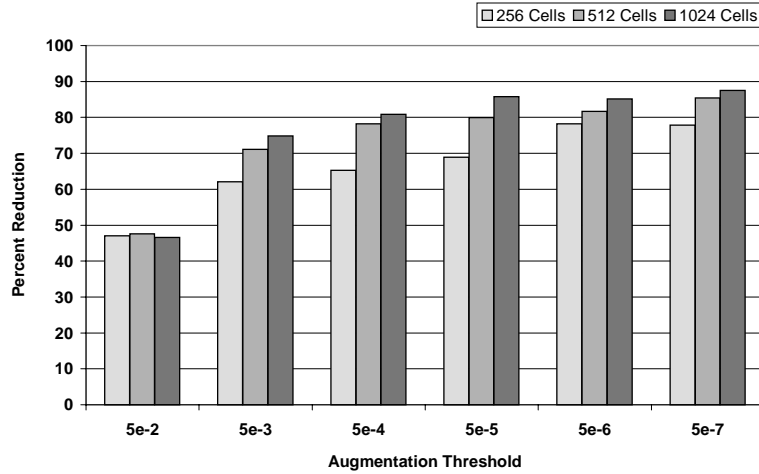


Fig. 15. Reduction in compute time for adapted vs non-adapted harmonic balance solutions.

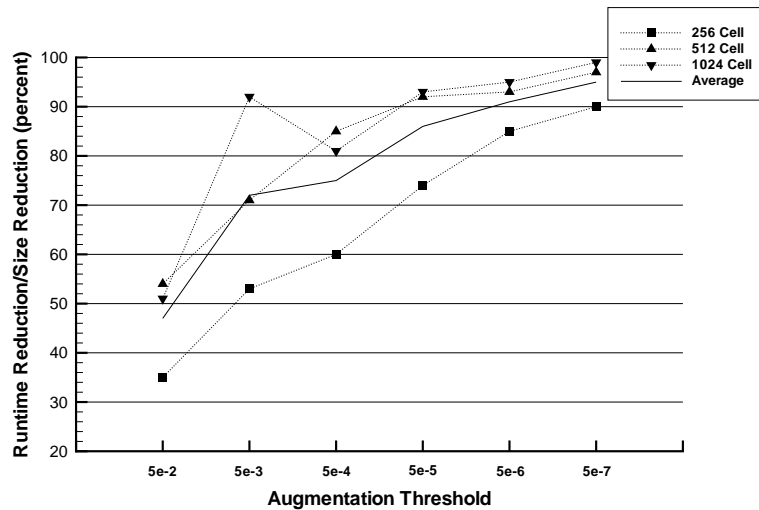


Fig. 16. Ratio of run time reduction to problem size reduction for adaptive harmonic balance solutions.

the high-fidelity adapted solution with a threshold of $5.0e-7$. All of the constant-frequency solutions compared favorably, though the 1-frequency solution showed a distinct phase shift relative to the other solutions. Maximum error ranged from 0.1% for the 1-frequency solution to 0.01% for the 5-frequency solution. Relative to the adapted solution, the 1, 3, and 5-frequency solutions took 0.37, 0.7, and 1.7 times as long to compute, respectively. To show the effect of increased E_{thresh} , the integrated load for the adapted solution with a threshold of $5.0e-3$ is also plotted. This solution had a maximum error approximately the same as the 3-frequency non-adapted solution ($\approx 0.05\%$), but took 25% less time to compute.

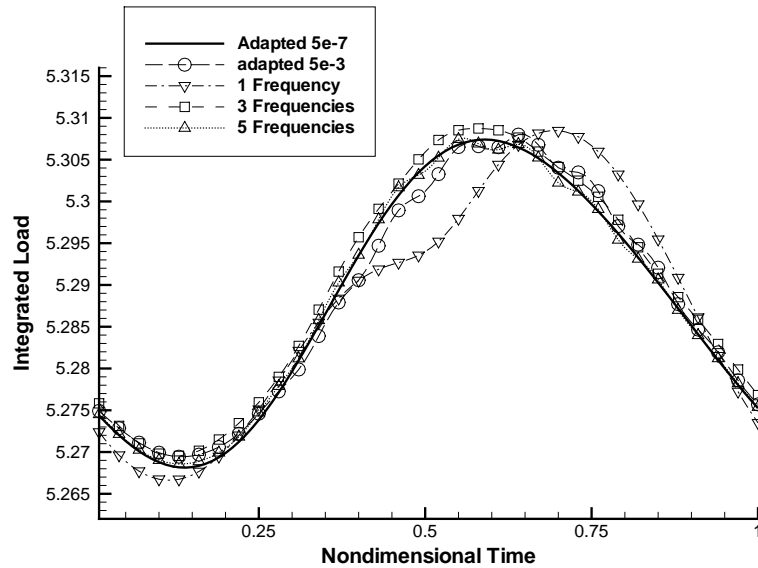


Fig. 17. Variation in time of integrated pressure for adapted fixed-frequency harmonic balance solutions on 1024-cell grid.

6. Conclusions

An adaptive harmonic balance method was presented and applied to an unsteady supersonic/subsonic diverging nozzle. This problem contained a periodic unsteady flow field with mostly continuous, low-frequency content flow, but with an isolated shock region with highly nonlinear, high-frequency content flow.

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 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 $5.0e-6$ and $5.0e-7$ averaged greater than 80% reductions in compute time relative to equivalent non-adapted solutions, on all grids. This reduced the difference in compute time between low-fidelity and high-fidelity solutions so that the highest fidelity solutions ran only 2–2.5 times longer than the lowest-fidelity solutions. When compared to the observed advantage of low-fidelity harmonic balance solutions over time-accurate calculations, it is clear that the adaptive harmonic balance method will allow higher-fidelity solutions (when more than one harmonic is important) at the same or reduced cost, and may be competitive with time-accurate calculations for some high-fidelity solutions.

