Available online at www.sciencedirect.com
ScienceDirect
Journal of Sound and Vibration 324 (2009) 1016–1038
JOURNAL OF SOUND AND VIBRATION
www.elsevier.com/locate/jsvi

High dimensional harmonic balance dealiasing techniques for a Duffing oscillator

A. LaBryer*, P.J. Attar

Department of Aerospace and Mechanical Engineering, University of Oklahoma, Norman, OK 73019, USA

Received 16 July 2008; received in revised form 9 December 2008; accepted 4 March 2009
Handling Editor: S. Bolton
Available online 8 April 2009

Abstract

A dealiasing study of the high dimensional harmonic balance method is conducted using a Duffing oscillator as a prototypical nonlinear dynamical system. Previous study has shown that aliasing can occur when the high dimensional harmonic balance method is used to solve systems that contain nonlinearities. We demonstrate that frequency filtering techniques, such as the Fourier smoothing method, can sufficiently reduce or eliminate the effects of aliasing. The drawback is that Fourier filtering requires redundant coordinate transformations between the time and frequency domains, resulting in unnecessary computational expense. As an alternative, temporal filters are constructed based upon compact finite difference schemes with spectral-like resolution. It is shown that the temporal filtering schemes successfully replicate the performance of the Fourier smoothing method with improved computational economy.

1. Introduction

The harmonic balance (HB) method is a computationally efficient approach for modeling nonlinear dynamical systems that have time-periodic solutions. Duffing's oscillator is an excellent prototype for such a system. A more complex example of a nonlinear dynamical system may involve a strong coupling between a structure and a fluid, otherwise known as a fluid–structure interaction problem. Time-periodic phenomena such as flutter, limit cycle oscillation, and buffet are of great importance in fluid–structure interaction problems and are most often studied. Another problem which has gathered recent attention due to the current emphasis on developing new micro aerial vehicle (MAV) technology is the aeroelastic response of flapping wings. While these problems are ultimately of great interest, the computational fluid and structural dynamics techniques that are commonly used to solve them contain too many dof to accommodate a general study of the HB method. Thus, it is more practical to channel the study of the HB methodology through a more manageable system, such as a Duffing oscillator. This is the approach that we will take in this work.

The classical HB approach involves substituting a temporal Fourier series expansion of the solution variables into the governing equations. The resulting equations are then expanded, and the terms associated

*Corresponding author. Tel.: +1 918 361 8514.
E-mail address: allen.r.labryer@ou.edu (A. LaBryer).

0022-460X/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.jsv.2009.03.005

with each harmonic are balanced in accordance with the uniqueness theorem of a trigonometric series. We denote the number of terms in the Fourier series expansion as $N_T = 2N_H+1$, where N_H is the number of harmonics retained. This method produces N_T equations to determine the N_T harmonic coefficients. Thus, an analytical solution can be obtained in terms of a Fourier series with N_H harmonics and N_T terms. In general, a more accurate representation of the true solution is achieved by retaining more harmonics.

The high dimensional harmonic balance (HDHB) approach is a modification of the classical HB method and was first presented by Hall and Thomas et al. [1,2]. The basis of the HDHB method is that instead of working in the Fourier domain as with the conventional HB approach, the problem is cast into the time domain. Rather than solving for Fourier coefficients, the dependent variables are discretized in time and stored at equally spaced time intervals for one period of oscillation. This novel modification circumvents the need to balance Fourier coefficients as required in the classical HB approach. Also, the HDHB approach allows for simple implementation of the HB methodology into large scale fluid and structural dynamic computational codes. It is worth noting here that other HB methods have been successful in computing periodic solutions for aeroelastic problems [3–6]. A more detailed discussion on the many variants of the HB methodology can be found in Dimitriadis' continuation study of higher-order HB solutions for nonlinear aeroelastic systems [7].

One shortcoming of using the HDHB approach for solving nonlinear systems is that it has a tendency to produce nonphysical solutions in addition to the physically meaningful ones that are sought [8]. This is due to the treatment of the nonlinear terms in the governing equations and can lead to numerical instability. A nonphysical solution can be identified by a lack of convergence in the Fourier series. This effect is known as aliasing.

Liu et al. and Hall [8] formally examined the effects of aliasing in a comparison of the classical HB and HDHB approaches for a Duffing oscillator. Through analytical expansion, they demonstrated that the classical HB method provides high accuracy and is free of nonphysical solutions. However, when many harmonics are included, the HB solution becomes analytically cumbersome. With the numerical HDHB method, highly accurate results can also be obtained for a large number of harmonics. In their work, Liu et al showed that the nonlinear equations produced by the HDHB approach contain all the terms produced by classical HB approach, plus some additional terms which cause the aliasing. Through numerical simulation, they revealed that aliasing in the HDHB method is more prevalent for cases with large motions, especially near regions of hysteresis. As more harmonics are used, the effects of aliasing are reduced with increased computational expense. They also showed that in the HDHB solution of the Duffing equation, approximately twice as many harmonics must be retained in order to achieve the same order of accuracy as the classical HB method.

The aim of this present study is to develop dealiasing techniques for the HDHB approach in order to improve its efficiency and accuracy for nonlinear dynamical systems. Dealiasing may be accomplished by eliminating or reducing the occurrence of nonphysical solutions via application of a filter to the solution. Results will be given here for both frequency and time domain filtering. To align with the efforts of Liu et al. [8], a Duffing oscillator will serve as the prototypical system to study the HDHB method.

A Duffing oscillator [9,10] is a nonlinear dynamical system that may exhibit both periodic and chaotic behavior. Within the regime of periodic solutions, the Duffing oscillator is well suited for HB analysis. The governing second-order differential equation for a forced Duffing oscillator is given by

$$m\ddot{x} + c\dot{x} + kx + \alpha x^3 = F\,\sin(\omega t), \tag{1}$$

where m, c, k, α, F, and ω are the mass, damping, linear stiffness, nonlinear stiffness, harmonic forcing amplitude, and excitation frequency of the system, respectively. Physically, a Duffing oscillator can be idealized as a damped mechanical oscillator having a nonlinear spring. That is, the restoring force of the system does not exactly obey Hooke's law. The linear stiffness k is generally assumed to be positive, while a positive value of α models a hardening spring and a negative value of α models a softening spring. Representing Eq. (1) in a non-dimensional form will enable us to investigate nonlinear phenomena in a more generic manner. We proceed by introducing the natural frequency ω_0, the damping ratio ζ, the non-dimensional forcing amplitude $\tilde{F}$, the non-dimensional frequency $\tilde{\omega}$, the non-dimensional time τ, and the state variable $\tilde{x}$ such that

$$\omega_0 = \sqrt{k/m}, \quad \zeta = c/2m\omega_0, \quad h = \sqrt{k/\alpha}, \quad \tilde{F} = F/kh, \quad \tilde{\omega} = \omega/\omega_0, \quad \tau = \omega_0 t, \quad \tilde{x} = x/h. \tag{2}$$

Inserting the parameters from Eq. (2) into Eq. (1) results in the following non-dimensional equation:

$$\ddot{\tilde{x}} + 2\zeta\dot{\tilde{x}} + \tilde{x} + \tilde{x}^3 = \tilde{F}\,\sin(\tilde{\omega}\tau). \tag{3}$$

For convenience, the $\sim$ notation is omitted for the remainder of this study and τ is replaced by t. All variables remain in non-dimensional form. Alternatively, Eq. (3) may be expressed as a system of two first-order differential equations. By defining $y = \dot{x}$, we have

$$\dot{\mathbf{X}} + \mathbf{R} - \mathbf{F} = \mathbf{0}, \tag{4}$$

where

$$\mathbf{X} = \begin{Bmatrix} x \\ y \end{Bmatrix}, \quad \dot{\mathbf{X}} = \begin{Bmatrix} \dot{x} \\ \dot{y} \end{Bmatrix}, \quad \mathbf{R} = \begin{Bmatrix} -y \\ 2\zeta y + x + x^3 \end{Bmatrix}, \quad \mathbf{F} = \begin{Bmatrix} 0 \\ F\,\sin(\omega t) \end{Bmatrix}, \quad \mathbf{0} = \begin{Bmatrix} 0 \\ 0 \end{Bmatrix}.$$

The non-dimensional system presented in Eq. (4) is the nonlinear dynamical system that will be solved using the HB methodology. Physically, variable x represents the displacement of the system, while variable y represents the velocity. Note that it is not necessary to express the Duffing equation as a system of first-order equations for the HDHB method. In fact, doing so increases the computational burden. However, rewriting the governing equation as a system of first-order equations will enable us to specify initial conditions for both displacement and velocity.

2. HB theory

Functions that are smooth and periodic in time may be expressed exactly by a Fourier series with appropriate Fourier coefficients. We will assume the solution to Eq. (4) to be smooth and periodic with period $T = 2\pi/\omega$, where ω is the fundamental frequency. Consequently, the solutions for displacement x and velocity y can be expanded in a Fourier series. For example,

$$x(t) = \hat{x}_0 + \sum_{k=1}^{N_H}[\hat{x}_{2k-1}\,\cos(k\omega t) + \hat{x}_{2k}\,\sin(k\omega t)] \quad \text{and} \quad y(t) = \hat{y}_0 + \sum_{k=1}^{N_H}[\hat{y}_{2k-1}\,\sin(k\omega t) + \hat{y}_{2k}\,\cos(k\omega t)]. \tag{5}$$

The only task that remains is to evaluate the Fourier coefficients for the solution variables x and y. This can be done analytically through the classical HB approach or numerically through the HDHB approach. Note that if the number of harmonics used in the expansion N_H is infinity, then Eq. (5) represents x and y exactly. Obviously, using an infinite number of harmonics is not computationally feasible. In practice, a finite number of harmonics are used. We will define the *harmonic approximation error* as the difference between the true solution where N_H is infinity and the approximate solution where N_H is a finite integer.

It is worth noting here that the HB methodology provides a fundamental advantage compared to time-marching methods, such as Runge–Kutta schemes, for time-periodic problems. Namely, time-marching methods include both the transient and the steady state response in the solution. In many cases, only the steady state response is desired. To its advantage for such problems, the HB method only includes the steady state response.

2.1. Classical HB method

The classical HB method is an analytical approach to determine the Fourier coefficients in Eq. (5). The first step in the classical HB method is the substitution of the Fourier expansions into the governing equations. The resulting equations are then expanded and the terms associated with each harmonic are balanced in accordance with the uniqueness theorem of a trigonometric series. This method produces $N_T = 2N_H+1$ equations to determine the N_T harmonic coefficients. Note that when nonlinear terms are present, evaluation of the Fourier coefficients becomes nontrivial. Such is the case with the cubic term in the Duffing equation.

