

Adaptive Harmonic Balance Solutions to Euler's Equation

Raymond C. Maple,* Paul I. King,† and Mark E. Oxley‡

Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio 45433

A new adaptive harmonic balance method is presented and applied to a variety of subsonic and supersonic one-dimensional flowfields containing strong moving shocks. The adaptive method augments 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 that cannot adequately resolve high-frequency terms. A study of the effect of augmentation thresholds, increments, and scheduling on the performance and accuracy of the adaptive method is presented, and optimal parameters identified. When optimal parameters are used, the new adaptive harmonic balance method produces solutions equivalent to a nonadapted harmonic balance solution, but with up to a 50 % reduction in run time.

Nomenclature

\tilde{a}	=	unsteady amplitude
C_n	=	vector of Fourier coefficients for the dependent variable Q , n th Fourier frequency
E_N	=	fraction of spectral energy in highest computed Fourier frequency
E_t	=	total energy, $\rho(e + \frac{1}{2}u^2)$
E_{thresh}	=	frequency augmentation threshold
e	=	specific internal energy
F	=	vector of flux terms for one-dimensional Euler's equation
\hat{F}	=	vector of harmonic balance flux terms for one-dimensional Euler's equation
\bar{F}	=	inverse Fourier transform of \hat{F}
M	=	Mach number
N	=	number of Fourier frequencies in an approximating series
n	=	Fourier frequency number
p	=	static pressure
Q	=	vector of dependent variables for one-dimensional Euler's equation
\hat{Q}	=	dependent variable vector for one-dimensional harmonic balance Euler's equation
\bar{Q}	=	inverse Fourier transform of \hat{Q}
R	=	solution residual, L_2 norm of change in solution over one time step
\hat{S}	=	vector of harmonic balance source terms for one-dimensional Euler's equation
t	=	time variable
u	=	x velocity
x	=	spatial variable
ρ	=	density
τ	=	pseudotime variable
ω	=	fundamental Fourier frequency

Introduction

UNSTEADY time-periodic flowfields are found in such applications as blade-vane interactions in turbomachinery and aero-

elastic limit-cycle oscillations in internal and external flows. Calculations of these flows have typically fallen into one of two categories, time-accurate solutions and linearized solutions. Time-accurate solutions capture nonlinear effects with no restrictions on unsteadiness, but can require considerable computation time. Linearized solutions, which are solved with steady-state techniques, are obtained quickly, but are limited to small-amplitude unsteadiness and do not capture nonlinear effects.

Recently, methods have been developed to speed up the calculation of problems with significant nonlinear effects. Hall et al.¹ introduced the harmonic balance method to solve nonlinear systems containing a finite number of fully coupled harmonics with no small perturbation assumption. In their method, the dependent variables are recast in the form of a truncated Fourier series in time. The time derivatives are replaced by harmonic frequency source terms, converting the original unsteady problem to a steady-state problem and allowing the application of convergence acceleration techniques.

Good harmonic balance solutions can be obtained with six or fewer frequencies in the approximating Fourier series if the motion of strong discontinuities is limited and the primary interest is in the time average and first harmonic of the flow solution.¹⁻⁴ However, if the flow contains strong traveling discontinuities and higher harmonics are important, a large number of series terms can be required.

The ability of the harmonic balance method to generate solutions containing strong moving discontinuities was examined by Maple et al.⁵ Accurate solutions for the inviscid Burgers equation were obtained if a sufficient number of Fourier frequencies were included in the approximating solution. It was noted, however, that the number of frequencies required for a good solution varied throughout the grid. In regions where the unsteadiness was continuous, few frequencies were necessary, whereas in regions with moving discontinuities, many more were required.

This paper will show that it is unnecessary to maintain a large number of Fourier frequencies throughout a grid to capture flow details in localized regions of strong moving discontinuities. A new adaptive technique is presented in which the number of Fourier frequencies at each grid cell is automatically varied to capture the local flow efficiently. The technique is demonstrated in the solution of the one-dimensional Euler equation applied to a variety of periodic flowfields containing strong moving discontinuities.

Governing Equations

The one-dimensional Euler equation is given in differential vector form by

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

Received 23 September 2002; revision received 12 April 2003; accepted for publication 12 April 2003. This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0001-1452/03 \$10.00 in correspondence with the CCC.

*Assistant Professor, Department of Aeronautics and Astronautics. Senior Member AIAA.

†Associate Professor, Department of Aeronautics and Astronautics. Senior Member AIAA.

‡Professor, Department of Mathematics and Statistics.

where

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

The perfect gas relation and assumption of constant ratio of specific heats closes the system of equations.

Equation (1) can be transformed into harmonic balance form.^{1,5} The basis of this transformation is the approximation of the vector of conservative variables, \mathbf{Q} , by a truncated Fourier series,

$$\mathbf{Q} \approx \sum_{n=-N}^N \mathbf{C}_n(x) e^{in\omega t} \quad (3)$$

where $i = \sqrt{-1}$, ω is the fundamental radian frequency of the series, and $\mathbf{C}_n(x)$ are vectors of complex Fourier coefficients that vary in space, but not time. Substituting Eq. (3) into Eq. (1), evaluating the time derivative and flux function \mathbf{F} , and collecting terms of like exponent results in the steady state equation

$$\frac{d\hat{\mathbf{F}}}{dx}(x) + \hat{\mathbf{S}}(x) = 0 \quad (4)$$

The vector $\hat{\mathbf{F}}$ contains harmonic balance flux terms, which are functions of $\hat{\mathbf{Q}}$, a concatenation of the Fourier coefficient vectors \mathbf{C}_n , for $0 \leq n \leq N$. The expressions for the components of $\hat{\mathbf{F}}$ are quite complex, consisting of gathered terms from multiple products of the series expansions for the conservative variables. In addition, evaluation requires the solution of a series representing the reciprocal of density. See Hall et al.¹ for a derivation of these terms. The vector $\hat{\mathbf{S}}$ arises from the time derivative in Eq. (1) and comprises the concatenation of vectors of the form $in\omega \mathbf{C}_n$ for $0 \leq n \leq N$. Because the original vector of conserved variables, \mathbf{Q} , is real, the Fourier coefficients corresponding to the negative frequencies in Eq. (3) are simply the complex conjugates of those for the positive frequencies, and only terms with $n \geq 0$ are retained.

Equation (4) represents a steady-state system of $2N + 1$ equations for $2N + 1$ unknowns. To allow a time-marching numerical solution method, a pseudotime derivative of the vector $\hat{\mathbf{Q}}$ can be added, which yields

$$\frac{\partial \hat{\mathbf{Q}}(x, \tau)}{\partial \tau} + \frac{\partial \hat{\mathbf{F}}(x, \tau)}{\partial x} + \hat{\mathbf{S}}(x, \tau) = 0 \quad (5)$$

When a steady-state solution to Eq. (5) is reached, the added term goes to zero, and Eq. (4) is recovered.

