

Minimizing Aliasing in Multiple Frequency Harmonic Balance Computations

Downloaded from: https://research.chalmers.se, 2024-07-27 08:11 UTC

Citation for the original published paper (version of record): Lindblad, D., Frey, C., Junge, L. et al (2022). Minimizing Aliasing in Multiple Frequency Harmonic Balance Computations. Journal of Scientific Computing, 91(2).

http://dx.doi.org/10.1007/s10915-022-01776-0

N.B. When citing this work, cite the original published paper.

research.chalmers.se offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all kind of research output: articles, dissertations, conference papers, reports etc. since 2004. research.chalmers.se is administrated and maintained by Chalmers Library



Minimizing Aliasing in Multiple Frequency Harmonic Balance Computations

Daniel Lindblad¹ · Christian Frey² · Laura Junge² · Graham Ashcroft² · Niklas Andersson¹

Received: 22 June 2020 / Revised: 8 December 2021 / Accepted: 17 December 2021 © The Author(s) 2022

Abstract

The harmonic balance method has emerged as an efficient and accurate approach for computing periodic, as well as almost periodic, solutions to nonlinear ordinary differential equations. The accuracy of the harmonic balance method can however be negatively impacted by aliasing. Aliasing occurs because Fourier coefficients of nonlinear terms in the governing equations are approximated by a discrete Fourier transform (DFT). Understanding how aliasing occurs when the DFT is applied is therefore essential in improving the accuracy of the harmonic balance method. In this work, a new operator that describe the fold-back, i.e. aliasing, of unresolved frequencies onto the resolved ones is developed. The norm of this operator is then used as a metric for investigating how the time sampling should be performed to minimize aliasing. It is found that a time sampling which minimizes the condition number of the DFT matrix is the best choice in this regard, both for single and multiple frequency problems. These findings are also verified for the Duffing oscillator. Finally, a strategy for oversampling multiple frequency harmonic balance computations is developed and tested.

Keywords Harmonic balance \cdot Aliasing \cdot Condition number \cdot Almost periodic Fourier transform \cdot APFT \cdot Duffing oscillator

Niklas Andersson niklas.andersson@chalmers.se

> Daniel Lindblad daniel.lindblad@chalmers.se

Christian Frey christian.frey@dlr.de

Laura Junge laura.junge@dlr.de

Graham Ashcroft graham.ashcroft@dlr.de

- ¹ Division of Fluid Dynamics, Department of Mechanics and Maritime Sciences, Chalmers University of Technology, Hörsalsvägen 7A, 412 96 Gothenburg, Sweden
- ² Department of Numerical Methods, Institute of Propulsion Technology, Deutsches Zentrum f
 ür Luftund Raumfahrt e.V. (DLR), Linder H
 öhe, 51147 Cologne, Germany

Mathematics Subject Classification 65L70

1 Introduction

Many processes found in engineering and scientific applications are time-periodic. Some examples are fluid flows inside turbomachines, vibrations of structures, or voltages within AC circuits. In order to understand these processes, they may be simulated by integrating their governing equations in time until a periodic solution is obtained. In cases when the initial condition is far from being periodic and/or transient phenomena decay slowly, it may however take a long time until the solution reaches a periodic state. In terms of a numerical simulation that integrates between discrete time steps, this translates into a large number of iterations, and consequently, a high computational cost [10, 17].

An alternative approach that can be more computationally efficient is to directly seek solutions that belong to some finite dimensional vector space spanned by a set of periodic basis functions. This paper considers the special case when these basis functions are sinusoids with different frequencies. In this case, the solution thus takes the form of a truncated Fourier series, for which the unknowns become the Fourier coefficients. There exist several methods in the literature for determining these Fourier coefficients. One approach is to integrate the equations obtained from substituting the Fourier series into the governing equations against each of the basis functions that span the vector space. This represents a Galerkin method, and will yield one equation per basis function. In cases when the governing equations are linear, the equations resulting from the integration uncouple with respect to the Fourier coefficients. If the governing equations on the other hand are nonlinear, each Fourier coefficient will in general be present in each of the final equations. When the Galerkin method is applied to linear problems it is sometimes referred to as the Linear Frequency Domain method [7, 15]. Some applications of the aforementioned approach to nonlinear problems can be found in e.g. [31, 39].

Unfortunately, the Galerkin integral can become very complicated to solve analytically in the nonlinear case [16, 31]. This is especially true when the dimension of the vector space is large and/or when the governing equations are complex in nature. Other methods have therefore been developed for nonlinear problems. One common approach is to require that the equations resulting from the substitution only are satisfied at a discrete set of time instances, rather than in a variational sense as in the Galerkin method. This leads to a class of methods which here will be referred to as harmonic balance methods. Harmonic balance methods can be formulated in both the time and the frequency domain, with the main difference being that the solution variables are time samples in the former, and Fourier coefficients in the latter [10, 26]. Within the electronics community, harmonic balance is most often formulated in the frequency domain [10], and the modern version of harmonic balance is acredited to Nakhla and Valch [34]. Within the fluid dynamics community, the time domain harmonic balance method was first introduced by Hall et al. [16] and the frequency domain formulation by McMullen [33]. A good review of the history and theory of harmonic balance can be found in a two part review paper by Gilmore and Steer [10, 11].

The harmonic balance method is based on applying a discrete Fourier transform (DFT) to either calculate the time derivative term (time-domain formulation), or the Fourier transform of the nonlinear terms (frequency domain formulation) [10, 26]. Applying a DFT to nonlinear problems can however introduce aliasing, which is known to reduce the accuracy of the harmonic balance method [24, 29, 32], and sometimes even lead to nonphysical solutions

[28]. Several methods have therefore been developed to address these problems. Huang and Ekici propose to add a time-spectral viscosity operator to the harmonic balance equations [19]. This operator acts by damping the sinusoids which it is applied to, and has been shown by Huang and Ekici to improve convergence in cases when aliasing otherwise prevents it [19]. Another approach investigated by La Bryer and Attair [28] is to filter the solution between each nonlinear iteration. The main idea of the filtering approach is to limit the amplitude of the sinusoids that have the highest frequencies, since these in general are the ones most contaminated by aliasing. In cases when the nonlinearity of the governing equations is known, exact frequency domain filters which keep the lower frequency sinusoids perfectly free from aliasing can in fact be constructed by using a sharp cut-off frequency based on Orszag's rule [28, 36]. It should also be noted that a frequency domain filter with a sharp cut-off frequency is equivalent to the oversampling strategy employed by Frey et al. in their harmonic balance solver [8].

In addition to filtering, a filtered inverse discrete Fourier transform has also been used to improve the convergence and stability of the harmonic balance method [3, 18]. This approach is particularly useful in cases when discontinuities and/or strong gradients lead to slow harmonic convergence (Gibbs phenomenon). The use of a filtered reconstruction operator is thouroughly discussed by Gottlieb and Shu [12]. For its application to Fourier spectral methods alongside so-called reprojection approaches the reader to referred to [9].

Harmonic balance methods were originally developed for problems where the solution contains one fundamental frequency and its harmonics. Several problems in nature can however be expected to have solutions which contain a more general set of frequencies. In these cases, the solution belongs to a wider class of functions known as almost periodic functions [1, 26]. A fundamental problem that arises when the harmonic balance method is applied to almost periodic functions is to construct the DFT. In the original harmonic balance method, the standard discrete Fourier transform based on M uniformly distributed points between 0 and the reciprocal of the lowest resolved frequency is almost exclusively employed. Here, $M \ge 2K + 1$, where K is the number of frequencies included in the Fourier series expansion. The reason for choosing this sampling is that it satisfies the assumptions of the Whittaker-Kotel'nikov–Shannon sampling theorem [23, 38, 40]. A uniform time sampling that satisfies the sampling theorem can also be constructed for almost periodic signals as long as the frequencies considered are commensurable. Such a sampling may however require a very large number of sampling points, which makes this approach unfavourable from a computational perspective [21, 26]. Other approaches have therefore been developed for cases when the solution is an almost periodic function. Some of these are based on introducing a basis for the frequency set, such that each frequency can be written as an integer linear combination of a finite set of fundamental frequencies [26]. The method developed by Frey et al. [8] takes into account only frequencies which are multiples of one the fundamental frequencies and applies the so-called *harmonic set approach*. The advantages of the harmonic set approach are that its computational cost scales linearly with the number of frequencies considered and that the standard discrete Fourier transform can be used within each harmonic set. The harmonic set approach does however neglect nonlinear coupling between different harmonic sets, except through common frequencies such as the zeroth frequency. These coupling terms may be accounted for by introducing several time variables, one for each fundamental frequency in the basis [14, 22]. This approach once again relies on the standard discrete Fourier transform with uniform time samples, but requires a substantially larger computational cost than the harmonic set approach.