Liu et al. [8] showed that the Fourier expansion of the cubic nonlinearity in Eq. (4) results in wavenumbers that are greater than the number of harmonics used in the Fourier expansion. One way to circumvent this problem is to expand the nonlinear terms using only the first N_H terms of the Fourier series, which will be done

here for the classical HB method. It should be noted that other techniques can be used to expand the nonlinear terms [7]. For the classical HB approach, the cubic term may be written as

$$(x(t))^3 \approx \hat{c}_0 + \sum_{k=1}^{N_H}[\hat{c}_{2k-1}\cos(k\omega t) + \hat{c}_{2k}\sin(k\omega t)], \tag{6}$$

where the Fourier coefficients of the cubic term are given by

$$\hat{c}_0 = \frac{1}{2\pi}\int_0^{2\pi}\left(\hat{x}_0 + \sum_{k=1}^{N_H}[\hat{x}_{2k-1}\cos(kt) + \hat{x}_{2k}\sin(kt)]\right)^3 \mathrm{d}t,$$

$$\hat{c}_{2k-1} = \frac{1}{\pi}\int_0^{2\pi}\left(\hat{x}_0 + \sum_{k=1}^{N_H}[\hat{x}_{2k-1}\cos(kt) + \hat{x}_{2k}\sin(kt)]\right)^3 \cos(kt)\,\mathrm{d}t,$$

$$\hat{c}_{2k} = \frac{1}{\pi}\int_0^{2\pi}\left(\hat{x}_0 + \sum_{k=1}^{N_H}[\hat{x}_{2k-1}\cos(kt) + \hat{x}_{2k}\sin(kt)]\right)^3 \sin(kt)\,\mathrm{d}t.$$

Now we can substitute Eqs. (5) and (6) into Eq. (4). An algebraic expansion and collection of terms associated with each harmonic yields a system of $2N_T$ equations to determine the $2N_T$ Fourier coefficients. The resulting classical HB system can be expressed in matrix form. Namely,

$$\omega\mathbf{A}\hat{\mathbf{Q}} + \hat{\mathbf{R}} - \hat{\mathbf{F}} = \mathbf{0}, \tag{7}$$

where

$$\hat{\mathbf{Q}} = \begin{bmatrix}\hat{\mathbf{q}}_0\\ \hat{\mathbf{q}}_1\\ \hat{\mathbf{q}}_2\\ \vdots\\ \hat{\mathbf{q}}_{2N_H}\end{bmatrix}_{N_T\times 2}, \quad \hat{\mathbf{R}} = \begin{bmatrix}\hat{\mathbf{r}}_0\\ \hat{\mathbf{r}}_1\\ \hat{\mathbf{r}}_2\\ \vdots\\ \hat{\mathbf{r}}_{2N_H}\end{bmatrix}_{N_T\times 2}, \quad \hat{\mathbf{F}} = \begin{bmatrix}0\\ 0\\ \hat{\mathbf{f}}_3\\ \vdots\\ 0\end{bmatrix}_{N_T\times 2}$$

with

$$\hat{\mathbf{q}}_k = \lfloor \hat{x}_k \;\; \hat{y}_k \rfloor_{1\times 2}, \quad \hat{\mathbf{r}}_k = \lfloor -\hat{y}_k \;\; 2\zeta\hat{y}_k + \hat{x}_k + \hat{c}_k \rfloor_{1\times 2}, \quad \hat{\mathbf{f}}_3 = \lfloor 0 \;\; \Phi \rfloor_{1\times 2}$$

and

$$\mathbf{A} = \begin{bmatrix}0 & & & & \\ & \mathbf{J_1} & & & \\ & & \mathbf{J_2} & & \\ & & & \ddots & \\ & & & & \mathbf{J_{N_H}}\end{bmatrix}_{N_T\times N_T}, \quad \mathbf{J}_k = \begin{bmatrix}0 & k\\ -k & 0\end{bmatrix} \quad \{k = 1, 2, \ldots, N_H\}.$$

The HB solution array $\hat{\mathbf{Q}}$ contains the Fourier coefficients for each solution variable. We denote the solution to the HB system (7) as HB1 when $N_H = 1$, HB2 when $N_H = 2$, and so on. The system can be solved exactly for the linear case when $\alpha = 0$. For the nonlinear case, solutions for $\hat{x}_k$ and $\hat{y}_k$ cannot be determined analytically. Instead, solutions for the harmonic amplitudes A_k can be found, along with the peak amplitude A, where

$$A_0 = \sqrt{\hat{x}_0^2}, \quad A_k = \sqrt{\hat{x}_{2k-1}^2 + \hat{x}_{2k}^2}, \quad A = \sum_{k=0}^{N_H} A_k. \tag{8}$$

Using Eq. (8), the classical HB approach generates N_H analytical expressions that can be used to generate solutions for the harmonic amplitudes A_n. Only the real valued solutions are physically meaningful.

This method yields highly accurate solutions that are free of aliasing. However, we note that implementing the classical HB method for dynamical systems generated via computational fluid dynamics or structural dynamics can become unwieldy, especially when many harmonics are used. Moreover, the classical HB approach can only handle systems with nonlinearities that are simple polynomial functions of the solution variables. When more complex nonlinearities are present, the HB system may become impossible to implement [8]. The HDHB approach overcomes these difficulties.

2.2. High dimensional HB method

The basis of the HDHB approach is that the Fourier coefficients can be related to the time domain variables through a discrete Fourier transform operator **E**. The time domain variables are represented at uniformly spaced time intervals for one period of oscillation. Namely,

$$\hat{\mathbf{Q}} = \mathbf{E}\tilde{\mathbf{Q}}, \quad \hat{\mathbf{R}} = \mathbf{E}\tilde{\mathbf{R}}, \quad \hat{\mathbf{F}} = \mathbf{E}\tilde{\mathbf{F}}, \tag{9}$$

where

$$\tilde{\mathbf{Q}} = \begin{bmatrix} \tilde{\mathbf{q}}(t_0) \\ \tilde{\mathbf{q}}(t_1) \\ \tilde{\mathbf{q}}(t_2) \\ \vdots \\ \tilde{\mathbf{q}}(t_{2N_H}) \end{bmatrix}_{N_T \times 2}, \quad \tilde{\mathbf{R}} = \begin{bmatrix} \tilde{\mathbf{r}}(t_0) \\ \tilde{\mathbf{r}}(t_1) \\ \tilde{\mathbf{r}}(t_2) \\ \vdots \\ \tilde{\mathbf{r}}(t_{2N_H}) \end{bmatrix}_{N_T \times 2}, \quad \tilde{\mathbf{F}} = \begin{bmatrix} \tilde{\mathbf{f}}(t_0) \\ \tilde{\mathbf{f}}(t_1) \\ \tilde{\mathbf{f}}(t_2) \\ \vdots \\ \tilde{\mathbf{f}}(t_{2N_H}) \end{bmatrix}_{N_T \times 2}$$

with

$$\tilde{\mathbf{q}}(t_i) = \lfloor x(t_i) \;\; y(t_i) \rfloor_{1\times 2}, \quad \tilde{\mathbf{r}}(t_i) = \lfloor -y(t_i) \;\; 2\zeta y(t_i) + x(t_i) + x(t_i)^3 \rfloor_{1\times 2},$$
$$\tilde{\mathbf{f}}(t_i) = \lfloor 0 \;\; F \sin(t_i) \rfloor_{1\times 2}, \quad t_i = \frac{2\pi i}{N_T} \;\; \{i = 0, 1, 2, \ldots, 2N_H\}.$$

The Fourier transform operator is given by

$$\mathbf{E} = \frac{2}{N_T} \begin{bmatrix} 1/2 & 1/2 & \cdots & 1/2 \\ \cos t_0 & \cos t_1 & \cdots & \cos t_{2N_H} \\ \sin t_0 & \sin t_1 & \cdots & \sin t_{2N_H} \\ \cos 2t_0 & \cos 2t_1 & \cdots & \cos 2t_{2N_H} \\ \sin 2t_0 & \sin 2t_1 & \cdots & \sin 2t_{2N_H} \\ \vdots & \vdots & & \vdots \\ \cos N_H t_0 & \cos N_H t_1 & \cdots & \cos N_H t_{2N_H} \\ \sin N_H t_0 & \sin N_H t_1 & \cdots & \sin N_H t_{2N_H} \end{bmatrix}_{N_T \times N_T}.$$

Note that the solution array $\tilde{\mathbf{Q}}$ contains the time domain solution variables for displacement and velocity and is discretized in time at N_T evenly spaced intervals for one period of oscillation. The time domain variables can also be related to the Fourier coefficients through the inverse of the Fourier transform operator **E**. That is,

$$\tilde{\mathbf{Q}} = \mathbf{E}^{-1}\hat{\mathbf{Q}}, \quad \tilde{\mathbf{R}} = \mathbf{E}^{-1}\hat{\mathbf{R}}, \quad \tilde{\mathbf{F}} = \mathbf{E}^{-1}\hat{\mathbf{F}}, \tag{10}$$

where

$$\mathbf{E}^{-1} = \begin{bmatrix} 1 & \cos t_0 & \sin t_0 & \cdots & \cos N_H t_0 & \sin N_H t_0 \\ 1 & \cos t_1 & \sin t_1 & \cdots & \cos N_H t_1 & \sin N_H t_1 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 1 & \cos t_{2N_H} & \sin t_{2N_H} & \cdots & \cos N_H t_{2N_H} & \sin N_H t_{2N_H} \end{bmatrix}_{N_T \times N_T}.$$

It is now possible to work entirely in terms of the time domain variables, which is much easier than directly solving for Fourier coefficients. The resulting HDHB system can be written as

$$\omega \mathbf{D}\tilde{\mathbf{Q}} + \tilde{\mathbf{R}} - \tilde{\mathbf{F}} = \mathbf{0}, \tag{11}$$

where the Fourier derivative operator is given by

$$\mathbf{D} = \mathbf{E}^{-1}\mathbf{A}\mathbf{E}. \tag{12}$$

The solution to the HDHB system (11) can be obtained numerically using an iterative root finding scheme such as the Newton–Raphson method [11]. Alternatively, methods such as pseudo-time-marching [12], used to solve steady state problems in computational fluid dynamics, may be used for high dimensional systems. Consistent with Liu et al. [8], we denote the solution to the HDHB system as HDHB1 when $N_H = 1$, HDHB2 when $N_H = 2$, and so on. The Fourier coefficients can readily be determined from the time domain solutions per Eq. (10). From there, the harmonic amplitudes A_n and the peak amplitude A can be obtained per Eq. (8). Note that we have taken full advantage of the assumption that the solution is periodic in time. Thus, the computational cost associated with a transient response is completely avoided.