Adaptive Split-Domain Harmonic Balance Solver

A Jameson–Schmidt–Turkel (JST)⁶-type cell-centered finite volume solver was written to solve Eq. (5). This solver is second-order accurate in space. To avoid oscillations at discontinuities, modified JST artificial dissipation^{6,7} was used in conjunction with globally applied second-order dissipation. The solver incorporated the split-domain harmonic balance approach of Maple et al.,⁵ that is, Eq. (5) was split into a homogeneous partial differential equation (PDE) containing the advection terms and an ordinary differential equation (ODE) containing the source term. The homogeneous PDE was transformed into the time domain to allow efficient calculation of the fluxes:

$$\frac{\partial \bar{\mathbf{Q}}}{\partial \tau} + \frac{\partial \bar{\mathbf{F}}}{\partial x} = 0 \quad (6a)$$

$$\frac{d\bar{\mathbf{Q}}}{d\tau} + \hat{\mathbf{S}} = 0 \quad (6b)$$

The vector $\bar{\mathbf{Q}}$ consists of $2N + 1$ instances of the original dependent vector $\mathbf{Q}(x, t)$, 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. It is related to the vector $\hat{\mathbf{Q}}$ through the discrete Fourier transform, for example, the ρ components of the vector $\bar{\mathbf{Q}}$ are the inverse discrete Fourier transform of the ρ components of the coefficients \mathbf{C}_n in $\hat{\mathbf{Q}}$. The flux vector $\bar{\mathbf{F}}$ is the standard Euler flux vector evaluated at the $2N + 1$ values of $\bar{\mathbf{Q}}$ making up the vector $\bar{\mathbf{Q}}$. Numerically, $\bar{\mathbf{F}}$ approximates the inverse discrete Fourier transform of the harmonic balance flux vector $\hat{\mathbf{F}}$. Equation (6a) represents $2N + 1$ independent instances of the original Euler equation with time replaced by pseudotime. The independent variables have been omitted for brevity.

The split-domain harmonic balance solution was advanced in pseudotime by integrating Eq. (6b) half a pseudotime step, transforming the intermediate solution to the time domain, integrating Eq. (6a) for one full pseudotime step, and finally transforming the result back to the frequency domain, where Eq. (6b) was integrated another half pseudotime step. This approach, often referred to as a Strang symmetric splitting approach,⁸ required only one fast Fourier transform (FFT)/inverse FFT calculation per cell, per iteration, and significantly improved the stability of the solver (compared to a nonsplit approach) when large numbers of Fourier frequencies were included in the solution.

Each of the equations in the system represented by Eq. (6b) is independent, with the form

$$\frac{d\mathbf{C}_m}{d\tau} + im\omega \mathbf{C}_m = 0, \quad m = 0, 1, \dots, N \quad (7)$$

The exact solution to these equations is easily obtained, with the updated value of the m th coefficient vector given by

$$\mathbf{C}_m^{n+1} = \mathbf{C}_m^n e^{-im\omega \Delta \tau} \quad (8)$$

In practice, it was found that solving the ODE (6b) exactly slowed or stalled convergence to steady state. Convergence properties were improved by solving the ODE numerically.

Pseudotime integration for both Eqs. (6a) and (6b) was accomplished with a three-stage Runge–Kutta (RK) scheme with good high-frequency smoothing properties and minimal storage requirements.⁹ Local time stepping was employed to accelerate convergence. The pseudo-time step in each cell was the smallest maximum step across all computed samples. The maximum stable Courant–Friedrichs–Lewy number (CFL) for the solver was 1.7.

Just as some regions in a flowfield may require more cells or grid points to capture adequately the flow physics, some regions also require more frequencies in a harmonic balance solution. Analogous to grid adaptation, which adjusts the number of grid points in a region to better resolve the flow, the adaptive harmonic balance approach automatically adjusts the number of Fourier frequencies included in the solution to resolve the flow at a given cell, on a cell-by-cell basis.

Frequency Augmentation

In the current implementation, frequency adaptation is accomplished by means of frequency augmentation. With this approach, the solution is begun with a user-specified initial number of frequencies. As the solution develops, frequencies are added in fixed increments to individual grid cells until a final frequency distribution and solution are obtained. Such an approach requires a reliable indicator of solution fidelity at each computational cell. Because 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. A reliable indicator is the fraction of spectral energy contained in the highest computed Fourier frequency, given by

$$E_N = \frac{|\mathbf{C}_N|^2}{\sum_{n=0}^N |\mathbf{C}_n|^2} \quad (9)$$

where all mathematical operations are performed elementwise, so that E_N contains the fraction of spectral energy of the highest computed frequency for each variable expanded in a series (ρ , ρu , etc.).

For the flowfields of interest, discontinuities in the time response are localized and do not take the form of pure impulses, and so E_N decreases as N is made larger. The use of E_N as an adaptation metric is dependent on this property.

The decision to augment frequencies at a cell was made by comparing each element of 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 five-point spatial average before thresholding. When any element of the smoothed E_N exceeds the threshold, additional frequencies are incorporated into the solution at that cell. Selection of the threshold value was based on the desired solution fidelity.

Once a cell was identified for augmentation, its frequency content was increased by a predetermined increment, chosen to minimize execution time. The time required to solve the split-domain harmonic balance equations is composed mainly of two components, the time required to solve the time-domain equations (6a) and the time required to perform the necessary Fourier transforms.

The computational time for the time-domain equations increases linearly with N , and so there is no performance advantage favoring any particular frequency increment, as long as the increment is relatively small. The computational time 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 Fast Fourier Transform in the West (FFTW) library used in this effort yield nearly as good performance when the number of terms can be factored into powers of small prime numbers, that is, when $2N + 1 = 2^a 3^b 5^c 7^d 11^e$ for some positive integers a, b, c, d , and e (Ref. 10). Frequency augmentation increments were chosen (see subsequent discussion) to minimize the combined cost of the flux and FFT calculations.

Adaptation Scheduling

The goal of adaptation scheduling is to adapt the solution at the times that yield the best performance. The frequency augmentation approach is most efficient when initiated with solutions that are representative of the final solution. If adaptation is attempted too early in the solution process, unnecessary 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, too much time could be spent refining flow details that change when the solution is adapted, resulting in increased execution times. Therefore, adaptation should occur soon after the major flow features are developed.

In the current work, two adaptation triggers were 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

$$R = \sqrt{\frac{\sum_{i=1}^{ni} (\Delta \hat{Q}_i \cdot \widetilde{\Delta \hat{Q}_i})}{ni}} \quad (10)$$

where ni is the number of cells in the computational grid, $\Delta \hat{Q}$ is the change in the solution vector in one iteration, and the tilde indicates complex conjugation. Adaptation was triggered when $\log_{10}(R)$ dropped by a user-specified amount. In the rare cases where convergence stalls and the specified residual drop is not achieved, the secondary trigger, which is based on the number of iterations completed, initiates adaptation.

Both the residual drop and iteration count were measured relative to a reference residual and iteration number that were initially set after the first iteration of the solution and reset whenever the solution was adapted. The adaptation algorithm was implemented with two sets of trigger values: one for the initial adaptation and a second all subsequent adaptations. The initial trigger values were set to allow time for the solution to develop from a poor initial guess. Subsequent trigger values were set to allow errors introduced by frequency augmentation to be removed from the solution and to further refine the solution a small amount. Selection of adaptation trigger values is discussed later.

Resampling at Frequency Transitions

Discretization of Eq. (6a) requires the addition and subtraction of the vectors \hat{Q} (and \bar{F}) within a discretization stencil. This presents a problem for cells in the discretization stencil that have different sample rates (numbers of frequencies). Not only do the vectors have different lengths, but the elements correspond to the state of the flow at different points in time. Cell quantities must be resampled so that the sample rate is consistent across the discretization stencil.

Various resampling methods can be implemented, such as truncation/zero padding in the frequency domain and linear interpolation in the time domain. To downsample a vector with truncation/zero padding, a time-domain vector is Fourier transformed, truncated to the desired size, and transformed back to the time domain. To up-sample, a vector is transformed, zero padded to increase its length, and transformed back to the time domain. For a linear operation such as addition or subtraction, this approach is equivalent to performing the same linear operation directly in the frequency domain.

One drawback with frequency-domain truncation/zero padding is that interpolated values are not bounded by the original data. For small sample rates, interpolation sometimes results in nonphysical values. When this occurs, linear interpolation in the time domain can be applied. With linear interpolation, values are obtained at the required sample rate from a linear interpolation of the calculated values.