A third approach, which is the main topic of this paper, is to employ the almost periodic Fourier transform (APFT) introduced by Kundert et al. [25, 26] in the harmonic balance

computation. The APFT differs from the standard discrete Fourier transform in two ways. First, it considers an arbitrary set of frequencies and is therefore well suited for cases when the solution is an almost periodic function. Secondly, it does not require that the time samples are distributed uniformly like in the standard discrete Fourier transform. Instead, Kundert et al. suggest that the time sampling is chosen so that it minimizes the condition number of the DFT matrix used in the APFT [25, 26]. Finding a time sampling that satisfies this criterion for an arbitrary set of frequencies is however a very complex problem, that to the authors' knowledge has not yet been solved analytically. In the original work on harmonic balance methods for arbitrary frequency sets, Chua and Ushida [2] employ a uniform discretization and use oversampling to avoid ill-conditioning. The same strategy has also been used in recent years by Ekici and Hall [5, 6]. Kundert et al. [25, 26] were the first to obtain a well conditioned DFT matrix based on 2K + 1 time samples. This was done using their Near-Orthogonal Selection Algorithm. Since the work of Kundert et al., several other algorithms have been developed to compute a time sampling that minimizes the condition number of the DFT matrix, see e.g. [13, 21, 27, 35, 37].

The rationale of choosing a time sampling that minimizes the condition number of the DFT matrix is that it limits the effect that a perturbation of the sampled signal can have on the resulting DFT. Since all unresolved frequencies manifest themselves as perturbations of the sampling, this implies that a low condition number can limit the amount of aliasing produced by the DFT [26]. In [26], it is also noted that the condition number does not say anything about how much the amplitude of a particular unresolved sinusoid affects the resolved ones. This implies that the condition number in itself can not be used to identify an alias free DFT. In order to do this, the mapping of the unresolved sinusoids onto the resolved ones must be defined [26]. In this work, an operator that describes this mapping has been developed. The benefit of this operator, hereinafter referred to as the *alias operator*, is that it directly defines how a particular time sampling affects aliasing. As will be shown later, the norm of this operator also provides an a-priori bound on the amount of aliasing that a particular time sampling gives rise to.

In this paper, the relation between the alias operator and the condition number of the DFT matrix employed in the APFT will be investigated. After this, the relation between the norm of the alias operator, the condition number of the DFT matrix and the actual alias error obtained in several harmonic balance computations will be investigated. Finally, the benefits of employing oversampling in multiple frequency harmonic balance computations based on the APFT will be investigated.

2 Method

2.1 Harmonic Balance Method

The purpose of this paper is to study how aliasing occurs in harmonic balance solutions of first order ordinary differential equations of the following form

$$\frac{dq}{dt} + f(q,t) = 0 \tag{1}$$

Here, $q \in \mathbb{R}^N$ is a vector that contains the unknown solution and $f : \mathbb{R}^N \times \mathbb{R} \to \mathbb{R}^N$ is a nonlinear function that describes the dynamics of the problem. The nonlinearity of f(q, t) implies that an almost periodic solution to Eq. (1) can contain an infinite number of sinusoids with different frequencies. Computing the amplitude of all these sinusoids is however not

feasible from a numerical perspective. Therefore, approximate solutions that are spanned by a limited number of sinusoids are considered instead

$$q(t) \approx \sum_{k: \ \omega_k \in \Lambda} \hat{q}_k e^{i\omega_k t} \tag{2}$$

Here, Λ represents the set of frequencies included in the series expansion

$$A = \{0, \omega_1, \omega_2, \dots, \omega_K, \omega_{-K}, \dots, \omega_{-1}\}$$
(3)

Note that the only requirement put on these frequencies is that $\omega_{-k} = -\omega_k$.

If a Galerkin projection of Eq. (1) onto the subspace spanned by the sinusoids in Eq. (2) is performed, the following is obtained

$$\boldsymbol{\Omega}_{\Lambda}\hat{\mathbf{q}}_{\Lambda} + \hat{\mathbf{f}}_{\Lambda}(\hat{\mathbf{q}}_{\Lambda},\Lambda) = 0 \tag{4}$$

Here, $\Omega_{\Lambda} = \Omega_{\Lambda} \otimes \mathbf{I}$, where \otimes is the Kronecker product, \mathbf{I} is an identity matrix of size $N \times N$, and

$$\Omega_{\Lambda} = \begin{bmatrix}
0 & 0 & \cdots & & 0 \\
0 & i\omega_{1} & 0 & \cdots & & 0 \\
\vdots & 0 & i\omega_{2} & 0 & \cdots & 0 \\
\vdots & 0 & \ddots & & 0 \\
\vdots & 0 & \ddots & & 0 \\
\vdots & \vdots & & & 0 \\
0 & 0 & 0 & 0 & 0 & i\omega_{-1}
\end{bmatrix}$$
(5)

The vectors $\hat{\mathbf{q}}_{\Lambda}$ and $\hat{\mathbf{f}}_{\Lambda}(\hat{\mathbf{q}}_{\Lambda}, \Lambda)$ in Eq. (4) further consist of 2K + 1 sub-vectors, in which the *k*th sub-vector contains \hat{q}_k and \hat{f}_k respectively.

Calculating the Galerkin projection of Eq. (1) can be very complicated in cases when the number of sinusoids considered in Eq. (2) is large, and/or the function f(q, t) is very complex in nature. In order to overcome this difficulty, the harmonic balance method approximates $\hat{\mathbf{f}}_{\Lambda}(\hat{\mathbf{q}}_{\Lambda}, \Lambda)$ by a discrete Fourier transform (DFT) of f(q, t) instead

$$\boldsymbol{\varrho}_{\Lambda}\hat{\mathbf{q}}_{\Lambda} + \mathbf{E}_{\Lambda}(\mathbf{t})f(\mathbf{E}_{\Lambda}^{-1}(\mathbf{t})\hat{\mathbf{q}}_{\Lambda}, \mathbf{t}) = 0$$
(6)

The DFT of f(q, t) in Eq. (6) is performed in two steps. In the first step, f(q, t) is evaluated at a set of time instances **t**

$$\mathbf{t} = \begin{bmatrix} t_0 \\ t_1 \\ \vdots \\ t_{M-1} \end{bmatrix}$$
(7)

This will in turn require the realization of Eq. (2) at these time instances

$$\mathbf{q} = \mathbf{E}_{\Lambda}^{-1}(\mathbf{t})\hat{\mathbf{q}}_{\Lambda} \tag{8}$$

Here, $\mathbf{E}_{\Lambda}^{-1}(\mathbf{t}) = E_{\Lambda}^{-1}(\mathbf{t}) \otimes \mathbf{I}$, and

$$E_{\Lambda}^{-1}(\mathbf{t}) = \begin{bmatrix} 1 & e^{i\omega_{1}t_{0}} & e^{i\omega_{2}t_{0}} & \dots & e^{i\omega_{-1}t_{0}} \\ 1 & e^{i\omega_{1}t_{1}} & e^{i\omega_{2}t_{1}} & \dots & e^{i\omega_{-1}t_{1}} \\ \vdots & \vdots & \vdots & & \vdots \\ 1 & e^{i\omega_{1}t_{M-1}} & e^{i\omega_{2}t_{M-1}} & \dots & e^{i\omega_{-1}t_{M-1}} \end{bmatrix}$$
(9)