Acknowledgements

The authors would like to thank the Dayton Area Graduate Studies Institute for partially funding this work under DAGSI Grant PR-AFIT-99-07.

References

- [1] J.M. Verdon, M.D. Montgomery, A.H. Chuang, Development of a linearized unsteady Euler analysis with application to wake/blade–row interactions. NASA Contractor Report NASA/CR 1999-208879, NASA, November 1999.
- [2] D. Hoyniak, W.S. Clark. Aerodynamic damping predictions using a linearized Navier–Stokes analysis, in: International Gas Turbine and Aeroengine Congress and Exhibit, Indianapolis IN, June 7–10, 1999, ASME 99-GT-207.
- [3] W.S. Clark, K.C. Hall, A time-linearized Navier–Stokes analysis of stall flutter, in: International Gas Turbine and Aeroengine Congress and Exhibit, Indianapolis IN, June 7–10, 1999, ASME 99-GT-383.
- [4] R. Florea, K.C. Hall, Sensitivity analysis of unsteady inviscid flow through turbomachinery cascades, in: 38th AIAA Aerospace Sciences Meeting and Exhibit, Reno NV, January 10–13, 2000, AIAA 2000-0130.
- [5] W. Ning, L. He, Computation of unsteady flows around oscillating blades using linear and nonlinear harmonic Euler methods, *J. Turbomach.* 120 (1998) 508–514.
- [6] L. He, W. Ning, Efficient approach for analysis of unsteady viscous flows in turbomachines, *AIAA J.* 17 (11) (1998) 2012–2012.
- [7] T. Chen, P. Vasanthakumar, L. He, Analysis of unsteady blade row interaction using nonlinear harmonic approach, *AIAA J. Propul. Power* 17 (3) (2001) 651–658.
- [8] K.C. Hall, J.P. Thomas, W.S. Clark, Computation of unsteady nonlinear flows in cascades using a harmonic balance technique, *AIAA J.* 40 (5) (2002) 878–886.
- [9] M. McMullen, A. Jameson, J.J. Alonso, Acceleration of convergence to a periodic steady state in turbomachinery flows, in: 39th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, January 8–11, 2001, AIAA 2001-0152.
- [10] J.P. Thomas, E.H. Dowell, K.C. Hall, Nonlinear inviscid aerodynamic effects on transonic divergence, flutter, and limit cycle oscillations, *AIAA J.* 40 (4) (2002) 638–646.
- [11] M. McMullen, A. Jameson, J.J. Alonso, Application of a non-linear frequency domain solver to the Euler and Navier–Stokes equations, in: 40th AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, January 14–17, 2002, AIAA 2002-0120.
- [12] R.C. Maple, P.I. King, P.D. Orkwis, Split-domain harmonic balance solutions to burger’s equation for large amplitude disturbances, *AIAA J.* 41 (2) (2003) 206–212.
- [13] R.C. Maple, P.I. King, M.E. Oxley, Adaptive harmonic balance solutions to Euler’s equation, in: 32nd AIAA Fluid Dynamics Conference and Exhibit, St. Louis, MO, June 24–26, 2002, AIAA 2002-2953.
- [14] G.R. Shubin, A.B. Stephens, H.M. Glaz, Steady shock tracking and newton’s method applied to one-dimensional duct flow, *J. Comput. Phys.* 39 (1981) 364–374.
- [15] G. Strang, On the construction and comparison of difference schemes, *SIAM J. Num. Anal.* 5 (3) (1968) 506–517.
- [16] E.F. Toro, *Fiemann Solvers and Numerical Methods for Fluid Dynamics*, Springer, New York, 1997 (Chapter 15).
- [17] R.J. LeVeque, H.C. Yee, A study of numerical methods for hyperbolic conservation laws with stiff source terms, *J. Comput. Phys.* 86 (1990) 187–210.
- [18] D.A. Schwer, W.H. Green Jr., Split-operator methods for computing steady-state reacting flow fields, in: 15th AIAA Computational Fluid Dynamics Conference, Anaheim, CA, June 11–14 2001, AIAA 2001-2635.
- [19] A. Jameson, W. Schmidt, E. Turkel, Numerical solutions of the Euler equations by finite volume methods using Runge–Kutta time-stepping schemes, in: 14th AIAA Fluid and Plasma Dynamics Conference, Palo Alto, CA, June 23–25, 1981, AIAA 81-1259.
- [20] R. Swanson, E. Turkel, Artificial dissipation and central difference schemes for the Euler and Navier–Stokes equations, in: 8th Computational Fluid Dynamics Conference, Palo Alto, CA, June 23–25, 1987, AIAA 87-1107.
- [21] R. Enander, Implicit explicit residual smoothing for the multidimensional Euler and Navier–Stokes equations, *SIAM J. Sci. Comput.* 18 (5) (1997) 1243–1254.
- [22] M. Frigo, A fast fourier transform compiler, in: Proceedings of the 1999 ACM SIGPLAN Conference on Programming Language Design and Implementation, 1999.
- [23] R.C. Maple, Multigrid for adaptive harmonic balance cfd, in: 16th AIAA Computational Fluid Dynamics Conference, Orlando, FL, June 23–26, 2003, AIAA 2003-3434.
- [24] K.A. Hoffman, S.T. Chiang, *Computational Fluid Dynamics*, vol. II, fourth ed., Engineering Education system, Wichita, 2000, pp. 133–151.
- [25] M.J. Zucrow, J.D. Hoffman, in: *Gas Dynamics*, vol. I, Wiley, New York, 1976, pp. 162–178, see also pp. 335–351.