Frequency-domain truncation/zero padding was selected as the primary resampling method due to the superior quality of the solution at sample rate transitions. The superiority of the method over linear interpolation is shown in Fig. 1. Figure 1 contains plots of

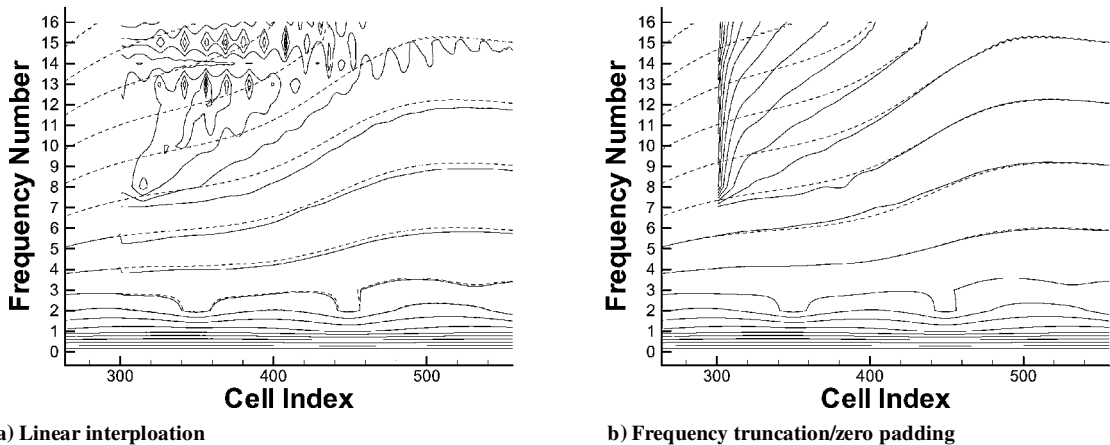


Fig. 1 Comparison of Fourier coefficient magnitude contours (momentum term) for different interpolation schemes at a transition from 7 to 16 frequencies with contours from a constant 16 frequency solution.

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). Differences in the computed Fourier coefficients, and thus in the reconstructed solution in the physical domain, are highlighted as differences in the contours. Also included are coefficient magnitude contours for a non-adapted solution with a constant 16 frequencies (dashed lines). Ideally, the contours for the solutions with a frequency transition will be the same as, or close to, these contours. As can be seen, the frequency-domain truncation/zero-padded solution is smoother and more closely matches the constant-sample-rate results than the linearly interpolated solution.

To achieve a smooth solution at frequency transitions, it was necessary to resample the transition boundaries at every RK integration stage. Attempts to time lag the transitions by freezing the resampled values at the first integration stage resulted in discontinuities at the transition cell, even when both sides of the transition had the same number of frequencies and no interpolation was 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. (6a) 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.

Results

Test Cases

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

$$M = 2.0 + \tilde{a} \sin(\omega t) \tag{11}$$

while maintaining a constant density and static pressure. The magnitude of the sinusoidal disturbance was varied to achieve various flow characteristics. Time-accurate pressure distributions at several snapshots in time for a flowfield with $\tilde{a} = 0.25$ are shown in Fig. 2. Because the resulting flow is hyperbolic in time with all flow information traveling in the downstream direction, these test cases provided a simplified environment for investigating frequency transitions.

The last test case consisted of a purely subsonic flowfield 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 $u(t) = \tilde{a} \sin(\omega t)$, a zero pressure gradient, and constant total enthalpy at a stationary (left) boundary. The open (right) end of the tube was modeled with a characteristic-variable far-field boundary condition.¹¹

These boundary conditions result in a complex, periodic, subsonic flowfield containing alternating right-running shock and expansion waves interacting with reflections from the open end of the tube. When the length of the tube and the magnitude and frequency of the imposed sinusoidal velocity are varied, the flowfield was adjusted so that a stationary periodic flowfield with moderately strong features was produced. Figure 3 contains time-accurate pressure plots for the subsonic test flowfield at several points in time throughout one period of the piston.

All harmonic balance solutions were converged to an overall residual R of $1.0e-9$, for a drop of approximately 6.5 orders of magnitude. For comparison purposes, all test cases were also solved with a validated time-accurate, second-order-in-space Roe solver.

The parameters of the supersonic test cases (SS1–SS4) and subsonic piston case (SU1) are summarized in Table 1.

Tuning the Adaptive Solver

Selection of Augmentation Threshold

The key parameter of the adaptive harmonic balance method is the augmentation threshold E_{thresh} . This threshold value indirectly

Table 1 Summary of test case parameters

Configuration	Unsteady amplitude	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	0.05π	5π	2.0

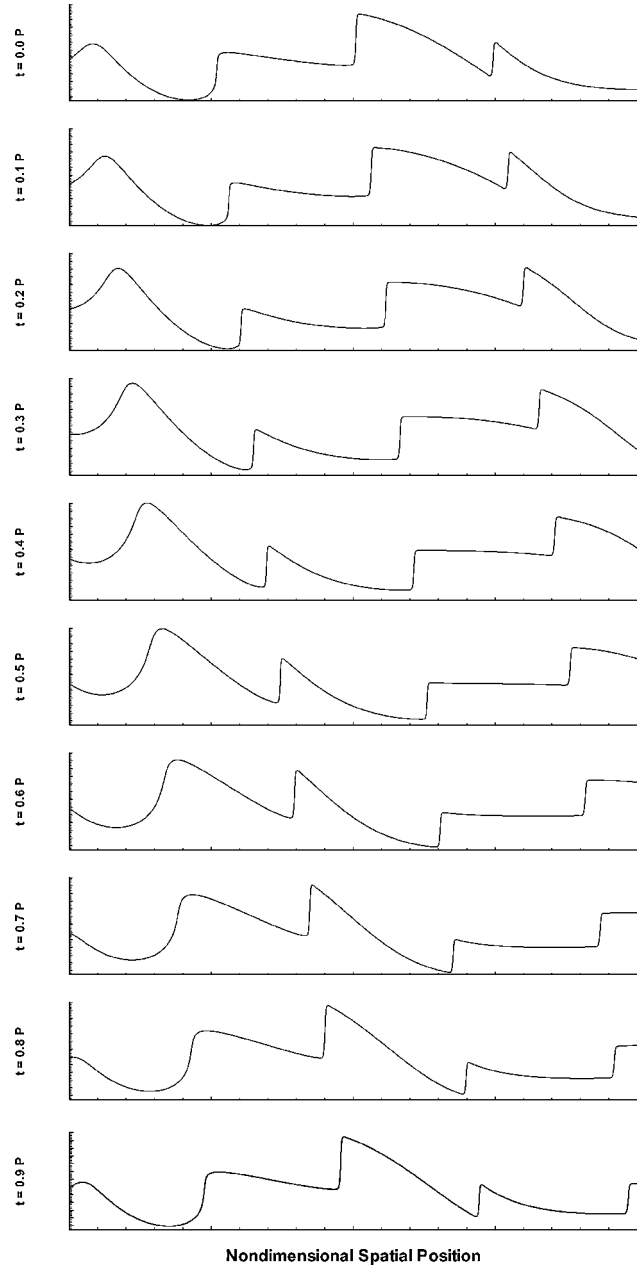


Fig. 2 Nondimensional static pressure for supersonic test case SS3, $M = 2.0 \pm 0.25$, generated with time-accurate Roe solver on 1000-cell grid, vertical axis range = 0.125–0.25.

controls the frequency distribution and, thus, the run time and accuracy of the resulting solution. To study the effect of E_{thresh} on the solution, and to determine a suitable range of values for E_{thresh} , solutions were generated for both the supersonic and subsonic test cases with thresholds ranging from $5.0e-4$ to $1.0e-10$.