Deringer

In the second step, the sampling of f(q, t) is transformed back into the frequency domain using the inverse of $\mathbf{E}_A^{-1}(\mathbf{t})$, here denoted $\mathbf{E}_A(\mathbf{t})$. For this inverse to be well defined, the columns in $\mathbf{E}_A^{-1}(\mathbf{t})$ must be linearly independent. Due to the structure of $\mathbf{E}_A^{-1}(\mathbf{t})$, this holds true if M = 2K + 1 time instances can be found such that the columns in $E_A^{-1}(\mathbf{t})$ are linearly independent. As it turns out, this is always possible. In order to prove this, note that for an equidistant sampling, the entries of $E_A^{-1}(\mathbf{t})$ take the form

$$E_{\Lambda}^{-1}(\mathbf{t})_{m,n} = \left(e^{i\Lambda_n \Delta t}\right)^{m-1} \tag{10}$$

For an equidistant time sampling, $E_{\Lambda}^{-1}(\mathbf{t})$ thus corresponds to a transposed Vandermonde matrix. Its determinant is therefore given by

$$\det\left(E_{\Lambda}^{-1}(\mathbf{t})\right) = \prod_{m < n} \left(e^{i\Lambda_m \Delta t} - e^{i\Lambda_n \Delta t}\right)$$
(11)

It is easy to verify that this expression is non-zero for almost all choices of Δt . This shows that for any given set of frequencies one can find an equidistant sampling point distribution **t** such that $E_{\Lambda}^{-1}(\mathbf{t})$, and thereby $\mathbf{E}_{\Lambda}^{-1}(\mathbf{t})$, is invertible.

The time sampling is usually chosen to minimize the condition number of $E_A^{-1}(\mathbf{t})$. Note that a small condition number automatically ensures that $E_A^{-1}(\mathbf{t})$ is invertible. Some authors use a numerical optimization procedure to minimize the condition number. It is important to note that the cost of this optimization procedure is usually negligible compared to solving the harmonic balance system.

In the literature, the discrete Fourier transform defined by the matrices $E_A^{-1}(\mathbf{t})$ and $E_A(\mathbf{t})$ is often referred to as the almost periodic Fourier transform (APFT) [26]. The APFT represents a generalization of the standard DFT to arbitrary sets of frequencies and sampling points, and will in fact be equivalent to the standard DFT when the frequencies in Λ are integer multiples of a single base frequency, and the time samples are distributed uniformly between zero and the reciprocal of this base frequency.

Equation (6) represents the frequency domain formulation of the harmonic balance method. An equivalent formulation in the time domain may also be expressed as follows

$$\mathbf{E}_{\Lambda}^{-1}(\mathbf{t})\boldsymbol{\mathcal{Q}}_{\Lambda}\mathbf{E}_{\Lambda}(\mathbf{t})\mathbf{q} + \mathbf{E}_{\Lambda}^{-1}(\mathbf{t})\mathbf{E}_{\Lambda}(\mathbf{t})f(\mathbf{q},\mathbf{t}) = 0$$
(12)

This formulation is easily obtained by premultiplying Eq. (6) from the left by $\mathbf{E}_{\Lambda}^{-1}(\mathbf{t})$ and then using Eq. (8). The matrix $\mathbf{E}_{\Lambda}^{-1}(\mathbf{t})\mathbf{E}_{\Lambda}(\mathbf{t})$ in the above equation will further be equal to the identity matrix when M = 2K + 1 time samples are employed. If more than M = 2K + 1 time samples on the other hand are used, this matrix will represent a projection that corresponds to the application of a modal filter.

The frequency domain formulation of the harmonic balance method presented in Eq. (6) will be used for the remainder of this paper. Due to the equivalence of Eq. (6) and (12), however, the results presented will also apply to the time domain formulation of the harmonic balance method.

2.1.1 Error Sources

It is well known that the harmonic balance method can give rise to two types of errors: aliasing and harmonic truncation [10, 11, 26]. Aliasing occurs because f(q, t) can contain more frequencies that those accounted for by the DFT in Eq. (6). Aliasing will on the other hand not occur if Eq. (4) is solved. This is because the Galerkin projection in this case

corresponds (up to a constant) to the standard formula for calculating Fourier coefficients. Independently on whether the harmonic balance or Galerkin method is used, however, the harmonic truncation error will always be present. This error occurs as a result of the fact that the contribution of the unresolved sinusoids to the solution, including their nonlinear coupling with the resolved sinusoids, is neglected.

Both the alias error and the harmonic truncation error can naturally be reduced by including more frequencies in Λ . This is however not always possible in practical applications since the computational cost of the harmonic balance method scales almost linearly with the size of Λ , here denoted # Λ . It is therefore often hard to avoid aliasing and harmonic truncation errors in the solution. This highlights the importance of understanding how these errors occur, and how they can be limited. In this paper, the focus is put on the aliasing error.

2.2 Alias Operator

As noted in the previous section, aliasing occurs when f(q, t) contains more frequencies than those accounted for by the DFT in Eq. (6). Let Λ'' denote the union of Λ and the set of all frequencies that are contained in f(q, t). In cases when f(q, t) is a polynomial in q, with coefficients that are finite Fourier sums in t, the set Λ'' will be finite. This in turn implies that the signal $f(\mathbf{q}, \mathbf{t})$ can be written as

$$f(\mathbf{q}, \mathbf{t}) = \mathbf{E}_{A''}^{-1}(\mathbf{t})\hat{\mathbf{f}}_{A''}$$
(13)

Note that $\Lambda \subsetneq \Lambda''$. From this, it follows that Eq. (13) can be rewritten as

$$f(\mathbf{q}, \mathbf{t}) = \mathbf{E}_{\Lambda}^{-1}(\mathbf{t})\hat{\mathbf{f}}_{\Lambda} + \mathbf{E}_{\Lambda'' \setminus \Lambda}^{-1}(\mathbf{t})\hat{\mathbf{f}}_{\Lambda'' \setminus \Lambda}$$
(14)

Based on this expression, the DFT of $f(\mathbf{q}, \mathbf{t})$ may now be expressed as

$$\tilde{\mathbf{f}}_{A} = \mathbf{E}_{A}(\mathbf{t}) f(\mathbf{q}, \mathbf{t})
= \mathbf{E}_{A}(\mathbf{t}) \left(\mathbf{E}_{A}^{-1}(\mathbf{t}) \hat{\mathbf{f}}_{A} + \mathbf{E}_{A'' \setminus A}^{-1}(\mathbf{t}) \hat{\mathbf{f}}_{A'' \setminus A} \right)
= \hat{\mathbf{f}}_{A} + \mathbf{E}_{A}(\mathbf{t}) \mathbf{E}_{A'' \setminus A}^{-1}(\mathbf{t}) \hat{\mathbf{f}}_{A'' \setminus A}$$
(15)

This shows that the Fourier coefficients obtained from the DFT of $f(\mathbf{q}, \mathbf{t})$ consist of two terms. The first term represents the correct (alias free) Fourier coefficients corresponding to the resolved frequencies, and the second term the fold-back of the remaining Fourier coefficients onto the resolved ones. The operator $E_A(\mathbf{t})E_{A''\setminus A}^{-1}(\mathbf{t})$ that defines the fold-back of the unresolved Fourier coefficients in Eq. (15) will be referred to as the *alias operator* for the remainder of this paper. Based on Eq. (15), it is easy to show that the norm of this operator puts the following bound on the alias error

$$\frac{\|\hat{\mathbf{f}}_{A} - \hat{\mathbf{f}}_{A}\|_{2}}{\|\hat{\mathbf{f}}_{A''\setminus A}\|_{2}} \leq \|\mathbf{E}_{A}(\mathbf{t})\mathbf{E}_{A''\setminus A}^{-1}(\mathbf{t})\|_{2}$$
$$= \|E_{A}(\mathbf{t})E_{A''\setminus A}^{-1}(\mathbf{t})\|_{2}$$
(16)

Provided that Λ'' is known, this relation may be used to obtain an a-priori bound on the alias error that a particular time-sampling **t** can give rise to. As will be shown next, however, determining Λ'' for a particular problem is not a trivial task.

Deringer

2.2.1 Selection of Unresolved Frequencies

The set Λ'' is completely defined by the function f(q, t) and the set Λ . In the present work, it is assumed that f(q, t) is a polynomial of degree p in q. This assumption allows the set Λ'' to be calculated by simple addition and subtraction of the frequencies in Λ . Clearly, this assumption is not valid for all problems. Despite this, however, it can still be a good approximation provided that f(q, t) itself approximates an analytic function, i.e., one that is locally given by a convergent power series (which then coincides with its Taylor series).