3. Evidence of aliasing

One shortcoming of the HDHB method is that when nonlinearities are present in the governing equations, it will produce nonphysical solutions in addition to the physically meaningful ones that are sought. This is due to the use of the discrete Fourier transform operator $\mathbf{E}$, which produces additional terms in the equations that are not present with the classical HB method. The additional terms arise from multiplication with higher-order harmonics that are left unresolved with the expansion in Eq. (6). As a result, the HDHB method produces identical solutions compared to the classical HB method, plus some additional solutions that are nonphysical.

The origin of the aliasing terms in the HDHB method can be seen more clearly by examining the complete Fourier expansion of the cubic term:

$$(x(t))^3 = \hat{c}_0 + \sum_{k=1}^{N_H}[\hat{c}_{2k-1}\cos(k\omega t) + \hat{c}_{2k}\sin(k\omega t)] + \sum_{k=N_H+1}^{2N_H}[\hat{c}_{2k-1}\cos(k\omega t) + \hat{c}_{2k}\sin(k\omega t)] \tag{13}$$

In the classical HB approach, the expansion of the cubic term is obtained by neglecting all wavenumbers greater than N_H (Eq. (6)). In the HDHB method, the entire Fourier expansion in Eq. (13) is included. The treatment of the cubic function in the time domain leads to the generation of wavenumbers greater than N_H in the Fourier domain. The Fourier transform operator $\mathbf{E}$ does not account for these higher-order terms, and as a result, spurious solutions emerge. This principle can be extended to other systems with nonlinear terms that involve multiplication of the solution variables. In general, the product of two or more discrete functions in the physical domain corresponds to a circular convolution in Fourier space. A nonlinearity that involves multiplication of two or more terms containing wavenumbers k within the domain $[0,N_H]$ will produce additional wavenumbers beyond N_H.

3.1. Evidence of aliasing in amplitude response curves

In this section, the effects of aliasing are examined by comparing amplitude response curves in the frequency domain. For the driven oscillator problem studied here, a single forcing frequency is specified along with the initial conditions for the solution array. Using the Newton–Raphson method, the system will converge to a solution, provided the initial conditions are within a certain radius of convergence. A peak amplitude curve can be generated by incrementally increasing or decreasing the frequency. At each step, the solution array becomes the initial conditions for the next computation. This procedure is known as *frequency marching*. For nonlinear dynamical systems, amplitude response curves are often accompanied with regions of hysteresis. The upper branch that is generated by increasing the forcing frequency is typically of great interest, and is commonly referred to as the *backbone curve*. Physically, the peak amplitudes on the upper branch of the backbone curve represent the maximum displacement of the system as a result of increasing the forcing frequency.

In Fig. 1, amplitude response curves produced by the classical HB and HDHB methods using two harmonics are compared. To align with the work of Liu et al. [8], the damping ratio ζ is set to 0.1 and the non-dimensional forcing amplitude F is set to 1.25. These parameters are held constant for the remainder of this study. Two sets of harmonic amplitude curves are produced by the HDHB system: one by marching the frequency from 0.1 to 2.8 at increments of $\Delta\omega = 0.01$, and another by marching the frequency from 2.8 to 0.1 at increments of $\Delta\omega = 0.01$.

Observe in Fig. 1 that the prescribed damping ratio and forcing amplitude result in a peak amplitude curve with a large region of hysteresis. A properly behaved Duffing oscillator subject to an increasing forcing frequency will produce a backbone curve similar to the one generated by the HB2 system. When the frequency passes the inflection near $\omega = 2.40$, the peak amplitude will jump down and continue on the lower branch. Similarly, an oscillator that begins on the lower branch will jump to the upper branch as the frequency decreases past the inflection near $\omega = 1.75$. Note that the unstable branch generated by the HB solution cannot be predicted with the HDHB system using the frequency marching procedure.

The HB2 solution only contains harmonic approximation error. We will define *aliasing error* as the difference between the HDHB solution and the HB solution for a given value of N_H. From Fig. 1, it is apparent that the HDHB2 solution contains not only harmonic approximation error, but also a large degree of aliasing error. Generally speaking, and as demonstrated in Fig. 1, the effects of aliasing are more pronounced for large motions and near regions of large hysteresis. Note that the HDHB2 solution obtained by

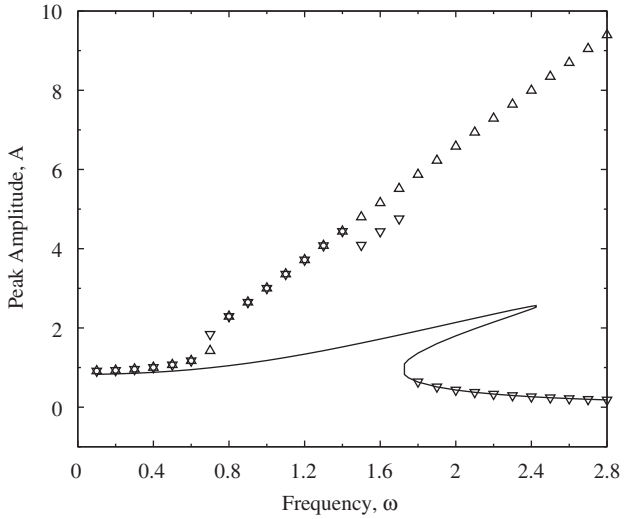

Fig. 1. Peak response amplitudes for the HDHB2 and HB2 systems. HDHB2 by increasing ω (upward triangle); HDHB2 by decreasing ω (downward triangle); HB2 (solid line).

increasing the frequency fails to drop down to the lower branch. Instead, the amplitudes continue to increase beyond the point of hysteresis.

In addition to frequency marching, other methods for initial condition generation are possible. Here, the nature of the solutions generated by the HDHB system is studied from a statistical perspective by generating initial conditions through Monte Carlo simulation. That is, initial conditions are randomly specified for a large number of computations in order to determine the probability of converging to any particular solution. The first harmonic amplitudes A_1 of displacement and velocity are randomly specified within the range of ± 5 units. The forcing frequency is set to $\omega = 2.0$ to produce the three branches of the hysteresis curve seen in Fig. 1. In Fig. 2, Monte Carlo histograms are presented for the solution convergence of the HDHB system using two and four harmonics. For each simulation, 10^4 sample initial conditions are specified.

For the HDHB2 and HDHB4 systems, the probability of converging to any particular solution is highly sensitive to initial conditions. In Fig. 2, the HDHB2 simulation yields 23 unique solutions while the HDHB4 simulation yields 130 unique solutions. Only the three highlighted solutions are physically meaningful. The rest of the solutions present are due to aliasing. In the HDHB2 simulation, the probability of converging to the lower, unstable, and upper branches are 0.025, 0.022, and 0.123, respectively. In the HDHB4 simulation, the probability of converging to the lower, unstable, and upper branches are 0.026, 0.020, and 0.467,

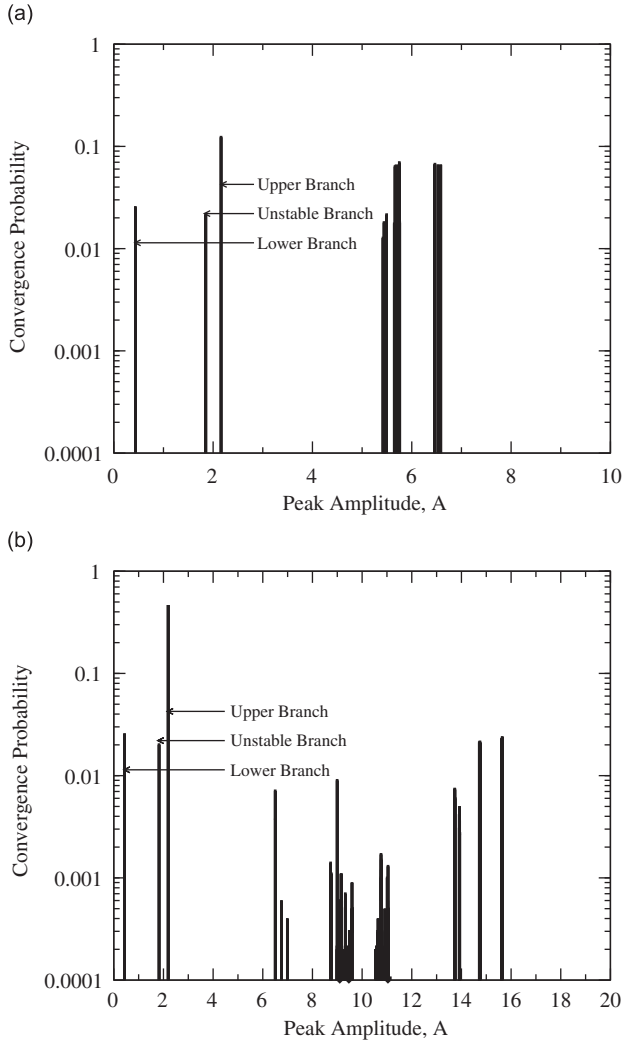

Fig. 2. Monte Carlo histograms for the solutions of the HDHB system at $\omega = 2.0$. (a) Results for the HDHB2 simulation. (b) Results for the HDHB4 simulation.

respectively. In sum, the probability of converging to a physical solution is 0.170 for the HDHB2 simulation and 0.467 for the HDHB4 simulation. In general, the attraction to physical solutions increases as more harmonics are retained. Finally, note that the peak amplitude of 6.58 at $\omega = 2.0$ in Fig. 1 is one of the spurious solutions predicted by the HDHB2 simulation in Fig. 2.