The effect of E_{thresh} on problem size (the average frequency content of the solution) and compute time is shown in Fig. 4. For both cases, the average frequency content increased linearly with

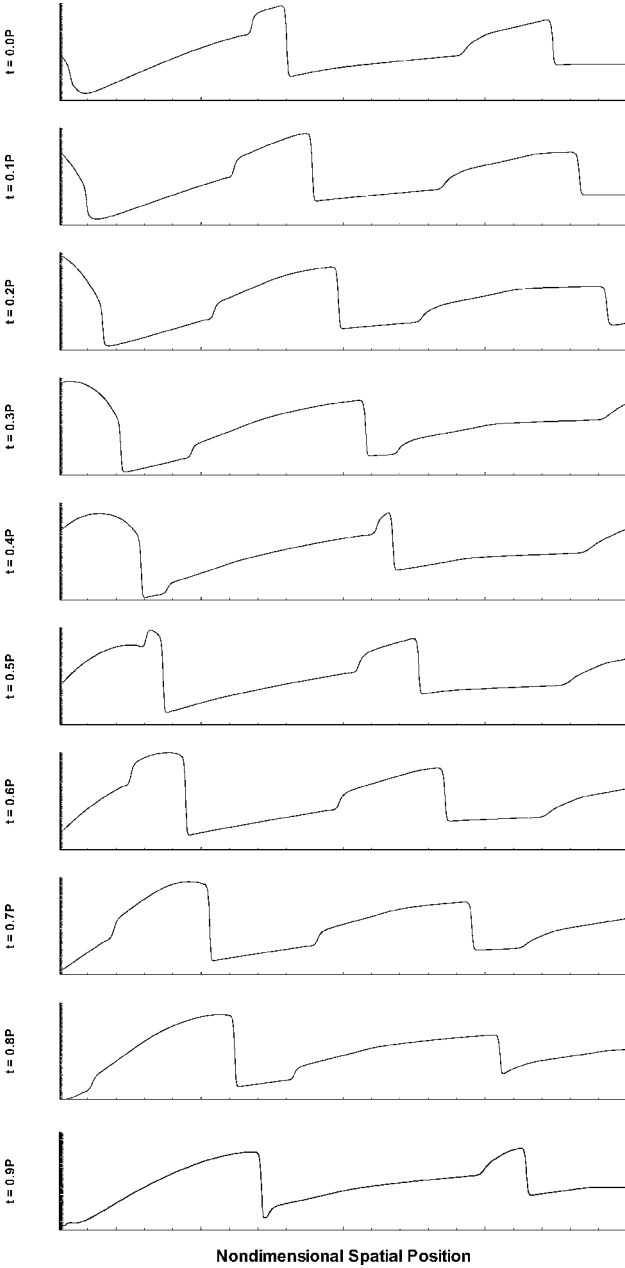


Fig. 3 Nondimensional static pressure for subsonic virtual piston test case SU2, generated with time-accurate Roe solver on 1000-cell grid, vertical axis range = 0.32–1.1.

decreases in $\log_{10}(E_{\text{thresh}})$. The size of the supersonic adapted harmonic balance problem grew approximately 8 times from the largest threshold to the smallest, whereas the size of the subsonic problem grew approximately 3.5 times. Compute 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 (Table 2).

An example of the change in frequency content and solution quality with changes in augmentation threshold is shown in Figs. 5–8. Figures 5–8 show pressures and frequency distributions for cases SS3 and SU1, respectively, at augmentation thresholds of $1.0e-4$, $1.0e-7$, and $1.0e-10$ on 500 cell grids. In Figs. 5 and 7, pressures from the harmonic balance solutions were reconstructed from the computed Fourier coefficients at time $t=0$ relative to the disturbance period. Note that the reconstruction could just as easily have been generated for any time with equal fidelity.

Table 2 Maximum Fourier frequency and stable CFL for case SU1, 500 grid cells, for decreases in augmentation threshold

E_{thresh}	Maximum N	Maximum 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

^aMaximum allowed for run.

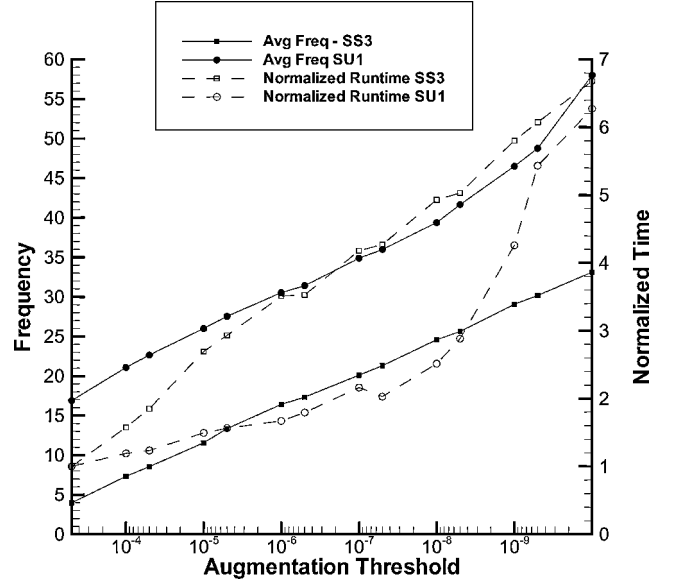


Fig. 4 Change in average frequency content and run time with changing augmentation threshold for cases SS3 and SU1 on 500-cell grid.

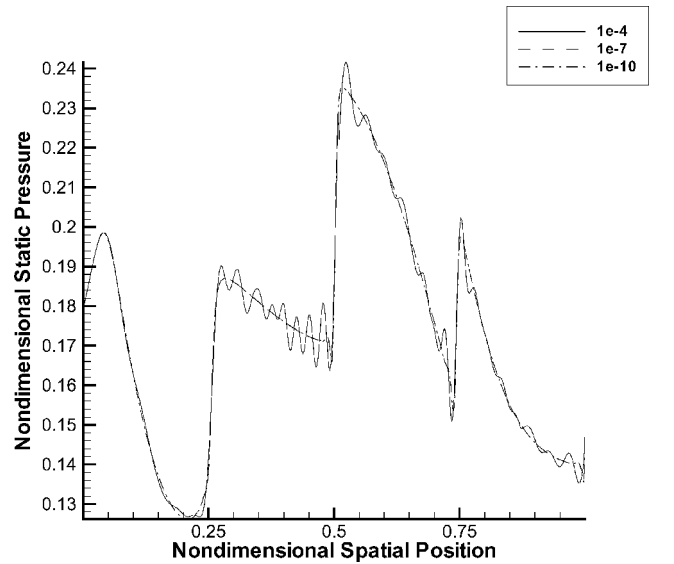


Fig. 5 Nondimensional static pressure for adapted solutions at three different augmentation thresholds, case SS3, 500 grid cells.

Figure 6 shows adapted frequency distributions for case SS3. From Fig. 6, 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 Fig. 6, there was a significant increase and refinement of the frequency distribution and solution quality with the change from $1.0e-4$ to $1.0e-7$. For smaller thresholds, there was no perceptible change in the reconstructed pressures, despite a large increase in frequency content.