The choice of Λ'' directly influences how well the norm of the alias operator will represent the actual alias error obtained in a harmonic balance simulation. To see this, note that Eq. (16) represents the largest possible alias error that any $\hat{\mathbf{f}}_{\Lambda''\setminus\Lambda}$ can give rise to. This means that the bound in Eq. (16) may not be representative if some elements in $\hat{\mathbf{f}}_{\Lambda''\setminus\Lambda}$ are negligible for a given problem. It is also important to keep in mind that Eq. (16) gives a relative bound on the alias error. As such, if Λ has been selected to contain all relevant frequencies for a given problem, then $\|\hat{\mathbf{f}}_{\Lambda''\setminus\Lambda}\|_2 \approx 0$ and consequently, the actual alias error will be small independent of the norm of the alias operator.

2.2.2 Relation to Condition Number

The condition number of $E_A^{-1}(\mathbf{t})$ has been successfully used by several authors in the past as a measure for selecting the time sampling in multiple frequency harmonic balance computations [13, 21, 25–27, 35, 37]. This raises the question of how the condition number relates to the alias operator defined in the present work. The condition number can be defined based on the l_2 norm as

$$\kappa(E_{\Lambda}^{-1}(\mathbf{t})) = \|E_{\Lambda}^{-1}(\mathbf{t})\|_{2} \|E_{\Lambda}(\mathbf{t})\|_{2}$$
(17)

This definition can be used to derive a bound on the norm of the alias operator as follows

$$\|E_{A}(\mathbf{t})E_{A''\setminus A}^{-1}(\mathbf{t})\|_{2} \leq \|E_{A}(\mathbf{t})\|_{2}\|E_{A''\setminus A}^{-1}(\mathbf{t})\|_{2}$$

= $\kappa(E_{A}^{-1}(\mathbf{t}))\underbrace{\frac{\|E_{A''\setminus A}^{-1}(\mathbf{t})\|_{2}}{\|E_{A}^{-1}(\mathbf{t})\|_{2}}}_{\alpha(A,A'',\mathbf{t})}$ (18)

Both matrices in the last equation have entries with norm 1. From this, it follows that there exists an upper bound on $||E_{A''\setminus A}^{-1}(\mathbf{t})||_2$, and a lower bound on $||E_A^{-1}(\mathbf{t})||_2$, both independent of **t**. Equation may therefore be rewritten as

$$\|E_{\Lambda}(\mathbf{t})E_{\Lambda''\setminus\Lambda}^{-1}(\mathbf{t})\|_{2} \le \kappa(E_{\Lambda}^{-1}(\mathbf{t})) \max_{\mathbf{t}} \left\{ \alpha\left(\Lambda,\Lambda'',\mathbf{t}\right) \right\}$$
(19)

This equation shows that the condition number in fact puts a bound on the aliasing operator. Note, however, that this does not imply that a time sampling which minimizes the condition number also minimizes the norm of the alias operator.

The complexity of the alias operator, condition number and α (Λ , Λ'' , **t**) in Eq. (19) makes it very hard to investigate how they relate analytically. Numerical optimizations were therefore used to search for time samplings which give a minimal condition number, minimal norm of the alias operator or maximal value of α (Λ , Λ'' , **t**). These results could then be used to investigate whether a time sampling that gives, e.g., a low condition number, also yields a low norm of the alias operator. For a detailed description of the optimization procedure used in the present work, see "Appendix B".

2.3 Almost Periodic Fourier Transform with Oversampling

In this work, the impact of oversampling on the alias error obtained in a harmonic balance computation has also been investigated. The motivation for doing this is that aliasing can never be eliminated when M = 2K + 1 time samples are employed. This can be proven by first noting that the columns in $E_A^{-1}(\mathbf{t})$ are required to be linearly independent, and therefore form a basis for \mathbb{C}^M . From this, it follows that there exists a nonzero matrix W such that $E_{A''\setminus A}^{-1}(\mathbf{t}) = E_A^{-1}(\mathbf{t})W$. But then, the alias operator in Eq. (15) must be equal to $E_A(\mathbf{t})E_{A''\setminus A}^{-1}(\mathbf{t}) = W \neq 0$, which proves the desired result.

When oversampling is employed, $\mathbf{E}_{\Lambda}(\mathbf{t})$ in Eq. (6) will represent a left inverse to $\mathbf{E}_{\Lambda}^{-1}(\mathbf{t})$. Contrary to the standard inverse of a square matrix, however, the left inverse is not uniquely defined. In cases when Λ'' is finite, it can therefore be tailored such that aliasing is eliminated. In order to prove this, one can start by noting that such a tailored inverse, here denoted $E_{\Lambda,\Lambda''}(\mathbf{t})$, must satisfy the following two conditions

$$E_{\Lambda,\Lambda''}(\mathbf{t})E_{\Lambda}^{-1}(\mathbf{t}) = \mathbf{I},$$
(20)

$$E_{\Lambda,\Lambda''}(\mathbf{t})E_{\Lambda''\setminus\Lambda}^{-1}(\mathbf{t}) = \mathbf{0}.$$
(21)

The above conditions can be combined as follows

$$E_{\Lambda,\Lambda''}(\mathbf{t}) \left[E_{\Lambda}^{-1}(\mathbf{t}) \ E_{\Lambda'' \setminus \Lambda}^{-1}(\mathbf{t}) \right] = \left[\mathbf{I} \ \mathbf{0} \right]$$
(22)

Note that the matrix within the brackets in the left hand side of the above equation corresponds to $E_{A''}^{-1}(\mathbf{t})$. From before, it is known that there exists a time sampling with M'' = #A'' time samples such that this matrix is invertible. From this, it follows that the above relation may be rewritten as

$$E_{A,A''}(\mathbf{t}) = \pi_{A'',A} E_{A''}(\mathbf{t}).$$
(23)

Here, $\pi_{\Lambda'',\Lambda}$ is the projection onto the harmonics in Λ .

The argument above shows that a left inverse that eliminates aliasing always can be defined if Λ'' is finite. Unfortunately, it is more computationally expensive to employ this left inverse since it requires that f(q, t) is evaluated at M'' > 2K + 1 time instances. This motivates the use of a left inverse which requires less than M'' time instances, but still has the potential to reduce aliasing compared to not using oversampling at all. In order to define such a left inverse, introduce a new set of frequencies Λ' that satisfies

$$\Lambda \subset \Lambda' \subseteq \Lambda'' \tag{24}$$

There exists a time sampling with M' = #A' time samples such that $E_{A'}^{-1}(\mathbf{t})$ is invertible. Therefore, the following left inverse may be defined

$$E_{\Lambda,\Lambda'}(\mathbf{t}) = \pi_{\Lambda',\Lambda} E_{\Lambda'}(\mathbf{t}) \tag{25}$$

The idea behind this left inverse is to eliminate aliasing with respect to the frequencies in $\Lambda' \setminus \Lambda$. This left inverse would also completely eliminate aliasing if a time sampling can be found such that the unresolved Fourier coefficients only contribute to the amplitude of the Fourier coefficients whose frequencies are in $\Lambda' \setminus \Lambda$.

An alternative choice for the left-inverse which has been suggested in the literature is the Moore–Penrose inverse of $E_A^{-1}(\mathbf{t})$ [5, 6, 21]. This left inverse represents a least squares projection onto the subspace spanned by the columns in $E_A^{-1}(\mathbf{t})$. From this, it follows the Moore–Penrose inverse can only eliminate aliasing if the columns in $E_{A''\setminus A}^{-1}(\mathbf{t})$ are orthogonal

to those in $E_{\Lambda}^{-1}(\mathbf{t})$ (with respect to the standard inner product on $\mathbb{C}^{M''}$). In cases when the frequencies in Λ are commensurable, and f(q, t) is a polynomial, it is possible to ensure this by selecting a uniform time sampling that satisfies Orszag's rule [21]

$$M'' \ge (p+1)\frac{\omega_{\max}}{\omega_{\text{base}}} + 1 \tag{26}$$

Here, ω_{max} is the largest frequency in Λ and ω_{base} is the greatest common divisor of all frequencies in Λ .