4. Dealiasing by filtering in the frequency domain

In order to improve the numerical properties of the HDHB solution for nonlinear systems, the aliasing error should be reduced. Complete aliasing removal can be accomplished by truncating all wavenumbers that are polluted with nonphysical terms [13,14]. This technique involves zeroing the Fourier coefficients for all wavenumbers greater than some cutoff wavenumber k_c. For example, Orszag [15] was the first to demonstrate that setting $k_c \leqslant (2/3)N_H$ will remove all aliasing terms for systems containing quadratic nonlinearities. Throughout the literature, this relationship is referred to as the *two-thirds rule*. In general, when the nonlinearity of the system is a polynomial of the solution variables, the relationship between the cutoff wavenumber k_c and number of harmonics used for computation N_H is given by

$$\frac{k_c}{N_H} \leqslant \frac{2}{\Phi + 1} \tag{14}$$

where Φ is the degree of the nonlinearity. This general concept of wavenumber truncation can be implemented by filtering the Fourier coefficients in the frequency domain. For future convenience, we introduce the scaled wavenumber $w = \pi k/N_H$ so that the domain of w is $[0,\pi]$ and the scaled cutoff wavenumber is $w_c = \pi k_c/N_H$. The filtered Fourier coefficients can be written as

$$\hat{\mathbf{Q}}' = G(w)\hat{\mathbf{Q}} \tag{15}$$

where $G(w)$ is commonly referred to as a low-pass Fourier filter or transfer function. Recall that the HDHB method is written and solved for in the time domain. In order to apply the Fourier filter, the time domain variables must first be transformed to the frequency domain using the Fourier transform operator $\mathbf{E}$. After the filtering takes place, the variables must be transformed back to the time domain with the inverse Fourier transform operator $\mathbf{E}^{-1}$. This filtering procedure must be performed on the solution array $\mathbf{Q}$ before each nonlinear computation and on the nonlinear forcing array $\mathbf{R}$ after each computation.

4.1. The one-half rule

Orszag's two-thirds rule for quadratic nonlinearities becomes the *one-half rule* for the Duffing oscillator, which contains a cubic term. The top half of all wavenumbers must be truncated. The Fourier filter in Eq. (15) becomes

$$G(w) = \begin{cases} 1 & \text{for } w \leqslant \pi/2, \\ 0 & \text{for } w > \pi/2. \end{cases} \tag{16}$$

Eq. (14) guarantees that this Fourier filter will completely eliminate all aliasing terms in the HDHB system. Note that sharp cutoff functions such as the one in Eq. (16) are susceptible to Gibbs-type phenomena. Also, the computational economy of such a filter is questionable. While wavenumbers w greater than $\pi/2$ are subject to aliasing, they also contain portions of the true solution. For such reasons, it is worth investigating other filters that decay smoothly and retain a portion of higher Fourier modes.

4.2. Fourier smoothing

Hou and Li [16] demonstrated that a Fourier smoothing method can be used to reduce the effects of aliasing with less harmonic approximation error than truncation. Rather than completely zeroing wavenumbers greater than w_c, it is possible to gradually damp out the highest frequency Fourier modes by choosing $G(w)$ to be

$$G(w) = \mathrm{e}^{-\alpha(w/\pi)^m}. \tag{17}$$

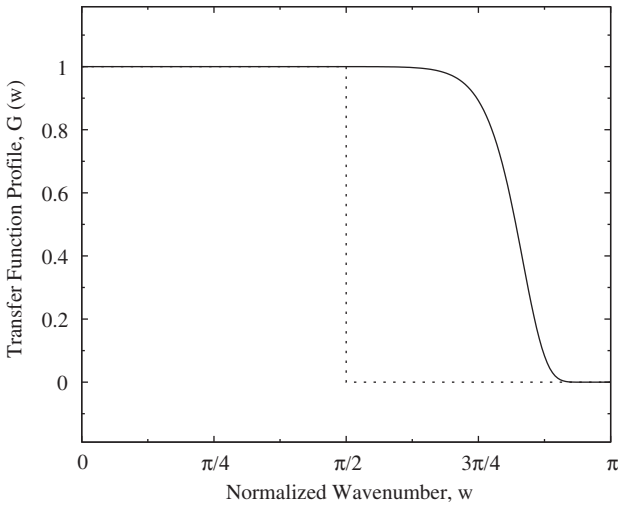

Fig. 3. Transfer function profiles for the frequency domain filtering schemes. One-half rule (dotted line); Fourier smoothing with $\alpha = 36$ and $m = 20$ (solid line).

Here, α is chosen to be 36 and m is chosen to be 20. This filter design allows $G(w)$ to remain very close to unity for $w < \pi$, and then smoothly and rapidly decay close to zero near $w = \pi$. The selection of α results in $G(\pi)$ evaluating to machine precision, i.e. 10^{-16}. The selection of m allows $G(\pi/2)$ to remain within 10^{-5} from unity. In Fig. 3, the profile of the Fourier smoothing filter is compared to that of the one-half rule.

From Fig. 3 we can see that the Fourier smoothing filter retains more harmonics than the one-half rule. Thus, a decrease in harmonic approximation error may be achieved with the possible inclusion of some aliasing error.

4.3. Frequency domain filtering results

In this section, the results for filtering the HDHB solution in the frequency domain are presented. The solution for the HDHB2 system is recomputed, but this time, the system is filtered using the one-half rule and Fourier smoothing. In Fig. 4, the filtered HDHB2 amplitude response curves are presented. The HB2 amplitude response curve is included as a baseline.

As demonstrated in Fig. 4, the filtered HDHB2 system converges only to physically legitimate peak amplitudes, which is a dramatic improvement compared to the unfiltered peak amplitude curve. Moreover, when the one-half rule is applied, the HDHB2 system yields identical results compared to the HB1 system. In general, the HB system using N_H harmonics produces the same results as the HDHB system using $2N_H$ harmonics when the one-half rule is applied. These results are consistent with the findings of Liu et al. [8].

The numerical accuracy of the HDHB system will now be examined in a more general manner. In order to compare the accuracy of the different filtering methods, a true solution must be established. Recall that with the HB methodology, a true solution can only be obtained by retaining an infinite number of harmonics. Here, the "true" solution is represented by the HDHB200 system filtered by the one-half rule, which is identical to the HB100 solution. Application of the one-half rule eliminates aliasing error while the use of two hundred harmonics renders the harmonic approximation error insignificant. The true peak amplitude at $\omega = 2.0$ is approximately $A = 2.1817764$.

We will compare the peak response amplitudes generated by frequency marching up the backbone curve to $\omega = 2.0$ as a function of the number of harmonics N_H. In Fig. 5, the corresponding peak amplitudes are compared for the unfiltered response, the one-half rule and Fourier smoothing. In addition, the percent errors of the peak displacement amplitudes compared to the true solution are presented.

As shown in Fig. 5, frequency marching of the unfiltered HDHB system results in the solution following a nonphysical branch of the backbone curve when three or less harmonics are retained. This divergent behavior

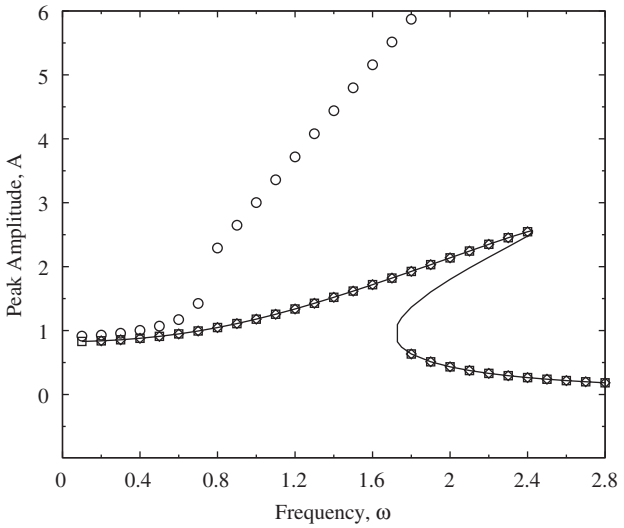

Fig. 4. Peak response amplitudes for the frequency filtered HDHB2 system. HDHB2 filtered by the one-half rule (diamond); HDHB2 filtered by Fourier smoothing (square); unfiltered HDHB2 (circle); HB2 (solid line).

results in large errors and numerical instability. Conversely, when four or more harmonics are retained, frequency marching of the unfiltered system results in physically meaningful amplitudes. The HDHB systems resulting from the one-half rule and Fourier smoothing filtering converge to the backbone for all harmonics.

Several important trends for frequency marching of the Duffing equation can be drawn from Fig. 5. When only a few harmonics are used, the unfiltered HDHB system yields pathological results. When more harmonics are retained, such as four in this example, the unfiltered system provides the best approximation to the true solution. Therefore, the unfiltered system results in the lowest harmonic approximation error when a sufficient number of harmonics are retained. The one-half rule results in the largest harmonic approximation error. With the one-half rule, approximately twice the number of harmonics must be retained in order to gain the same amount of accuracy provided by the unfiltered system. Fourier smoothing sufficiently reduces aliasing so that only physical amplitudes are predicted for frequency marching. For a given value of N_H, Fourier smoothing provides a better approximation to the true solution as compared to the one-half rule.

Using the same procedure described in Section 3.1, we now investigate the statistical behavior of solutions generated by the frequency filtered HDHB system by Monte Carlo simulation for 10^4 initial conditions at $\omega = 2.0$. In Fig. 6, histograms are presented for the HDHB4 system filtered by the one-half rule and by Fourier smoothing.

Observe in Fig. 6 that the HDHB4 system filtered by the one-half rule admits only three physical solutions. The HDHB4 system filtered by Fourier smoothing admits eight unique solutions, five of which are of the spurious variety. While the elimination of nonphysical solutions is ideal, the five yielded by Fourier smoothing are a considerable improvement over the 127 nonphysical solutions yielded by the unfiltered HDHB4 system in Fig. 2. For the simulations which used the one-half rule filtering, the probability of converging to the lower, unstable, and upper branches are 0.026, 0.026, and 0.948, respectively. In the Fourier smoothing simulations, the probability of converging to the lower, unstable, and upper branches are 0.025, 0.025, and 0.798, respectively. In sum, the probability of converging to a physical solution is 1.000 for the one-half rule simulation and 0.848 for the Fourier smoothing simulation. These statistics are much improved when compared to the 0.467 probability of converging to a physical solution for the unfiltered HDHB4 simulation.

To obtain a more comprehensive understanding of the numerical and statistical behavior of the filtered and unfiltered HDHB solutions, Monte Carlo simulations are reproduced for various values of N_H. Each simulation results in a certain number of unique solutions at $\omega = 2.0$, along with the corresponding probability of each solution occurring. The results are presented in Fig. 7 where up to 10 harmonics are retained.