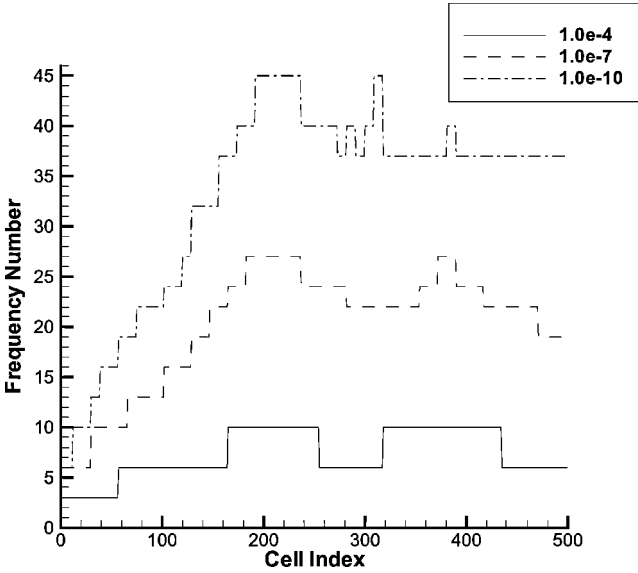


Fig. 6 Adapted frequency distributions for three different augmentation thresholds, case SS3, 500 grid cells.

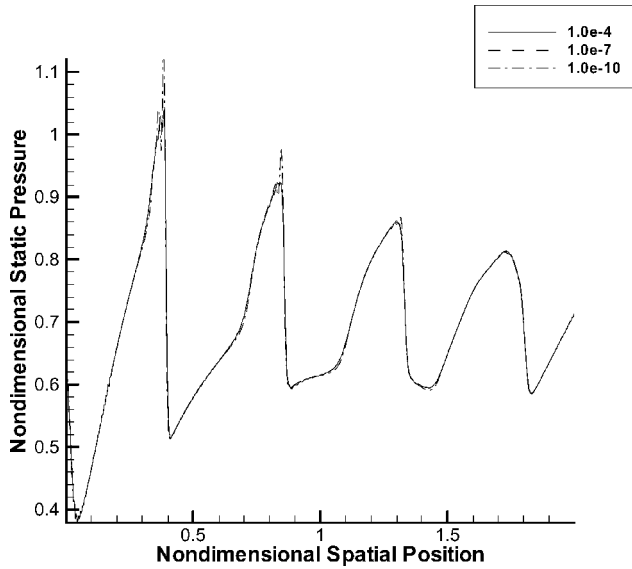


Fig. 7 Nondimensional static pressure for adapted solutions at three different augmentation thresholds, case SU1, 500 grid cells.

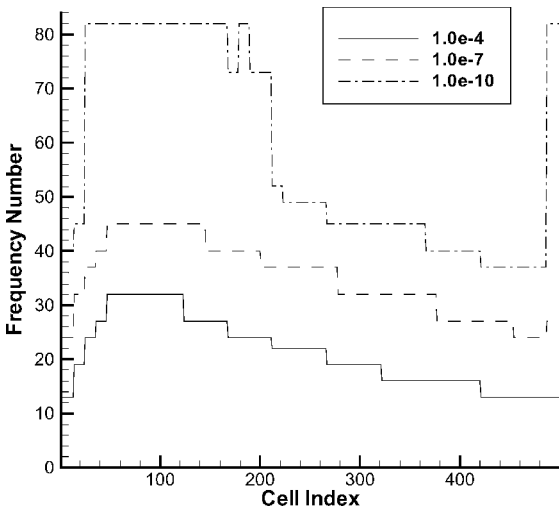


Fig. 8 Adapted frequency distributions for three different augmentation thresholds, case SU1, 500 grid cells.

The adapted frequency content of the subsonic case also increased significantly with each decrease in E_{thresh} (Fig. 8). Unlike the supersonic case, the subsonic pressure distributions for thresholds of $1.0\text{e-}4$ and $1.0\text{e-}7$ are nearly identical (Fig. 7). The increased frequency content resulting from the decrease in threshold did not improve the solution. In fact, a threshold decrease to $1.0\text{e-}10$ produced a degraded solution with high-frequency oscillations upstream of the strongest shocks.

The results for both the subsonic and supersonic cases are consistent with earlier findings of Maple et al.,⁵ who found that there is a frequency distribution for which further augmentation does not increase accuracy. Across all of the test cases, this distribution was attained for augmentation thresholds at or above $5.0\text{e-}8$, suggesting this value as a suitable lower bound for E_{thresh} .

The maximum acceptable value for E_{thresh} will vary widely depending on the flow being modeled, the grid used, and the desired solution quality. A conservative upper bound for a high-fidelity solution would be on the order of $5.0\text{e-}7$, although a larger threshold may be acceptable for some problems, as it was for case SU1, where a threshold of $1.0\text{e-}4$ produced comparable results (Fig. 7). If a lower-fidelity solution is acceptable, for example, only the time-average solution is of interest, then a much higher threshold would suffice.

Augmentation Increment Size

To find the frequency augmentation increments for best overall performance, the split-domain solver was run for a fixed number of iterations for every number of frequencies, N from 3 to 100. The execution times for each run were normalized by the run time for $N = 3$ and are shown in Fig. 9. For values of N that result in a poor factorization as already discussed, the cost of the FFTs dominates the overall run time. However, for values of N with good factorization, the normalized run time is nearly linear, which indicates that the cost of the FFTs is small relative to the cost of the flux calculation. These best-factored series lengths resulted in augmentation increments of 1 for $1 \leq N \leq 7$, approximately 3 for $7 \leq N \leq 30$, and approximately 5 for $30 \leq N \leq 100$.

Adaptation Trigger

A series of solutions was computed for various values of the residual-based adaptation triggers, κ_1 (initial) and κ_2 (subsequent). To ensure that augmentation was based on these triggers, the iteration-count-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

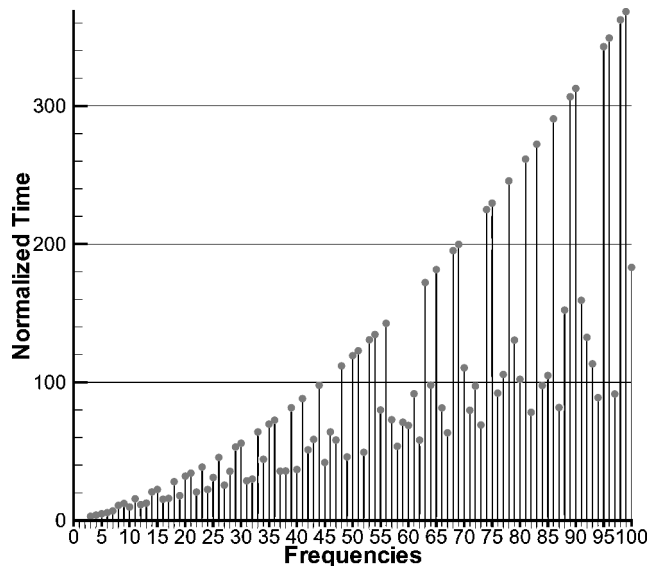


Fig. 9 Effect of series length on split-domain solver run time.

for supersonic and subsonic solutions. For the supersonic cases, it was best to adapt after only a small reduction in residual. The hyperbolic nature of the supersonic flowfields 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 before adding new frequencies. The optimum trigger values for supersonic cases were found to be $\kappa_1 = 0.25$ and $\kappa_2 = 0$. In comparison, the 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 coefficients of the lower Fourier frequencies were dominated by spatially low-frequency behavior and were slowest to converge. Adding higher Fourier frequencies introduced small errors with higher spatial frequency that were quickly removed by the solver, but had relatively little effect on the lower Fourier frequency coefficients. Thus, it was advantageous to refine the low-frequency solution before augmenting the frequencies. Conceptually, this is similar to a full multigrid scheme, in which a solution is first refined on a coarse grid before being transferred to a fine grid for further refinement.

The best run time for the subsonic case was 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.