It might also be possible to construct a (possibly non-uniform) time sampling such that the Moore–Penrose inverse eliminates aliasing in the general case, but no such algorithm is known to the authors. As shown in "Appendix A", it is however possible to redefine the inner product on $\mathbb{C}^{M''}$ such that the columns in $E_{A''\setminus A}^{-1}(\mathbf{t})$ become orthogonal to those in $E_{A}^{-1}(\mathbf{t})$. There, it is also shown that the Moore–Penrose inverse defined with respect to this new inner product is equivalent to the left inverse defined by Eq. (23).

3 Results

3.1 Aliasing for $M = #\Lambda$

This section considers the case when the number of sampling points corresponds to the theoretical minimum of M = #A. To begin with, it is demonstrated how $\kappa(E_A^{-1})$ and $||E_A E_{A'' \setminus A}^{-1}||_2$ relate in this case. This is done for both the case when A contains a single as well as multiple fundamental frequencies. After this, it is investigated how $\kappa(E_A^{-1})$ and $||E_A E_{A'' \setminus A}^{-1}||_2$ influence the alias error obtained when the Duffing oscillator is simulated with the harmonic balance method. This is done for both the single and multiple frequency case as well.

3.1.1 Relation Between $\kappa(E_A^{-1})$ and $||E_A E_{A'' \setminus A}^{-1}||_2$

In Sect. 2.2.2 it was shown that $||E_A E_{A''\setminus A}^{-1}||_2$ is bounded by $\kappa(E_A^{-1})$ through Eq. (19). This fact can be illustrated by a diagonal line in a $\kappa(E_A^{-1}) - ||E_A E_{A''\setminus A}^{-1}||_2$ plane, if the slope of this line is taken to be the maximum value of $\alpha(A, A'', \mathbf{t})$. In the same plane, the minimal values of $\kappa(E_A^{-1}) \ge 1$ and $||E_A E_{A''\setminus A}^{-1}||_2 \ge 0^1$ can also be represented by a vertical and a horizontal line respectively. The region between these three lines will then contain all possible values of $\kappa(E_A^{-1})$ and $||E_A E_{A''\setminus A}^{-1}||_2$ that a time sampling can give rise to. In Fig. 1a, this region is illustrated for the case when A contains a single fundamental frequency and A'' is generated by a second degree polynomial (p = 2). The case when A'' is generated by a third degree polynomial (p = 3) is further depicted in Fig. 1b. The limit values of $\kappa(E_A^{-1})$, $||E_A E_{A''\setminus A}^{-1}||_2$ and $\alpha(A, A'', \mathbf{t})$ shown in these figures have all been computed with the OPT algorithm [13], as described in "Appendix B".

A large set of random time samples was also generated to investigate how $\kappa(E_A^{-1})$ and $||E_A E_{A'' \setminus A}^{-1}||_2$ relate. The resulting values of $\kappa(E_A^{-1})$ and $||E_A E_{A'' \setminus A}^{-1}||_2$ are illustrated by gray markers in Fig. 1. These random samplings indeed confirm that a time sampling is limited to

¹ The norm of the alias operator can only be identically zero when oversampling is employed since at most # Λ different Fourier coefficients can be computed when $M = #\Lambda$.



Fig. 1 Relation between norm of alias operator and condition number for the case $\Lambda = \{0, 1, 2, -2, -1\} \cdot 2\pi$

the region between the three lines in Fig. 1. They also indicate that the maximal norm of the alias operator is proportional to $\kappa(E_A^{-1})$, as would be expected from Eq. (19). For the single frequency case shown in Fig. 1, it did in fact turn out that both the minimum value of $\kappa(E_A^{-1})$ and the minimum value of $\|E_A E_{A''\setminus A}^{-1}\|_2$ were obtained with M = #A uniformly distributed time samples between 0 and T. The results in Fig. 1 thus point to the fact that the standard DFT is optimal in terms of aliasing for the single frequency problem. Several other tests with different numbers of harmonics in A and different sets A'' have confirmed this statement.

different numbers of harmonics in Λ and different sets Λ'' have confirmed this statement. The relation between $\kappa(E_{\Lambda}^{-1})$ and $||E_{\Lambda}E_{\Lambda''\setminus\Lambda}^{-1}||_2$ for the case when Λ contains two fundamental frequencies and Λ'' is generated by a second degree polynomial (p = 2) is depicted in Fig. 2a. The case when Λ'' is generated by a third degree polynomial (p = 3) is further depicted in Fig. 2b. The results shown in Fig. 2 very much resemble the single frequency case in Fig. 1. A couple of differences can however be noted. To begin with, the minimal value of $||E_{\Lambda}E_{\Lambda''\setminus\Lambda}^{-1}||_2$ can be seen to increase when going from the single frequency to the multiple frequency case. For both cases, it can also be seen that the minimal value of $||E_A E_{A'' \setminus A}^{-1}||_2$ increases when the nonlinearity of the problem increases. Both these trends can be explained by the fact that the number of unresolved frequencies increases when the frequencies in Λ are no longer harmonically related and/or when the nonlinearity of the problem increases. In addition, it can be seen in Fig. 2 that the time sampling which minimizes $\kappa(E_{\Lambda}^{-1})$ (blue square marker) and the one that minimizes $||E_{\Lambda}E_{\Lambda''\setminus\Lambda}^{-1}||_2$ (yellow circle marker) are not the same in the multiple frequency case. The difference between these two samplings can however be seen to be very small. One possible explanation of this difference could be that the OPT algorithm has not converged. Alternatively, there may be a trade off between $\kappa(E_{\Lambda}^{-1})$ and $||E_A E_{A'' \setminus A}^{-1}||_2$ for the multiple frequency case. The trend of the randomly generated samplings shown in Fig. 2 does however suggest that this trade-off is very small. Overall, Figs. 1 and 2 suggest that $\kappa(E_{\Lambda}^{-1})$ is a reasonable metric for selecting a time

Overall, Figs. 1 and 2 suggest that $\kappa(E_A^{-1})$ is a reasonable metric for selecting a time sampling when M = #A time samples are used in the single and multiple frequency case. These figures also show that it is important that the optimization algorithm used to minimize $\kappa(E_A^{-1})$ is able to come close to the global optimum, since the time sampling will otherwise not minimize the norm of the alias operator.

3.1.2 Duffing Oscillator

The Duffing oscillator is a model for a damped and driven oscillator with a nonlinear stiffening spring. The equation governing the displacement of the weight in the presence of two driving



Fig. 2 Relation between norm of alias operator and condition number for the case $\Lambda = \{0, 1, \sqrt{2}, -\sqrt{2}, -1\} \cdot 2\pi$

forces may be written as

$$\ddot{x} + 2\zeta \dot{x} + x + x^3 = F_1 \sin(\omega_1 t) + F_2 \sin(\omega_2 t)$$
(27)

Here, ζ and F_i respectively denote the damping coefficient and the amplitude of the *i*th driving force. Furthermore, a dot represents differentiation with respect to time. Equation (27) may be rewritten as a first order ordinary differential equation by introducing the velocity of the weight, $y = \dot{x}$, as an additional degree of freedom. This yields the following system of equations

$$\frac{dq}{dt} + f(q,t) = 0 \tag{28}$$

where

$$q = \begin{bmatrix} x \\ y \end{bmatrix}, \quad f(q,t) = \begin{bmatrix} -y \\ 2\zeta y + x + x^3 - F_1 \sin(\omega_1 t) - F_2 \sin(\omega_2 t) \end{bmatrix}$$
(29)

Equation (28) is discretized in time using the frequency domain harmonic balance method presented in Eq. (6). The resulting nonlinear system of equations have been implemented in the Python programming languange and solved using the fsolve routine available in SciPy [20]. The exact Jacobian is also provided to fsolve to avoid problems associated with finite-difference approximations.