From Fig. 7 it can be seen that the number of spurious solutions produced by the unfiltered system dramatically increases with the number of harmonics. However, the probability of converging to physical

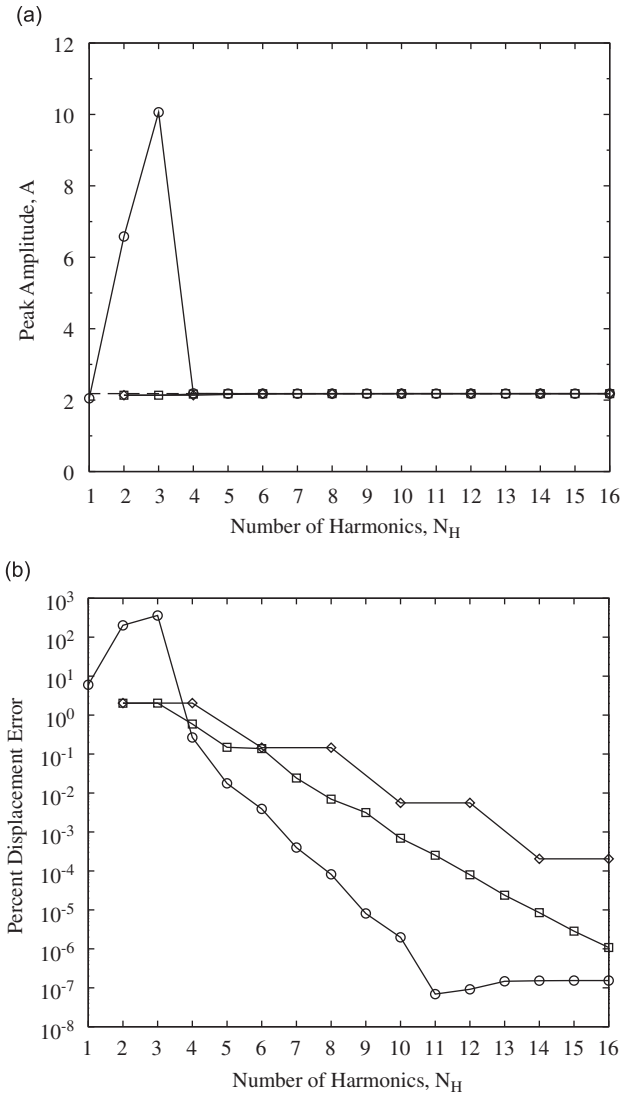

Fig. 5. Numerical accuracy comparison of the filtered and unfiltered HDHB systems. (a) Peak response amplitudes generated by marching up the backbone curve to $\omega = 2.0$. (b) Percent error of the peak amplitudes compared to the true response. True response (dashed line); one-half rule (diamond); Fourier smoothing (square); unfiltered (circle).

amplitudes improves as more harmonics are used. The one-half rule consistently admits three physical amplitudes for all cases. Interestingly, the Fourier smoothing filter eliminates all spurious solutions for certain harmonics, i.e. for two, three, five, and nine harmonics. Even when Fourier smoothing does not completely eliminate spurious solutions, the occurrence is always reduced compared to the unfiltered system, usually by at least a full order of magnitude. The probability of the Fourier smoothing method admitting physical amplitudes is always less than or equal to the one-half rule, but greater than the unfiltered system.

4.4. Higher-order nonlinearities

Higher-order nonlinearities can be modeled by including additional terms in the polynomial expansion of the restoring force. For example, a fifth-order nonlinearity can be modeled by adding an x^5 term to the **R** vector in Eq. (4). Physically, this additional term increases the hardening behavior of the spring.

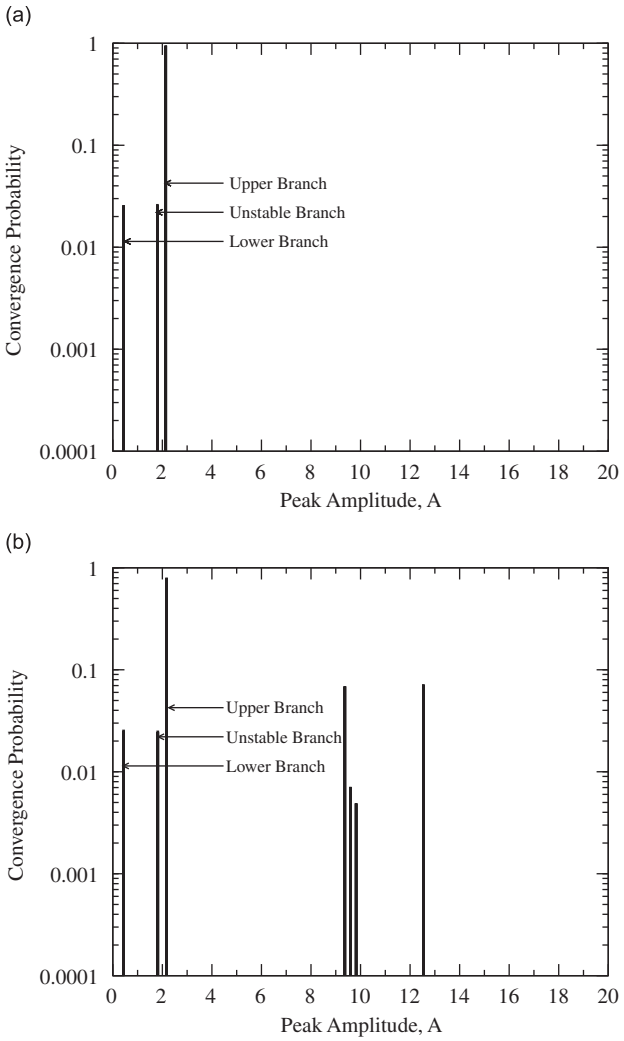

Fig. 6. Monte Carlo histograms for the solutions of the frequency filtered HDHB4 system at $\omega = 2.0$. (a) Results for the one-half rule. (b) Results for Fourier smoothing.

The HDHB formulation for a fifth-order nonlinearity is aliased to a greater extent than the HDHB formulation for the Duffing oscillator. Qualitatively, the effects of aliasing while backbone marching are similar. The unfiltered system converges to nonphysical solutions for lower harmonics. When a certain number of harmonics are retained, the system behavior changes such that frequency marching yields physically correct solutions. For example, in Fig. 5, this occurs when N_H goes from 3 to 4. When a fifth-order nonlinearity is included, this transition occurs when N_H goes from 5 to 6.

Suitable extensions of the dealiasing techniques developed in this section may be employed for the system with a fifth-order nonlinearity. Orszag's two-thirds rule becomes the one-third rule in accordance with Eq. (14). Using the same design criteria as in Section 4.2, the coefficients for the Fourier smoothing method become $\alpha = 36$ and $m = 12$. The one-third rule and Fourier smoothing method successfully eliminate the effects of aliasing for backbone marching. The one-third rule completely eliminates aliasing terms while the Fourier smoothing method provides a better approximation to the true solution.

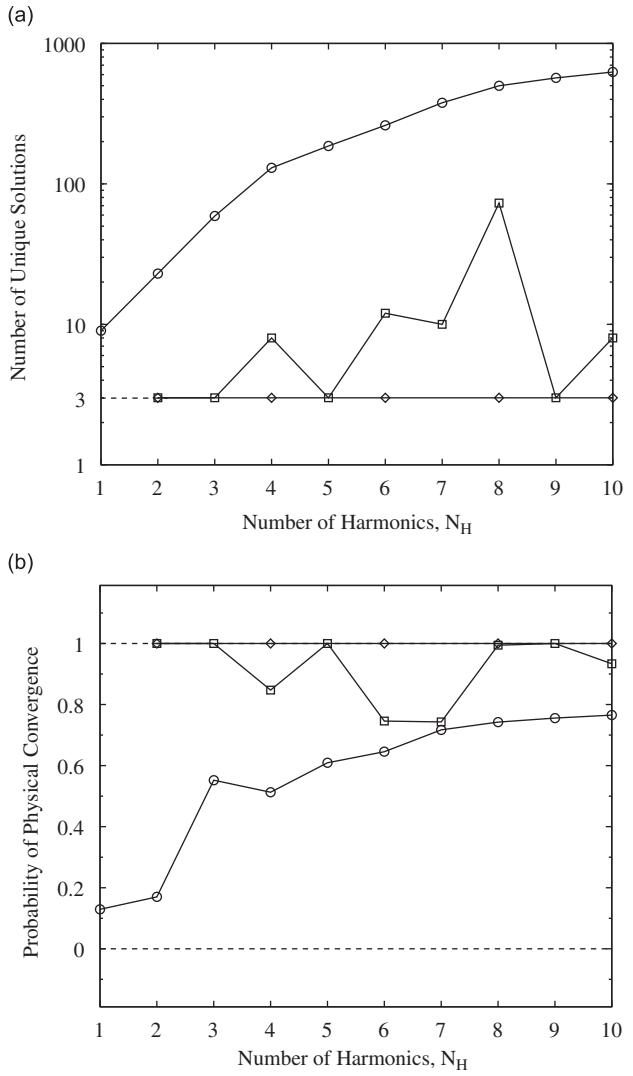

Fig. 7. Solution characteristics for the filtered and unfiltered HDHB systems at $\omega = 2.0$. (a) Number of unique solutions. (b) Probability of converging to a physically meaningful solution. One-half rule (diamond); Fourier smoothing (square); unfiltered (circle).

4.5. Frequency domain filtering summary

The one-half rule completely eliminates the occurrence of spurious solutions by truncating all wavenumbers that contain aliasing terms. This technique is effective yet wasteful, as the truncated wavenumbers also contribute to the accuracy of the physically meaningful solutions. The Fourier smoothing method sufficiently reduces the occurrence of aliasing terms while retaining a greater portion of wavenumbers. When initial conditions are generated by backbone marching, the Fourier smoothing method consistently converges to physical solutions and provides improved harmonic approximation error compared to the one-half rule.

From an implementation and efficiency standpoint, one of the benefits of the HDHB method is that the equations are written and solved entirely in terms of time domain variables. In order to filter using the one-half or Fourier smoothing frequency domain filters, the time domain variables must first be transformed to the frequency domain where the filtering takes place and then transformed back to the time domain. For systems with many dof, this can become computationally expensive. Instead, a filtering procedure which operates entirely upon the time domain variables would be more efficient. Such a procedure will be now outlined.

5. Dealiasing by filtering in the time domain

In this section we design temporal filters that will attempt to "mimic" the results of the Fourier smoothing filter using a Padé type compact finite difference scheme. This circumvents the need to transform variables between the time and frequency domains. The underlying goal is to optimize filtering performance while economizing computational expense.