Accuracy

Figures 10–13 address the accuracy of the adapted harmonic balance method. Figures 10 and 12 compare pressure distributions for adapted and nonadapted (fixed frequency) harmonic balance solutions with a conventional time-accurate solution at time $t = 0$ relative to the disturbance period for cases SS4 and SU1, respectively. Figures 11 and 13 contain isolevels of Fourier coefficient magnitude in the position-frequency plane for the ρu Fourier series for cases SS4 and SU1, respectively, along with the final adapted frequency envelope for each case. The contours in the lower frequencies correspond to larger coefficient magnitudes, whereas the contours in the

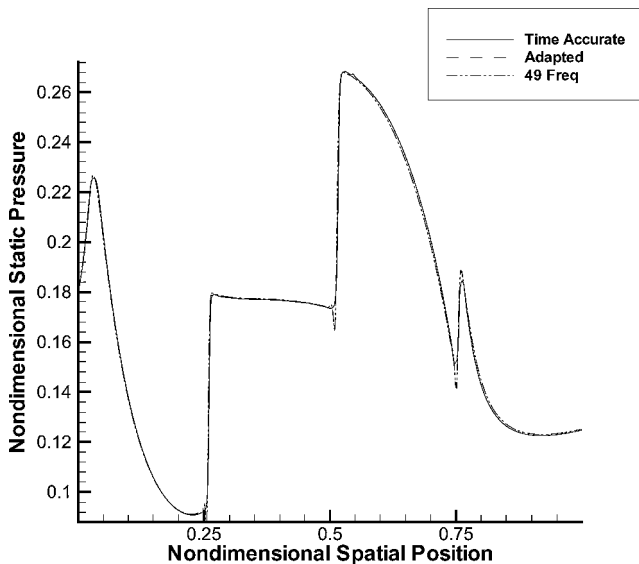


Fig. 10 Computed nondimensional static pressure for adapted (augmentation threshold $7.0e-8$) and nonadapted, 49 frequency harmonic balance solutions, case SS4, 601 cells, compared with 1000-cell time-accurate solution.

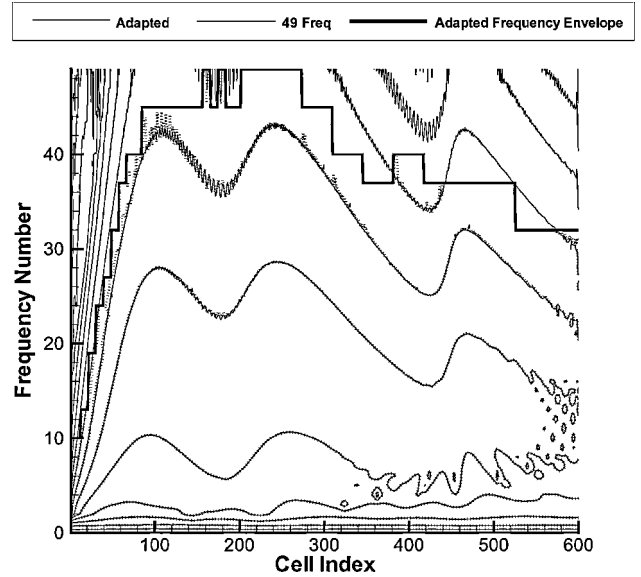


Fig. 11 Isolevels of Fourier coefficient magnitude (ρu) for adapted (augmentation threshold $7.0e-8$) and nonadapted, 49-frequency harmonic balance solutions, case SS4, 601 grid cells.

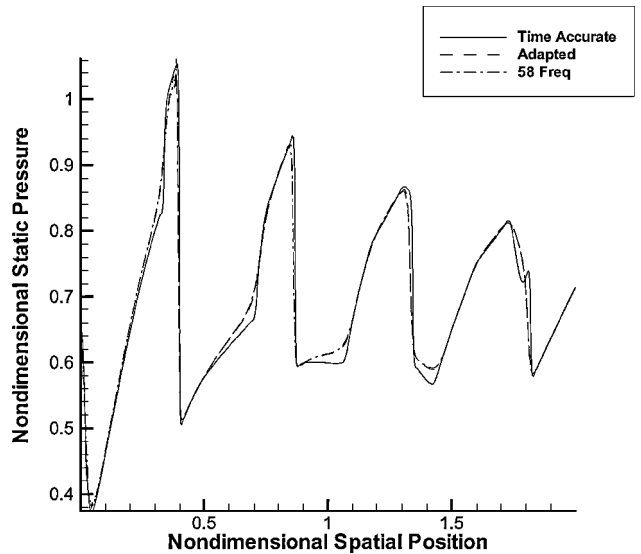


Fig. 12 Computed nondimensional static pressure for adapted (augmentation threshold $7.0e-8$, variable CFL) and nonadapted, 58 frequency harmonic balance solutions, case SU1, 601 cells, compared with 1000-cell time-accurate solution.

higher frequencies correspond to smaller magnitudes. Differences between the adapted and nonadapted harmonic balance solutions that are too small to be easily seen in the reconstructed pressure distributions are easily seen as differences in the coefficient magnitude contours.

The pressure distributions for case SS4 (Fig. 10) are typical of all of the supersonic cases, that is, the reconstructed pressures for the adapted and nonadapted solutions are equivalent and compare favorably with the time-accurate calculation. Both include slight second-order Gibbs effects at the shocks that normally would be removed with increased artificial dissipation. The Fourier coefficient magnitudes of the ρu term (Fig. 11) confirm 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 higher frequencies included in the constant-frequency (nonadapted) approach but not in the adapted approach have no significant impact on accuracy.

For case SU1, the reconstructed pressure distributions for the adapted and nonadapted harmonic balance solutions (Fig. 12) are

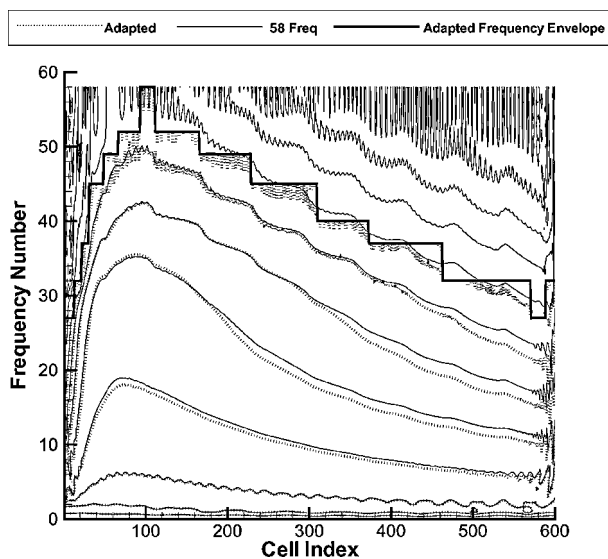


Fig. 13 Isolvels of Fourier coefficient magnitude (ρu component) for adapted (augmentation threshold $7.0e-8$, variable CFL) and non-adapted, 58-frequency harmonic balance solutions, case SU1, 601 grid cells.

comparable. There are slight differences in the two solutions, however, as can be seen by examining the coefficient magnitude contours in Fig. 13. Unlike the supersonic case, there are small but noticeable differences between the adapted and nonadapted harmonic balance Fourier coefficients. These differences result from the use of a variable CFL in the adapted solution. Even with a split-domain harmonic balance formulation, it was sometimes necessary to decrease CFL when a large number of Fourier frequencies were included in the solution. To minimize run time, the CFL was linearly varied between a user-specified maximum and minimum as the sample rate increased. This allowed the use of a larger CFL during the early stages of the solution when frequency content was low and a larger time step was possible and in cells that required fewer frequencies in the final solution.