The implementation of the frequency domain harmonic balance solver for the Duffing oscillator has been validated against the publically available MATLAB tool NLvib [24]. To begin with, only a single driving force is considered ($F_2 = 0$). The presence of a single driving force admits the standard DFT to be used in the harmonic balance method. Figure 3a shows a comparison between the Python solver and NLvib for $\zeta = 0.1$, $F_1 = 1.25$ and K = 5. Both solutions were obtained by starting with a driving frequency close to 0 and then gradually increase the frequency. Each new frequency computation is then started from the previous solution. Figure 3a shows that the Python solver and NLvib yield identical results for this case. Note, however, that the hysteresis branch could only be computed with NLvib since no arc length continuation method was implemented in the Python solver. It should also be noted that both solutions were computed with M'' = 21 time instances to ensure that Orszag's rule (Eq. (26)) was satisfied for the cubic nonlinearity present in the governing equations.

Figure 3b shows results obtained with the Python solver for a multiple frequency problem. In this case, two incommensurable driving frequencies with a ratio of $\omega_2/\omega_1 = \sqrt{2}$ are used.



Fig. 3 Validation of frequency domain harmonic balance implementation of Duffing's oscillator against NLvib v1.0 [24]. Comparisons are made in terms of displacement amplitude $(A(\omega) = \|\hat{\mathbf{x}}_{\omega}\|_2)$ for different frequencies

The damping factor and amplitudes of the driving forces are further set to $\zeta = 0.1$ and $F_1 = F_2 = 0.1$ respectively. No higher harmonics of the driving frequencies are considered in the computation. The presence of a cubic nonlinearity in Eq. (29) will however generate a large number of additional frequencies. To avoid aliasing, the DFT defined by Eq. (23) is used. This requires the use of M'' = #A'' = 25 sampling points, which were selected to minimize the condition number of the DFT matrix. Unfortunately, NLvib does not support the APFT. A direct comparison between the two solvers was therefore not possible in the multiple frequency case. Instead, two NLvib computations were run at each of the two driving frequencies and with M'' = 5 time samples to obtain alias free solutions. These results are provided for reference in Fig. 3b. This figure clearly shows that the APFT solution and the reference single-frequency solution deviate. To ensure that this deviation is due only to the nonlinear coupling resolved by the APFT, the Python solver was rerun with $F_2 = 0$ and all other settings intact. The results from this computation can be seen to agree perfectly with the corresponding NLvib solution.

As shown previously, both $||E_{\Lambda}E_{\Lambda''\setminus\Lambda}^{-1}||_2$ and $\kappa(E_{\Lambda}^{-1})$ put a bound on the amount of aliasing produced by a DFT. These metrics will now be compared to the actual amount of aliasing obtained with the aforementioned frequency domain harmonic balance solver for the Duffing oscillator. The comparison has been performed for the two sets of frequencies that were previously used for the validation. For each of these sets, 90 random time samplings were generated by drawing M = #A samples from a uniform distribution and then checking the condition number of the corresponding DFT matrix. If the condition number was below 10, the sampling was saved. Otherwise, another sampling was generated until 90 samplings had been obtained in total.

For both the single and multiple frequency cases the corresponding alias free discrete Fourier transforms that were used for the validation were first used to compute a set of reference solutions using the same frequency stepping procedure as described previously. These reference solutions were then used as initial conditions for the simulations that employed the random samplings. For each frequency, several solutions in which the random samplings were shifted in time were computed. These time shifts were employed to account for the fact that the phase of the solution can affect the amount of aliasing obtained. Once all frequency and time shift combinations had been computed for a given random time sampling,



(a) Alias error plotted against $\kappa(E_{\Lambda}^{-1})$ (b) Alias error plotted against $||E_{\Lambda}E_{\Lambda''\setminus\Lambda}^{-1}||_2$

Fig. 4 Numerical alias error for the Duffing oscillator plotted against $\kappa(E_{\Lambda}^{-1})$ and $||E_{\Lambda}E_{\Lambda''\setminus\Lambda}^{-1}||_2$. $\Lambda = \{0, 1, 2, 3, 4, 5, -5, -4, -3, -2, -1\} \cdot \omega$, $F_1 = 0.25$, $\zeta = 0.1$, $\omega_1 = \omega$, M = 11, $\omega \in (0, 1.3)$

the following numerical alias error was calculated

Alias Error =
$$\max_{\omega} \left\{ \max_{\Delta t} \left\{ \frac{||\hat{\mathbf{x}}_{A,APFT} - \hat{\mathbf{x}}_{A,OS}||_2}{||\hat{\mathbf{x}}_{A,OS}||_2} \right\} \right\}$$
(30)

Here, APFT and OS respectively denote the random sampling and the alias-free oversampled solution. The vector $\hat{\mathbf{x}}_{\Lambda}$ in this equation further contains all Fourier coefficients of the displacement.

The numerical alias error is plotted against $\kappa(E_A^{-1})$ with circular markers for the single frequency case in Fig. 4a. This figure shows that $\kappa(E_A^{-1})$ puts a bound on the amount of aliasing obtained in a harmonic balance simulation, which is consistent with the results presented in Fig. 1. The red diamond marker and blue square marker in Fig. 4a further denote the alias free reference solution and a solution obtained with the standard DFT. It is interesting to note that the standard DFT gives the lowest numerical alias error among all samplings with M = #A, which is consistent with the results in Fig. 1. From Fig. 4a it can also be noted that $\kappa(E_A^{-1}) = 1$ for both the oversampled reference solution as well as the standard DFT based on M = #A time samples. This result points to the fact that $\kappa(E_A^{-1})$ can not distinguish an alias free solution from an aliased one. This is on the other hand possible when the numerical alias error is plotted against $\|E_A E_{A'' \setminus A}^{-1}\|_2$, as shown in Fig. 4b. This figure also shows that $\|E_A E_{A'' \setminus A}^{-1}\|_2$ puts a bound on the numerical alias error, which is consistent with Eq. (16).

In Fig. 5a and b the numerical alias errors obtained for the multiple frequency case have been plotted against $\kappa(E_A^{-1})$ and $||E_A E_{A''\setminus A}^{-1}||_2$ respectively. These figures very much resemble their single frequency counterparts in the sense that both $\kappa(E_A^{-1})$ and $||E_A E_{A''\setminus A}^{-1}||_2$ can be seen to bound the numerical alias error. The square and triangle markers in Fig. 5 correspond to two samplings which give the lowest value of $\kappa(E_A^{-1})$ and $||E_A E_{A''\setminus A}^{-1}||_2$ respectively. These two results can be seen to be very close to each other. This once again indicates that the trade off between $\kappa(E_A^{-1})$ and $||E_A E_{A''\setminus A}^{-1}||_2$ is very small in the multiple frequency case, and thus that it is sufficient to optimize for $\kappa(E_A^{-1})$ in order to obtain a DFT that minimizes aliasing. Note however that this statement only holds true when M = #A sampling points are employed since the minimal value of $\kappa(E_A^{-1})$ can be the same with or without oversampling.



(a) Alias error plotted against $\kappa(E_{\Lambda}^{-1})$ (b) Alias error plotted against $||E_{\Lambda}E_{\Lambda''\setminus\Lambda}^{-1}||_2$

Fig. 5 Numerical alias error for the Duffing oscillator plotted against $\kappa(E_A^{-1})$ and $||E_A E_{A'' \setminus A}^{-1}||_2$. $A = \{0, 1, \sqrt{2}, -\sqrt{2}, -1\} \cdot \omega$, $F_1 = F_2 = 0.01$, $\zeta = 0.1$, $\omega_1 = \omega$, $\omega_2 = \sqrt{2}\omega$, M = 5, $\omega \in (0, 2)$

3.2 Aliasing for $M' > #\Lambda$

Two strategies for employing oversampling with the APFT were presented in Sect. 2.3. In the first strategy, $E_A(\mathbf{t})$ is constructed by explicitly accounting for some of the unresolved frequencies (Eq. (25)), and in the second strategy, $E_A(\mathbf{t})$ is computed as the Moore–Penrose inverse of $E_A^{-1}(\mathbf{t})$. In this section, it will be investigated whether the additional knowledge about A'' built into the first definition of $E_A(\mathbf{t})$ is beneficial from an aliasing perspective. This will be done by investigating how the alias error obtained when the Duffing oscillator is simulated with the harmonic balance method varies with M' for both a single and a multiple frequency case. In relation to this, it is also shown how the numerical alias error relates to the norm of the alias operator and the condition number of the DFT matrix by plotting their variation with M' as well.