5.1. Temporal filter construction

Lele demonstrated that compact finite difference schemes can be constructed for low-pass filtering applications [17]. The fundamental idea is that by smoothing a discretized function in time, high wavenumbers in the frequency domain can be effectively damped out. We proceed by considering the uniformly discretized HDHB solution vector components from Eq. (9), which we will denote here as $\mathbf{q}_i$. The general form for the compact finite difference approximation on a seven-point stencil is

$$\beta\mathbf{q}'_{i-2}+\alpha\mathbf{q}'_{i-1}+\mathbf{q}'_i+\alpha\mathbf{q}'_{i+1}+\beta\mathbf{q}'_{i+2}=a\mathbf{q}_i+\frac{b}{2}(\mathbf{q}_{i+1}+\mathbf{q}_{i-1})+\frac{c}{2}(\mathbf{q}_{i+2}+\mathbf{q}_{i-2})+\frac{d}{2}(\mathbf{q}_{i+3}+\mathbf{q}_{i-3}), \tag{18}$$

where $\mathbf{q}'_i$ denotes the smoothed, or filtered, value of $\mathbf{q}_i$. The compact finite difference scheme presented in Eq. (18) is a generalization of the Padé scheme. Here, we denote the stencil size (number of temporal grid points in the differencing scheme) as N_S. The stencil size may be altered to achieve various levels of formal accuracy, with the constraint $N_S \leqslant N_T$. Since the HDHB solution vectors are periodic in time, periodic boundary conditions may be assumed. Note that when α and β are zero, Eq. (18) provides an explicit differencing scheme. That is, the smoothed values $\mathbf{q}'_i$ can be determined directly from the unfiltered values $\mathbf{q}_i$. When α or β are nonzero, Eq. (18) provides an implicit differencing scheme. If α is nonzero and β is zero, the result is a cyclic tridiagonal linear algebraic system. If α and β are nonzero, the result is a cyclic pentadiagonal linear algebraic system. Highly efficient algorithms exist for solving cyclic tridiagonal and pentadiagonal implicit schemes [18,19] which are variants of the Thomas algorithm [11].

Next, we consider what temporal smoothing does to the simple complex sinusoid $q(t)=e^{j\omega t}$. Substitution of the sinusoid into Eq. (18), along with Euler's identity, results in a transfer function $G(\omega)$ such that $\mathbf{q}'_i = G(\omega)\mathbf{q}_i$. Realizing that the frequency ω corresponds to the scaled wavenumber $w=\pi k/N_H$, we obtain

$$G(w)=\frac{a+b\cos(w)+c\cos(2w)+d\cos(3w)}{1+2\alpha\cos(w)+2\beta\cos(2w)}. \tag{19}$$

The temporal filter design process can be accomplished by imposing specific conditions on the transfer function $G(w)$. For low-pass filters, we generally require $G(\pi)=0$. This constraint automatically results in $dG(\pi)/dw=0$. In addition, the location of the cutoff can be controlled by requiring $G(w_1)=\upsilon_1$ and $G(w_2)=\upsilon_2$. Note that special care must be taken when posing conditions on the transfer function. Certain values may result in overshoot behavior. Insofar as possible, we desire $0\leqslant G(w)\leqslant 1$ for w in $[0,\pi]$. To best approximate the Fourier smoothing filter in Eq. (17), we have used $G(7\pi/10)=0.97168343$ and $G(8\pi/10)=0.66030611$.

The relations for the coefficients a, b, c, d and α, β are generated by matching the Taylor series coefficients of various orders. The truncation error of the Taylor series approximants effectively result in the damping of higher wavenumbers. The relationship for the sixth-order approximation is

$$a=\tfrac{1}{16}(11+10\alpha-10\beta),\quad b=\tfrac{1}{32}(15+34\alpha+30),\quad c=\tfrac{1}{16}(-3+6\alpha+26\beta),\quad d=\tfrac{1}{32}(1-2\alpha+2\beta). \tag{20}$$

A different family of schemes is generated when posing the additional constraint $d^2G(\pi)/dw^2=0$. When this constraint is imposed in addition to the fourth-order approximation, the resulting relationship is

$$a=\tfrac{1}{4}(2+3\alpha),\quad b=\tfrac{1}{16}(9+16\alpha+10\beta),\quad c=\tfrac{1}{4}(\alpha+4\beta),\quad d=\tfrac{1}{16}(6\beta-1). \tag{21}$$

The goal of our temporal filter design is to attempt to best replicate the Fourier smoothing transfer function given in Eq. (17). Results will be given here for three temporal filtering schemes constructed on a seven-point stencil, including an explicit, implicit tridiagonal and implicit pentadiagonal scheme. Since here, $N_S=7$,

Table 1
Temporal filtering coefficients for the seven point Padé type compact finite difference schemes.

Temporal scheme	α	β	a	b	c	d
Explicit	–	–	0.6875	0.46875	−0.1875	0.03125
Implicit tridiagonal	0.43026998	–	0.95641874	0.92591186	−0.026148757	0.0043581261
Implicit pentadiagonal	0.66099766	0.17096259	0.99574824	1.3303493	0.33621200	0.0016109707

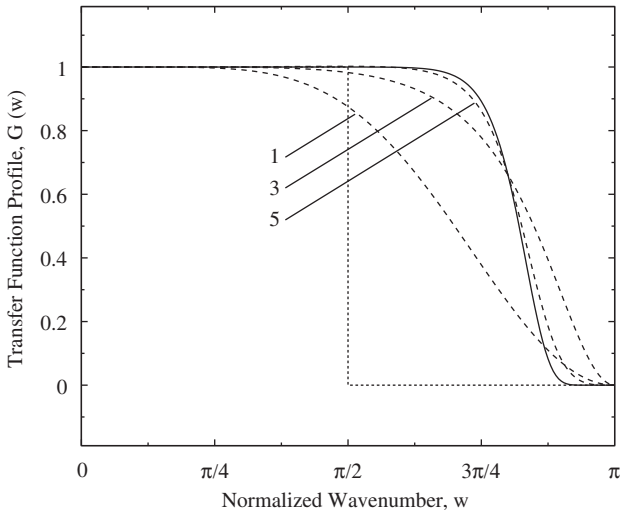

Fig. 8. Transfer function profiles for the temporal and frequency domain filtering schemes. One-half rule (dotted line); Fourier smoothing with $\alpha = 36$ and $m = 20$ (solid line); explicit temporal scheme (dashed line 1); implicit tridiagonal temporal scheme (dashed line 3); implicit pentadiagonal temporal scheme (dashed line 5).

the values of N_T and N_H for the HDHB method are limited to $N_T \geqslant 7$ and $N_H \geqslant 3$. For the explicit scheme, the sixth-order coefficients a, b, c and d are readily determined by setting α and β to zero in Eq. (20). The relationship for the implicit tridiagonal scheme is obtained by solving Eqs. (19) and (20) with $\beta = 0$ along with the constraint $G(8\pi/10) = 0.66030611$. The coefficients for the implicit tridiagonal scheme are determined by solving the fourth-order relationship in Eq. (21) along with the constraints $G(7\pi/10) = 0.95$ and $G(8\pi/10) = 0.66030611$. Note that the value $G(7\pi/10) = 0.97168343$ is adjusted in order to avoid undesirable overshoot behavior in the transfer function. The coefficients for the three temporal schemes are listed in Table 1. The temporal transfer function profiles are shown in Fig. 8, along with that of the one-half rule and Fourier smoothing filters.

As with the frequency filtering procedure, the temporal filtering procedure must be performed on the solution array **Q** before each nonlinear computation and on the nonlinear forcing array **R** after each computation. There is no need to transform between the time and frequency domains.

5.2. Temporal domain filtering results

In this section, the results for filtering the HDHB system in the time domain are presented. The solution for the HDHB3 system is computed using the explicit, implicit tridiagonal, and implicit pentadiagonal temporal filtering schemes described in Table 1. The amplitude response curves are presented in Fig. 9. For comparison, the response curves for the HB2 system, the unfiltered HDHB3 system, and the HDHB3 system filtered by Fourier smoothing in the frequency domain are included.

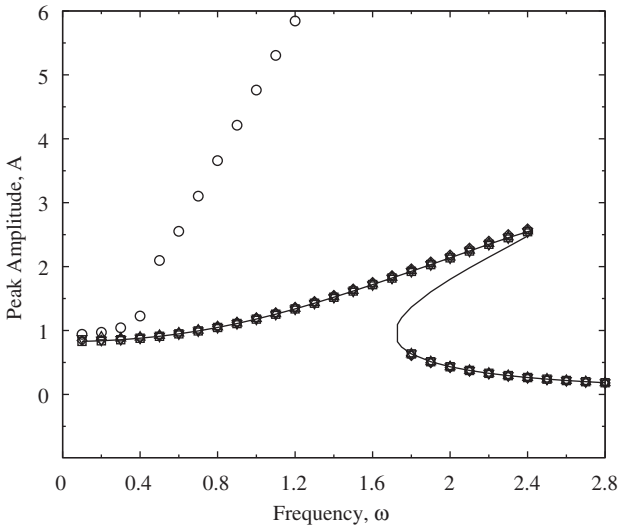

Fig. 9. Peak response amplitudes for the filtered HDHB3 system. HDHB3 filtered by the explicit temporal scheme (diamond); HDHB3 filtered by the implicit tridiagonal temporal scheme (triangle up); HDHB3 filtered by the implicit pentadiagonal temporal scheme (triangle down); HDHB3 filtered by frequency domain Fourier smoothing (square); unfiltered HDHB3 (circle); HB2 (solid line).

From Fig. 9, the filtered HDHB3 system successfully converges to physical response amplitudes for each of the presented temporal filtering schemes. For this case, the temporal filters appear to qualitatively and quantitatively match the results yielded by Fourier smoothing in the frequency domain.

The numerical accuracy of the temporally filtered HDHB system will now be examined by comparing the results with the true response amplitude given in Section 4.3. Again, we compare the peak response amplitudes generated by frequency marching up the backbone curve to $\omega = 2.0$ as a function of the number of harmonics N_H. In Fig. 10, the peak amplitudes are compared for the explicit, implicit tridiagonal and implicit pentadiagonal temporal filters. The response of the unfiltered HDHB system and the HDHB system filtered by Fourier smoothing in the frequency domain are included for comparison.

From Fig. 10, the temporal filtering schemes successfully result in physical convergence of the HDHB system for all harmonics when the initial conditions are generated by backbone marching. From the displacement error curves, the overall accuracy of the implicit schemes is greater than that provided by the explicit scheme. This is true for all harmonics except for $N_H = 3$, where it is suspected that the explicit scheme enjoys a fortuitous cancellation of harmonic and aliasing error. Moreover, the pentadiagonal implicit scheme generally provides greater accuracy than the tridiagonal implicit scheme. Note that for up to approximately $N_H = 10$, the implicit pentadiagonal scheme closely matches, if not outperforms, the Fourier smoothing method. For higher harmonics, the percent displacement error decreases for the temporal schemes at a slower rate than the Fourier smoothing method. This effect is expected due to the finite order of the temporal filtering schemes and becomes less pronounced as the number of stencil points N_S for the temporal schemes is increased.