The operator-splitting technique employed in the split-domain formulation introduces a small discretization error. Because of this error, the steady-state solution computed by the split-domain harmonic balance method has a small dependence on time step size. Application of CFL scaling resulted in larger time steps in some cells and, thus, slightly changed the steady-state solution. The two solutions could be made nearly identical by freezing the CFL of the adapted solution at the same value used for the nonadapted solution, but this would increase run time and result in only minute changes to the reconstructed pressures.

Both adapted and nonadapted harmonic balance solutions for case 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. For this case, the steady-state solutions corresponding to a few time samples have a shock at the exit boundary. The far-field boundary condition applied at that boundary was not designed for such an extreme gradient, and the pressures at the boundary were over-predicted. This is evident in a comparison of the computed pressure at that boundary for the harmonic balance and time-accurate solutions (Fig. 14). The pressure recovery immediately following the period of negative velocity (inflow) was much sharper for the time-accurate solution than for the harmonic balance solution. The result was a delay and attenuation of the left-running expansion wave reflected from the right boundary. 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.

Grid Density

The maximum significant Fourier frequency in any grid cell is highly dependent on grid density. High-frequency coefficients vary

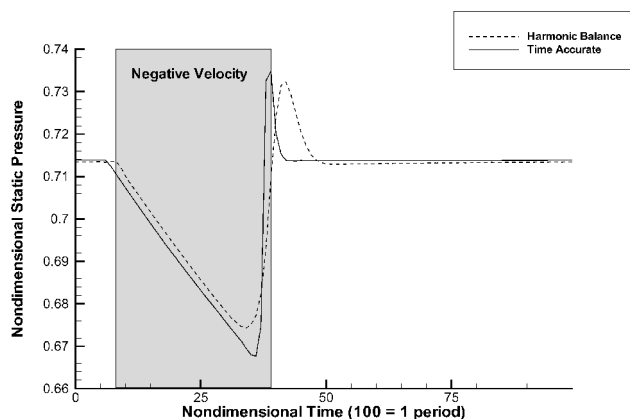


Fig. 14 Comparison of unsteady pressure at tube exit for time-accurate and harmonic balance solutions.

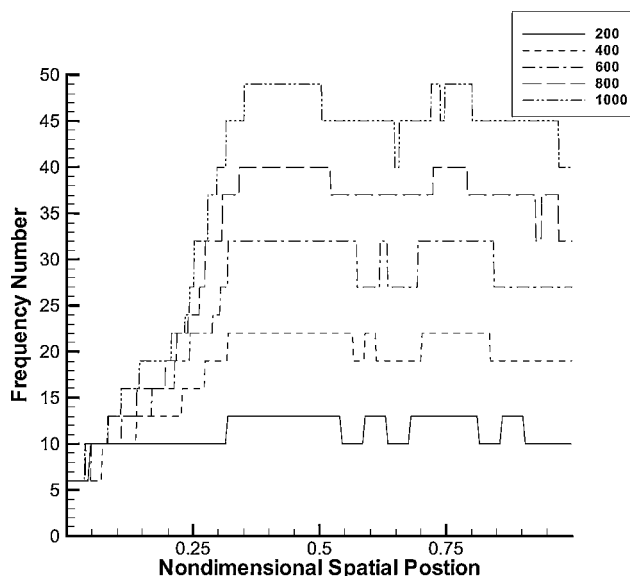


Fig. 15 Variation in adapted frequency distributions with changing grid density for test case SS3, augmentation threshold $5.0e-8$.

rapidly in the spatial dimensions and require a fine grid for accurate resolution. For a grid density insufficient to support a given Fourier frequency, the coefficient 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, automatically matching the included frequency content to that supported by the computational grid. The variation in adapted frequency distribution with grid density is shown in Figs. 15 and 16, which show the final frequency distributions at several grid densities for cases SS3 and SU1, respectively.

Performance

To judge the performance benefit of the adaptive harmonic balance approach, two performance metrics were examined. The first was problem size reduction, as measured by the difference between the maximum and average number of included frequencies in an adapted solution. The second metric was the run time required to converge each solution to steady state. For each of the test cases SS2, SS3, SS4, and SU1, an adapted solution was calculated with an augmentation threshold of $7.0e-8$ and optimal triggers. A non-adapted solution was also calculated based on the highest frequency content in the adapted solution.

The reduction of frequency content and run time for the adapted solution is shown in Fig. 17. In all cases, adaptation resulted in reduced run times, despite relatively small reductions in frequency

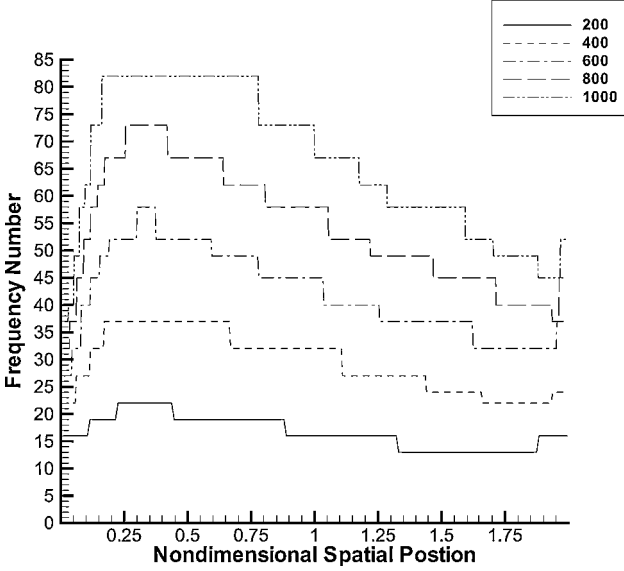


Fig. 16 Variation in adapted frequency distributions with changing grid density for test case SU1, augmentation threshold $5.0e-8$.

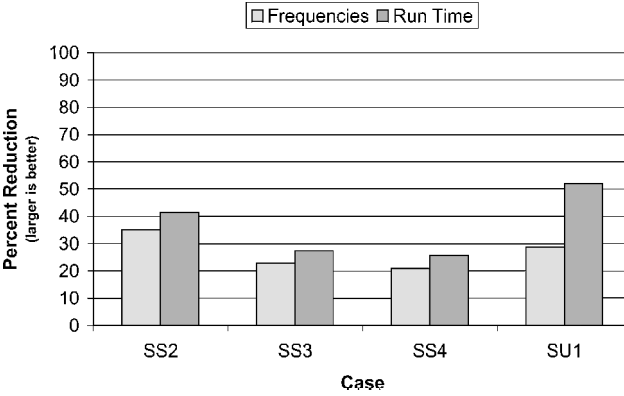


Fig. 17 Reduction of average frequency content and run time for adaptive vs nonadaptive harmonic balance.

content. The subsonic case showed the most improvement, with more than a 50% reduction in run time compared to the nonadapted case. Much of this reduction was due to the variable CFL. The adapted solution was computed with a CFL ranging from 1.7 to 1.35. The nonadapted solution required a fixed CFL of 1.3. The higher CFL, coupled with the efficiency gained in solving the large-scale flow structures with fewer frequencies, resulted in a run time reduction 20% higher than the problem size reduction.

The adapted and nonadapted supersonic cases all converged with the maximum CFL of 1.7, so that the adapted solutions had no time step advantage. The number of iterations required to converge the adapted solutions increased relative to the nonadapted solutions by as much as 15%, but the reduced average-frequency content still allowed a small but significant reduction in run time.