3.2.1 Duffing Oscillator

The first case that has been used to investigate the impact of oversampling with the APFT is the Duffing oscillator with a single driving force. The alias free reference solution for this case was computed with a uniform sampling that satisfies Orszag's rule (Eq. (26)) for cubic nonlinearities. The two definitions of $E_A(\mathbf{t})$ presented in Sect. 2.3 were evaluated using two different sets of time samples. The first set consists of M' uniformly distributed points between 0 and the reciprocal of the lowest frequency. The second set consists of M' non-uniformly distributed points created by perturbing each point in the first sampling set by a random value drawn from the uniform distribution U(-0.15T/M', 0.15T/M'). When $E_A(\mathbf{t})$ was computed based on Eq. (25), Λ' was selected to contain the M' lowest frequencies in Λ'' .

Once all samplings had been generated for all values of M', the numerical alias error was calculated according to the procedure outlined in Sect. 3.1.2. The results from this calculation are presented in Fig. 6a. This figure shows that the two definitions of $E_A(\mathbf{t})$ give equivalent results for the uniform sampling. This is because both these discrete Fourier transforms are equivalent to the standard DFT in this case. This fact also explains why both definitions of $E_A(\mathbf{t})$ eliminate aliasing once Orszag's rule is satisfied. From Fig. 6a it can also be seen that aliasing is not eliminated for any value of M' when a non-uniform sampling is used and $E_A(\mathbf{t})$



Fig. 6 Variation of numerical alias error, $||E_{\Lambda}E_{\Lambda''\setminus\Lambda}^{-1}||_2$ and κ with number of sampling points. $\Lambda = \{0, 1, 2, 3, 4, 5, -5, -4, -3, -2, -1\} \cdot \omega$, $F_1 = 1.0, \zeta = 0.1, \omega_1 = \omega, \omega \in (0, 2.2)$

is selected to be the Moore–Penrose inverse of $E_{\Lambda}^{-1}(\mathbf{t})$. This is on the other hand possible when all the information about Λ'' is built into the definition of $E_{\Lambda}(\mathbf{t})$ for the non-uniform sampling.

The variation of $||E_A E_{A''\setminus A}^{-1}||_2$ with M' for the single frequency case is further presented in Fig. 6b. This figure shows that the norm of the alias operator follows the same trend as the numerical alias error when the number of sampling points increase. In particular, it can be seen that both $||E_A E_{A''\setminus A}^{-1}||_2$ and the numerical alias error becomes zero for the same number of sampling points. This highlights the fact that $||E_A E_{A''\setminus A}^{-1}||_2$ is able to identify an alias free sampling. As noted previously, this is on the other hand not possible if only the condition number is considered. To demonstrate this, $\kappa(E_A^{-1})$ is plotted against M' in Fig. 6c.

In general, Fig. 6 indicates that both the definition of $E_A(\mathbf{t})$ as well as the choice of time sampling can have a substantial impact on the amount of aliasing obtained in a harmonic balance simulation. This raises the question regarding which combination is best. Based only on the results shown in Fig. 6, it appears that it is beneficial to define $E_A(\mathbf{t})$ based on Eq. (25). It also appears that the uniform sampling, which gives the lowest value of $\kappa(E_A^{-1})$, is a better choice than a non-uniform sampling. This indicates that the best choice in terms of aliasing is to construct $E_A(\mathbf{t})$ based on Eq. (25), and then select a time sampling which minimizes $\kappa(E_A^{-1})$. Here, care must however be taken in the definition of the condition number. In this work, it is suggested to compute the condition number based on the square matrix $E_{A'}^{-1}(\mathbf{t})$, and not based on Eq. (17) in which some columns/rows of $E_{A'}^{-1}/E_{A'}(\mathbf{t})$ are neglected. The first reason for this is that, in practice, $E_A(\mathbf{t})$ is computed from a numerical inverse of $E_{A'}^{-1}(\mathbf{t})$ (see Eq. (25)). For this preliminary step to be well-posed it is thus necessary that $\kappa(E_{A'}^{-1})$ is small. The second reason is that a minimal value of the condition number has been seen to make the APFT robust against aliasing for the non-oversampled cases, independently on the choice of Λ'' . In the suggested approach, $E_{\Lambda}(\mathbf{t})$ is thus computed by first computing a square matrix $E_{\Lambda'}(\mathbf{t})$ that gives as little aliasing as possible, and then explicitly eliminate aliasing with respect to some of the unresolved Fourier coefficients, i.e., those that correspond to frequencies in $\Lambda' \setminus \Lambda$.

In order to investigate the usefulness of the suggested approach for computing $E_A(\mathbf{t})$, the Duffing oscillator with two incommensurable driving frequencies was once again considered. The alias free reference solution for this case was computed using the DFT defined in Eq. (23). The time sampling for this case was selected by minimizing $\kappa (E_{A''}^{-1})$ using the OPT algorithm [13]. For each M', two different discrete Fourier transforms were then constructed. The first one was calculated as the Moore–Penrose inverse of $E_A^{-1}(\mathbf{t})$, using a time sampling which minimized $\kappa (E_A^{-1})$, and the second one using Eq. (25) and a time sampling that minimized $\kappa (E_{A'}^{-1})$. These minima were also obtained with the OPT algorithm.

The variation of the numerical alias error with M' for the multiple frequency case is presented in Fig. 7a. From this figure, it can be seen that for this case the suggested approach effectively eliminates aliasing once $M' \ge 17$. Calculating $E_A(\mathbf{t})$ as the Moore–Penrose inverse of $E_A^{-1}(\mathbf{t})$ does on the other hand not eliminate aliasing for any M'. These results once again indicate that the definition of $E_A(\mathbf{t})$ can have a significant effect on aliasing. The fact that aliasing also is effectively eliminated despite the fact that not all the unresolved frequencies are accounted for in the construction of $E_A(\mathbf{t})$ is also interesting, since this suggests that there could exist a generalization of Orszag's rule for multiple frequency problems. That is, it is not necessary to construct $E_{A'}(\mathbf{t})$ in Eq. (25) such that all the unresolved sinusoids can be distinguished, only such that the unresolved ones alias onto those with frequencies in $\Lambda' \setminus \Lambda$.

The variation of $||E_A E_{A''\setminus A}^{-1}||_2$ with M' for the multiple frequency case is shown in Fig. 7b. This figure shows that the trend of $||E_A E_{A''\setminus A}^{-1}||_2$ once again follows that of the numerical alias error as the number of sampling points increase.

A comment on the point corresponding to M' = 15 in Fig. 7 should also be made. In this point, it can be seen that the numerical alias error is the largest for the case when $E_A(\mathbf{t})$ is computed based on Eq. (25). At first, this might seem counter intuitive since more information about the unresolved sinusoids is included in the DFT compared to e.g. M' = 5. A possible explanation for this can be found in Fig. 7c. This figure shows that the largest value of $\kappa(E_{A'}^{-1})$ is obtained for M' = 15. As such, it can be expected that $E_{A'}(\mathbf{t})$ is the least robust against aliasing for M' = 15. Thus, if not all the unresolved frequencies alias onto the frequencies in $A' \setminus A$, which can not be guaranteed, this sampling may very well generate more aliasing than another one that uses less sampling points and have a lower condition number. It is however not known at this point whether the relatively large value of $\kappa(E_{A'}^{-1})$ for M' = 15 is due to the fact that the OPT algorithm was unable to find the global optima or not.

4 Conclusions

It is well known that aliasing may occur when the harmonic balance method [16, 33, 34] is applied to nonlinear problems. In the present paper, aliasing is studied in detail for the case when the harmonic balance method is used to solve first order ordinary differential equations which are nonlinear functions of the solution variables. It is shown that aliasing occurs as a result of the fact that the harmonic balance method approximates the Fourier coefficients of a nonlinear function by a discrete Fourier transform (DFT). Although this is already well



Fig. 7 Variation of numerical alias error, $||E_A E_{A'' \setminus A}^{-1}||_2$ and κ with number of sampling points. $\Lambda = \{0, 1, \sqrt{2}, -\sqrt{2}, -1\} \cdot \omega, F_1 = F_2 = 0.04, \zeta = 0.1, \omega_1 = \omega, \omega_2 = \sqrt{2}\omega, \omega \in (0, 2)$

known, the aliasing of each unresolved frequency onto the resolved ones is here given a precise meaning by introducing a new operator that governs aliasing. This operator, referred to as the *alias operator*, is derived for the general case when the almost periodic Fourier transform (APFT) [26] is used in the harmonic balance method. As such, the alias operator can be used to study aliasing in both single frequency as well as multiple frequency harmonic balance computations.