Using the same procedure described in Sections 3.1 and 4.3, we investigate the statistical behavior of the solutions generated by the temporally filtered HDHB system using Monte Carlo simulation for 10^4 initial conditions at $\omega = 2.0$. In Fig. 11, histograms are presented for the HDHB4 system filtered by the explicit, implicit tridiagonal and implicit pentadiagonal temporal schemes.

Observe in Fig. 11 that the temporal filtering schemes significantly decrease the number of spurious solutions compared to the 127 generated by the unfiltered HDHB4 system in Fig. 2. The number of unique solutions admitted by the explicit, implicit tridiagonal, and implicit pentadiagonal schemes are 7, 43, and 19, respectively. This behavior is quite similar to that of the Fourier smoothing method in Fig. 6, but with one subtle difference. The peak amplitudes of the spurious solutions admitted by the temporal filtering scheme are orders of magnitude larger compared to Fourier smoothing.

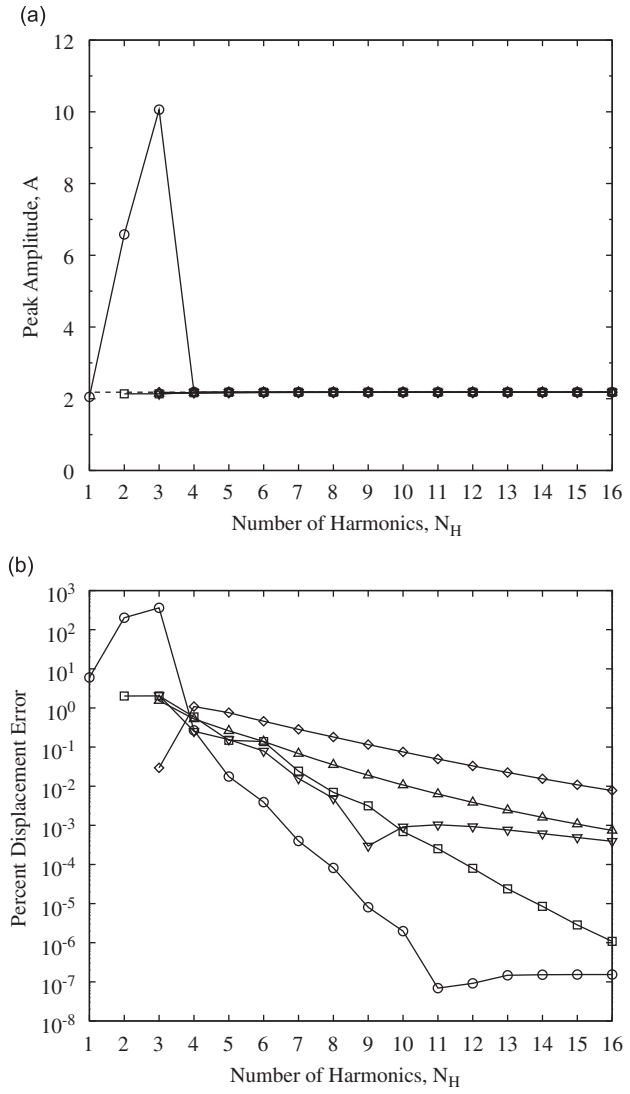

Fig. 10. Numerical accuracy comparison of the filtered and unfiltered HDHB systems. (a) Peak response amplitudes generated by marching up the backbone curve to $\omega = 2.0$. (b) Percent error of the peak amplitudes compared to the true response. True response (dashed line); explicit temporal scheme (diamond); implicit tridiagonal temporal scheme (triangle up); implicit pentadiagonal temporal scheme (triangle down); frequency domain Fourier smoothing (square); unfiltered (circle).

The probability of the temporal filtering schemes converging to a physically legitimate solution can be deduced from Fig. 11. In the explicit, implicit tridiagonal, and implicit pentadiagonal simulations, the probability of converging to the (lower, unstable, upper) branches are (0.029, 0.025, 0.932), (0.028, 0.023, 0.788), and (0.027, 0.017, 0.539), respectively. In sum, the probability of converging to a physical solution for the explicit, implicit tridiagonal, and implicit pentadiagonal simulations are 0.986, 0.839, and 0.583, respectively.

Next, the Monte Carlo simulations are reproduced for various values of N_H. The results are presented in Fig. 12 where up to10 harmonics are retained. For comparison, the results for the unfiltered HDHB system, along with that of the HDHB system filtered by Fourier smoothing in the frequency domain, are included.

In Fig. 12, first note that the temporal filtering schemes always reduce the number of spurious solutions compared to the unfiltered system. Also, the probability of converging to a physical solution is always

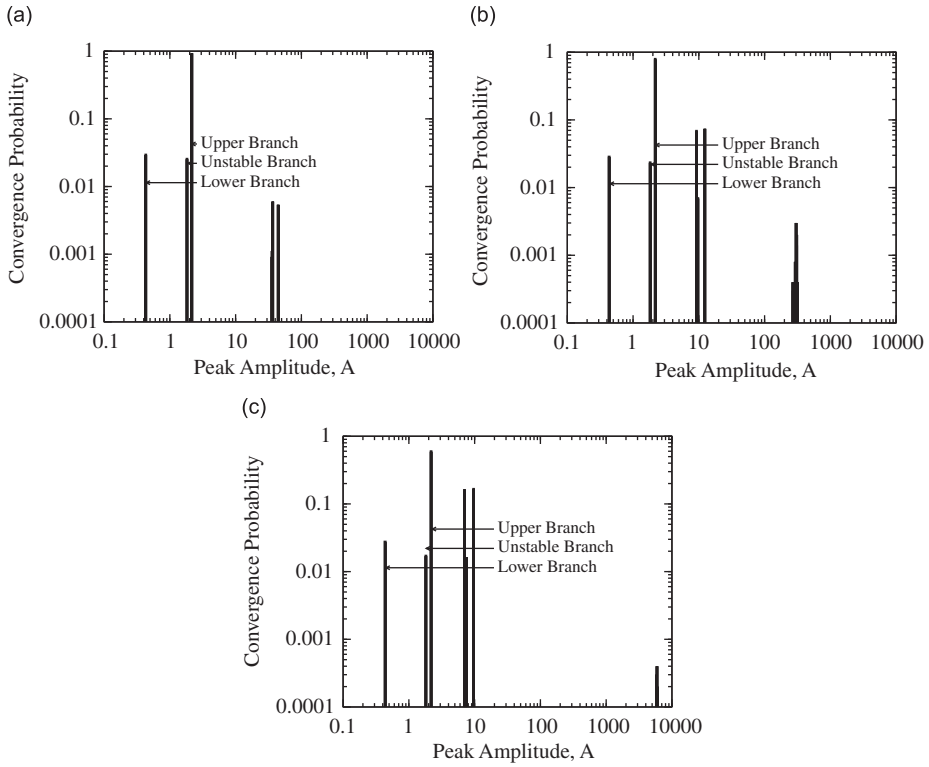

Fig. 11. Monte Carlo histograms for the solutions of the temporally filtered HDHB4 system at $\omega = 2.0$. (a) Results for the explicit temporal filtering scheme. (b) Results for the implicit tridiagonal temporal filtering scheme. (c) Results for the implicit pentadiagonal temporal filtering scheme.

improved over the unfiltered system. For lower harmonics, i.e. $N_H \leqslant 4$, the explicit temporal scheme and Fourier smoothing are more effective at decreasing the number of spurious solutions compared to the two implicit temporal schemes. For higher harmonics, it is difficult to deduce any trends regarding the number of spurious solutions. It is also worth noting that the explicit temporal scheme generally has a higher probability of converging to physical solutions compared to the Fourier smoothing method or implicit temporal schemes.

5.3. Temporal domain filtering summary

The Padé type temporal filtering results compare favorably with the frequency domain Fourier smoothing results. When the initial conditions are generated by backbone marching, the temporally filtered HDHB system only converges to physical solutions. Overall, the implicit temporal schemes provide greater numerical accuracy compared to the explicit scheme. For lower numbers of retained harmonics, i.e. $N_H \leqslant 10$, the implicit pentadiagonal scheme provides similar numerical accuracy compared to that of the Fourier smoothing method. It also appears that when compared to the Fourier smoothing method, the solutions generated by filtering in the time domain converge (with increasing N_H) to the true solution at a slower rate. This is more than likely due to the finite order of the temporal filters. Including more stencil points minimizes this effect.

6. Computational economy

In this section, the computational economy of each filtering method is compared. The ultimate goal is to select the filter that provides the best accuracy per unit computation time. It can be shown that the majority of the expense in the HDHB methodology is consumed by solving the nonlinear system of algebraic equations in Eq. (11). In this study we have solved the system using a Newton–Raphson method. The same sort of

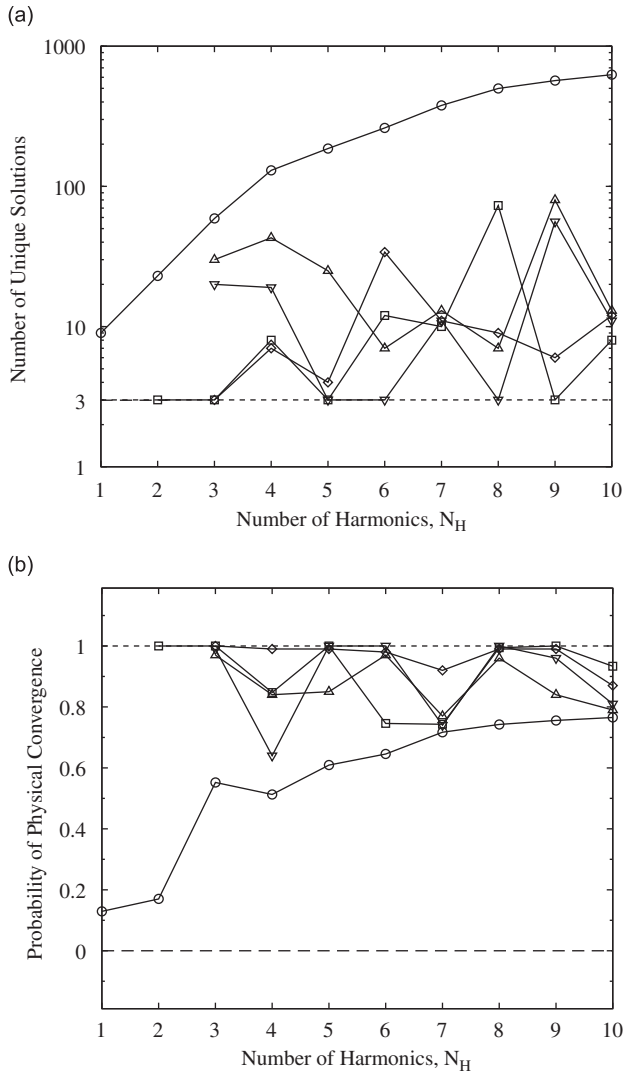

Fig. 12. Solution characteristics for the filtered and unfiltered HDHB systems at $\omega = 2.0$. (a) Number of unique solutions. (b) Probability of converging to a physically meaningful solution. Explicit temporal scheme (diamond); implicit tridiagonal temporal scheme (triangle up); implicit pentadiagonal temporal scheme (triangle down); frequency domain Fourier smoothing (square); unfiltered (circle).

computational economy analysis presented here could be performed for other solution methods such as pseudo-time-marching.