Multigrid

The adaptive harmonic balance technique is compatible with multigrid convergence acceleration. For verification, full approximation scheme (FAS) multigrid was implemented in the solver. Conservative, volume-weighted restriction and linear prolongation operators were applied to the Fourier coefficients in the frequency domain. Residuals required for multigrid calculations were obtained by dividing the difference between a temporarily updated solution and the current solution by the local pseudotime step:

$$\hat{R}_i^n = \frac{(\hat{Q}_i^{n+1} - \hat{Q}_i^n)}{\Delta \tau_i} \quad (12)$$

Table 3 Run-time performance comparison for adapted vs nonadapted harmonic balance solutions, with and without multigrid acceleration, case SS1, 1025 cells

Method	Seven frequencies	Adapted
Nonmultigrid, s	102.3	88.4
Multigrid, s	50.9	45.7

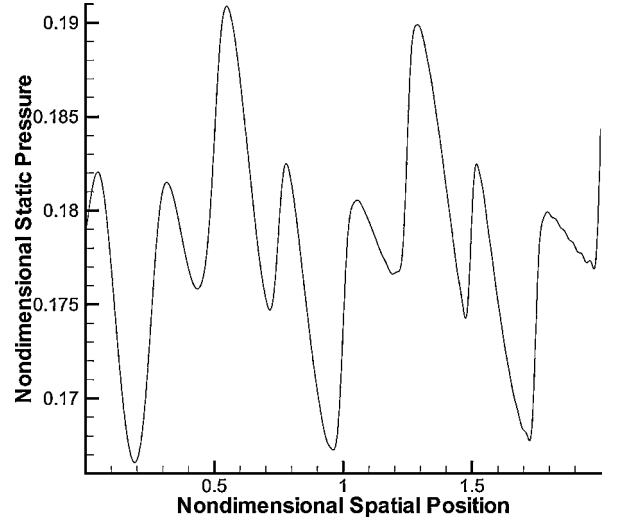


Fig. 18 Computed nondimensional static pressure for adapted (augmentation threshold $7.0e-8$) harmonic balance solutions, case SS1, 1026 cells.

The temporarily updated solution was later replaced by the multigrid-corrected solution.

A low-disturbance-magnitude supersonic case (case SS1) was solved on a 1024-cell grid using a five-level FAS multigrid V cycle (Fig. 18). This case produced a solution that was relatively smooth over most of the grid. The adapted frequency distribution contained just two frequencies at the inlet and gradually increased in one-frequency increments to seven frequencies at the exit.

To simplify the FAS multigrid implementation, the adaptive harmonic balance algorithm was modified slightly. Under the modified algorithm, the primary adaption trigger was based on cycle count rather than residual drop. For this case, adapting every three multigrid cycles was effective.

Run times for both adapted and nonadapted solutions, with and without multigrid, are given in Table 3. In both the adapted and nonadapted cases, run times with multigrid acceleration are shorter. The nonadapted solution benefits slightly more than the adapted solution. Whereas these speedups are significant, much better acceleration should be possible. For example, the linear restriction and prolongation operators used in this implementation are not well suited to the oscillatory nature of the Fourier coefficients obtained from the harmonic balance solution. The authors believe that development of improved restriction and prolongation operators designed specifically for the harmonic balance equations is an area of research that could greatly improve on the acceleration observed here.

Conclusions

A new adaptive harmonic balance method was presented and applied to a variety of subsonic and supersonic one-dimensional flowfields containing strong moving shocks. The adaptive method is a frequency augmentation approach that automatically increases, on a cell-by-cell basis, the frequency content of the solution. Increases are based on E_N , the fraction of spectral energy contained in the highest computed Fourier frequency. For thresholds in the range of $5e-7$ – $5e-8$, E_N reliably identified grid cells requiring additional frequencies that could be adequately resolved by the computational grid.

Adaptation scheduling was found to have a significant impact on the run-time performance of the adaptive solver. 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. Proper selection of augmentation increments was also found to be critical. When increments were selected that minimized the time required for FFT calculations, run times were reduced by as much as a factor of three.

For all test cases, the adaptive method with best parameters produced solutions equivalent to a nonadapted harmonic balance solution, but with as much as 2.1 times speedup. This speedup was obtained despite the fact that the test cases contained strong discontinuities throughout most of the solution domain and, thus, had a relatively high average adapted-frequency content.

The authors acknowledge that the adaptive harmonic balance solution approach is not the ideal approach for all problems and that, for the test cases presented in the paper, a conventional time-accurate approach would be more efficient. The test cases addressed were chosen to stress the algorithm, not to be good benchmark cases. Two factors make them poor benchmark cases, the high overall frequency content throughout the grid and the high grid density that supports the high-frequency content. The adaptive approach does not require that all possible frequencies be included: Frequency content is naturally capped by grid density and can be additionally controlled through the user-defined adaptation threshold E_{thresh} .

Maximum benefit from the adaptive approach is achieved when the average-frequency content is low and regions of high-frequency content are limited to a portion of the domain, either through the physics of the flow or through grid design. Under these conditions, the adaptive approach provides higher fidelity at the same cost as current harmonic balance approaches, or if desired, it provides the same fidelity at a reduced cost.

Acknowledgments

This work was partially funded by Dayton Area Graduate Studies Institute Grant PR-AFIT-99-07, under the auspices of Frank Moore. The authors also thank Paul Orkwis of the University of Cincinnati and Mitch Wolff of Wright State University for their assistance in this work. The views expressed in this article 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 U.S. Government.

References

- ¹Hall, K. C., Thomas, J. P., and Clark, W. S., "Computation of Unsteady Nonlinear Flows in Cascades Using a Harmonic Balance Technique," *AIAA Journal*, Vol. 40, No. 5, 2002, pp. 879–886.
- ²McMullen, M., Jameson, A., and Alonso, J. J., "Acceleration of Convergence to a Periodic Steady State in Turbomachinery Flows," AIAA Paper 2001-0152, Jan. 2001.
- ³Thomas, J. P., Dowell, E. H., and Hall, K. C., "Nonlinear Inviscid Aerodynamic Effects on Transonic Divergence, Flutter, and Limit Cycle Oscillations," *AIAA Journal*, Vol. 40, No. 4, 2002, pp. 638–646.
- ⁴McMullen, M., Jameson, A., and Alonso, J. J., "Application of a Non-Linear Frequency Domain Solver to the Euler and Navier–Stokes Equations," AIAA Paper 2002-0120, Jan. 2002.
- ⁵Maple, R. C., King, P. I., and Orkwis, P. D., "Split-Domain Harmonic Balance Solutions to Burger's Equation for Large Amplitude Disturbances," *AIAA Journal*, Vol. 41, No. 2, 2003, pp. 206–212.
- ⁶Jameson, A., Schmidt, W., and Turkel, E., "Numerical Solutions of the Euler Equations by Finite Volume Methods Using Runge–Kutta Time-Stepping Schemes," AIAA Paper 81-1259, June 1981.
- ⁷Swanson, R., and Turkel, E., "Artificial Dissipation and Central Difference Schemes for the Euler and Navier–Stokes Equations," AIAA Paper 87-1107, June 1987.
- ⁸Strang, G., "On the Construction and Comparison of Difference Schemes," *SIAM Journal on Numerical Analysis*, Vol. 5, No. 3, 1968, pp. 506–517.
- ⁹Enander, R., "Implicit Explicit Residual Smoothing for the Multidimensional Euler and Navier–Stokes Equations," *SIAM Journal on Scientific Computing*, Vol. 18, No. 5, 1997, pp. 1243–1254.
- ¹⁰Frigo, M., and Johnson, S. G., "FFTW: An Adaptive Software Architecture for the FFT," *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech, and Signal Processing*, Inst. of Electrical and Electronics Engineers, 1998, pp. 1381–1384.
- ¹¹Janus, J. M., "The Development of a Three-Dimensional Split Flur Vector Euler Solver with Dynamic Grid Applications, M.S. Thesis, Dept. of Aerospace Engineering, Mississippi State Univ., Mississippi State, MS, Aug. 1984.

P. Givi
Associate Editor