Typically, several Newton–Raphson iterations are required to converge to the solution. Each filtering technique requires a certain number of floating point operations (FLOPs) per iteration (FPI). For simplicity, it is assumed that each filtering technique requires the same number of iterations to converge to the solution. In addition, any computational cost outside the Newton–Raphson procedure is neglected. These are inexact, yet reasonable assumptions that allow for the direct comparison of the accuracy provided by each filtering technique as a function of FPI. The number of FPIs required to solve the HDHB system by the Newton–Raphson method for each filtering technique can be approximated by

$$\mathrm{FPI} \approx \lambda_3 N_T^3 + \lambda_2 N_T^2 + \lambda_1 N_T. \tag{22}$$

where the coefficients λ_n are determined for each filtering technique. When N_T is large, λ_3 dominates, while λ_2 and λ_1 become less significant. The coefficients for the various filtering methods are presented in Table 2.

Note in Table 2 that temporal filtering only increases λ_2 and λ_1, while frequency filtering also increases λ_3. At the low-end stencil limit of the temporal methods, i.e. $N_H = 3$ and $N_T = 7$, frequency filtering is more expensive in terms of FPI than the explicit and implicit tridiagonal schemes, but less expensive than the implicit pentadiagonal scheme. When the number of harmonics used for computation is greater than or equal to four, i.e. $N_H \geqslant 4$ and $N_T \geqslant 9$, frequency filtering becomes far more expensive than any temporal filtering scheme.

Computational economy implies minimizing the number of FPIs required to produce a given displacement error. The economy for the various filtering methods can be compared by combining the results of Eq. (22), Table 2 and Fig. 10. In Fig. 13, the percent displacement error of the peak amplitudes for each filtering method is shown as a function of FPI.

From Fig. 13, the seven point temporal filtering schemes provide the best computational economy when a moderate number of harmonics are retained. Compared to the Fourier smoothing method, the explicit scheme is more economical for $N_H \leqslant 5$, the implicit tridiagonal scheme is more economical for $N_H \leqslant 7$, and the implicit pentadiagonal scheme is more economical for $N_H \leqslant 10$. Within the range of $4 \leqslant N_H \leqslant 10$, the seven point temporal implicit pentadiagonal scheme provides the best overall computational economy while sufficiently decreasing the occurrence of spurious solutions. For large values of N_H, the computational economy of the temporal filtering schemes can be further improved by increasing the stencil size.

Table 2
Approximate FPI coefficients for the various filtering methods.

Filtering method	λ_3	λ_2	λ_1
Unfiltered	18	60	23
One-half rule	50	84	27
Fourier smoothing	50	84	27
Temporal explicit (7pt)	18	116	51
Temporal implicit tridiagonal (7pt)	18	172	79
Temporal implicit pentadiagonal (7pt)	18	308	147

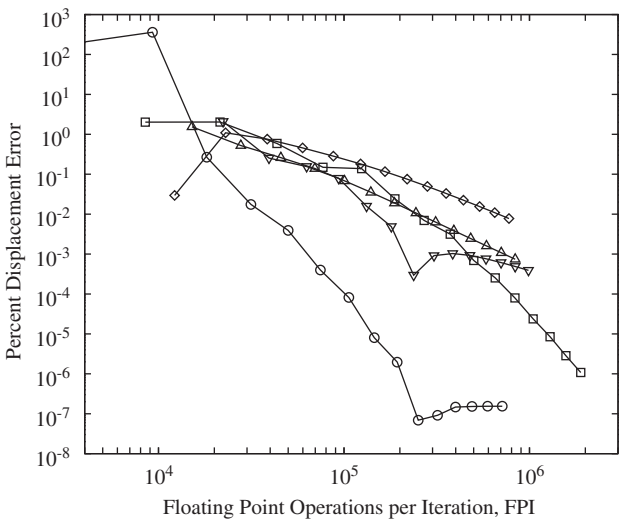

Fig. 13. Computational economy comparison for the temporal and frequency domain filtering schemes. Explicit temporal scheme (diamond); implicit tridiagonal temporal scheme (triangle up); implicit pentadiagonal temporal scheme (triangle down); frequency domain Fourier smoothing (square); unfiltered (circle).

7. Conclusions

Using a Duffing oscillator as a prototypical nonlinear dynamical system, it is demonstrated that aliasing terms in the HDHB method arise due to the treatment of nonlinearities in the system and result in numerical instability. It is shown that the aliasing terms can be completely removed, or sufficiently reduced, through application of low-pass filtering techniques in the frequency domain. The sharp-cutoff frequency domain filter based upon the one-half rule completely removes the aliasing terms, but results in poor harmonic approximation error. The Fourier smoothing frequency domain filter sufficiently reduces the occurrence of aliasing terms while retaining a greater portion of wavenumbers, thus providing a better approximation to the true solution. The drawback to filtering in the frequency domain is that coordinate transformations between the time and frequency domains are required, which result in unnecessary computational expense. Filtering in the frequency domain is particularly expensive for higher dimensional systems such as those encountered in computational fluid and structural dynamics. As an alternative, temporal filters based on a Padé type compact finite difference scheme with spectral-like resolution are designed to replicate the performance of the Fourier smoothing filter. Three temporal filtering schemes (explicit, implicit tridiagonal and implicit pentadiagonal) are constructed on a seven-point stencil. Monte Carlo simulations demonstrate that the temporal filtering schemes and the Fourier smoothing method decrease the number of spurious solutions yielded by the HDHB system and increase the probability of converging to a physical solution. The peak displacement errors yielded by the temporal filtering schemes and the Fourier smoothing method compare favorably, with the implicit pentadiagonal scheme giving the closest representation. When computational expense in terms of FLOPs for a given solution error is considered, the temporal filters outperform the Fourier smoothing method. Of the three temporal filters, the implicit pentadiagonal scheme provides the best overall computational economy.

The temporal and frequency filtering techniques for the HDHB solution of a Duffing oscillator presented in this study may serve as a framework for solving more complex nonlinear dynamical systems, including those encountered in large scale computational fluid and structural dynamics.

References

[1] J.P. Thomas, E.H. Dowell, K.C. Hall, Nonlinear inviscid aerodynamic effects on transonic divergence, flutter, and limit cycle oscillations, *AIAA Journal* 40 (4) (2002) 638–646.
[2] K.C. Hall, J.P. Thomas, W.S. Clark, Computation of unsteady nonlinear flows in cascades using a harmonic balance technique, *AIAA Journal* 40 (5) (2002) 879–886.
[3] P. Beran, C. Pettit, A direct method for quantifying limit cycle oscillation response characteristics in the presence of uncertainties, AIAA Paper 2004-1695.
[4] Y. Kim, S. Noah, Y. Choi, Periodic response of multi-disk rotors with bearing clearances, *Journal of Sound and Vibration* 144 (3) (1991) 381–395.
[5] M. McMullen, A. Jameson, J. Alonso, Acceleration of convergence to a periodic steady state in turbomachinery flows, AIAA Paper 2001-0152.
[6] M. McMullen, The Application of Non-linear Frequency Domain Methods to the Euler and Navier–Stokes Equations, PhD Dissertation, Stanford University, 2001.
[7] G. Dimitriadis, Continuation of higher-order harmonic balance solutions for nonlinear aeroelastic systems, *Journal of Aircraft* 45 (2) (2008) 523–537.
[8] L. Liu, J.P. Thomas, E.H. Dowell, P. Attar, A comparison of classical and high dimensional harmonic balance approaches for a Duffing oscillator, *Journal of Computational Physics* 215 (2006) 298–320.
[9] M.D. Greenberg, *Advanced Engineering Mathematics*, Prentice-Hall, Englewood Cliffs, NJ, 1998.
[10] L.N. Virgin, *Introduction to Experimental Nonlinear Dynamics*, Cambridge University Press, Cambridge, 2000.
[11] W. Cheney, D. Kincaid, *Numerical Mathematics and Computing*, Brooks/Cole, Belmont, CA, 2004.
[12] R.C. Maple, P.I. King, P.D. Orkwis, J.M. Wolff, Adaptive harmonic balance method for nonlinear time-periodic flows, *Journal of Computational Physics* 193 (2004) 620–641.
[13] J.P. Boyd, *Chebyshev and Fourier Spectral Methods*, Springer, Berlin, 1989.
[14] C. Canuto, M.Y. Haussaini, A. Quarteroni, T.A. Zang, *Spectral Methods in Fluid Dynamics*, Springer, Berlin, 1987.
[15] S.A. Orszag, On the elimination of aliasing in finite difference schemes by filtering high-wavenumber components, *Journal of the Atmospheric Sciences* 28 (1971) 1074.
[16] T.Y. Hou, R. Li, Computing nearly singular solutions using pseudo-spectral methods, *Journal of Computational Physics* 226 (2006) 379–397.

[17] S.K. Lele, Compact finite difference schemes with spectral-like resolution, *Journal of Computational Physics* 103 (1992) 16–42.
[18] C. Temperton, Algorithms for the solution of cyclic tridiagonal systems, *Journal of Computational Physics* 19 (1975) 317–323.
[19] S. Sebben, B.R. Baliga, Some extensions of tridiagonal and pentadiagonal matrix algorithms, *Numerical Heat Transfer, Part B: Fundamentals* 2B (1995) 323–351.