Using the norm of the newly introduced alias operator as a metric, the best sampling strategy for single and multiple frequency harmonic balance computations is investigated. It is found that for single frequency problems, the widely adopted uniform sampling approach is optimal. This sampling is also known to minimize the condition number of the DFT matrix [26]. For the general case with multiple fundamental frequencies, the time sampling that minimizes the condition number of the DFT matrix seems to be close, but not identical, to the time sampling that minimizes the norm of the alias operator. A low condition number and norm of the alias operator is also shown to give a smaller alias error in numerical simulations of the DFT matrix is a reasonable metric for selecting a time sampling, as has previously been done by e.g. [13, 21, 25–27, 35, 37].

Another important finding from the present work is that the alias error is only bounded by, but not strongly correlated with, the condition number. This does in turn imply that the time sampling must give a condition number which is close to the global optimum if the alias error is to be minimized. Therefore, in cases when an optimization algorithm is used to find the time sampling, it is important that it converges well. In the present work, it is found that the OPT algorithm introduced in [13] performs very well. Employing the APFT in combination with oversampling is also considered in this work. In this case, the DFT matrix is defined as a left inverse of the inverse DFT matrix. Previous studies have used the Moore–Penrose inverse for this purpose [5, 6, 21]. In the present work, however, it is shown that this left inverse in general can not eliminate aliasing in the multiple frequency case. Based on this finding, a new left inverse that can eliminate aliasing completely is introduced. This new left inverse is also shown to perform better in general than the Moore–Penrose inverse on the Duffing Oscillator. Based on these observations, it is suggested that the newly introduced left inverse is used in favor of the Moore–Penrose inverse when oversampling is employed in multiple frequency harmonic balance computations based on the APFT.

The above findings for the case of oversampling apply both to the abstract metric defined by the norm of the aliasing operator and to the measured alias error in the numerical experiments. In contrast, the conditon number could not be used as a criterion to distinguish samplings that were optimal with regard to aliasing. It is therefore concluded that the norm of the alias operator contains valuable information about the quality of a given sampling strategy.

It should finally be noted that both the alias error and the harmonic truncation error encountered in the harmonic balance method can be reduced by increasing the number of frequencies in the computation. In many practical applications, this is unfortunately not feasible, and good control over the alias error is thus important to achieve accurate results. In cases when more frequencies can be afforded, however, one must be mindful of the computational complexity of the APFT, which scales as $\mathcal{O}(M^2)$. To overcome this limitation, the Non Uniform Fast Fourier Transform (NUFFT) of type three could be used instead [4, 30]. The theory developed in this paper would however need to be adapted for the NUFFT.

Acknowledgements The first author would like to express his gratitude towards the German Aerospace Center (DLR) in Cologne for hosting him during a 6 month research visity in the spring of 2019.

Funding Open access funding provided by Chalmers University of Technology. The first author has recieved funding from the Swedish Energy Agency, Grant Agreement No. 2017-00116.

Availability of data and material The datasets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

Appendix A Moore–Penrose Inverse with Redefined Inner Product

The goal of this section is to prove that the novel Fourier transformation matrix defined in Eq. (25) is in fact a generalization of the standard pseudoinverse in that it is the Moore–Penrose inverse with respect to a generalized inner product on the space of sampled functions.

Let $\Lambda'' = \{0, \omega_1, \omega_2, \dots, \omega_{K''}, \omega_{-K''}, \dots, \omega_{-1}\}$ be a spectrum that contains a strictly smaller spectrum Λ , the latter representing the *resolved* frequencies. Assume that $\mathbf{t} = [t_0, \dots, t_{M''-1}]^T$ is a sampling point distribution such that $M'' = \#\Lambda'' = 2K'' + 1$ and such that $E_{\Lambda''}^{-1}(\mathbf{t})$ is invertible. Consider the space of real-valued discrete functions on \mathbf{t} . Such a function f can be viewed as an element in $\mathbb{R}^{M''}$, where each component f_i corresponds to the value of f at the sampling point t_i . Instead of the standard inner product on $\mathbb{R}^{M''}$ we define an inner product on the space of sampled functions by pulling back the standard inner product of $\mathbb{C}^{M''} \simeq \mathbb{R}^{2M''}$ using the Fourier transformation $E_{\Lambda''}(\mathbf{t})$, i.e., we set

$$\langle f, g \rangle := \langle E_{\Lambda''}(\mathbf{t}) f, E_{\Lambda''}(\mathbf{t}) g \rangle, \tag{31}$$

for two discrete functions f and g. That is, $E_{A''}(\mathbf{t})$ becomes an isometry with respect to these inner products, i.e., it preserves lengths and angles. An immediate consequence of (31) is that

$$(E_{A''}(\mathbf{t}))^* = E_{A''}^{-1}(\mathbf{t}).$$

Here, the symbol * refers to the adjoint with respect to the corresponding inner products. It is defined as

$$\langle (E_{A''}(\mathbf{t}))^* \hat{f}, g \rangle = \langle \hat{f}, E_{A''}(\mathbf{t})g \rangle,$$

for all Fourier coefficient vectors \hat{f} and discrete functions g. Moreover, note that we can write the oversampled reconstruction operator $E_A^{-1}(\mathbf{t})$ as the composition of the reconstruction operator $E_{A''}^{-1}(\mathbf{t})$ and the natural embedding operator $\iota_{A,A''} : \mathbb{R}^{2M} \to \mathbb{R}^{2M''}$, which, loosely speaking, just fills up the Fourier coefficient vector with zeros,

$$(\iota_{\Lambda,\Lambda''}\hat{f})_{\omega} = \begin{cases} \hat{f}_{\omega}, & \text{if } \omega \in \Lambda, \\ 0, & \text{otherwise.} \end{cases}$$

Hence,

$$E_{\Lambda}^{-1}(\mathbf{t}) = E_{\Lambda''}^{-1}(\mathbf{t}) \iota_{\Lambda,\Lambda''},$$

The adjoint of $\iota_{\Lambda,\Lambda''}$ is the projection onto the first *K* harmonics, denoted by $\pi_{\Lambda'',\Lambda}$, i.e., the operator that "forgets" the Fourier coefficients with frequencies in $\Lambda'' \setminus \Lambda$. The Moore–Penrose inverse of $E_{\Lambda}^{-1}(\mathbf{t})$ is thus given by

$$(E_{\Lambda}^{-1}(\mathbf{t}))^{+} = \left((E_{\Lambda}^{-1}(\mathbf{t}))^{*} E_{\Lambda}^{-1}(\mathbf{t}) \right)^{-1} (E_{\Lambda}^{-1}(\mathbf{t}))^{*} = \left(\pi_{\Lambda'',\Lambda} (E_{\Lambda''}^{-1}(\mathbf{t}))^{*} E_{\Lambda''}^{-1}(\mathbf{t}) \iota_{\Lambda,\Lambda''} \right)^{-1} \pi_{\Lambda'',\Lambda} (E_{\Lambda''}^{-1}(\mathbf{t}))^{*} = \left(\pi_{\Lambda'',\Lambda} \iota_{\Lambda,\Lambda''} \right)^{-1} \pi_{\Lambda'',\Lambda} E_{\Lambda''}(\mathbf{t}) = \pi_{\Lambda'',\Lambda} E_{\Lambda''}(\mathbf{t}),$$
(32)

where we have used that $\pi_{\Lambda'',\Lambda}\iota_{\Lambda,\Lambda''}$ is the identity. Moreover, the kernel of the Moore–Penrose inverse is given by

$$\ker\left((E_{\Lambda}^{-1}(\mathbf{t}))^{+}\right) = \ker\left(\pi_{\Lambda'',\Lambda}E_{\Lambda''}(\mathbf{t})\right)$$

Deringer

$$E_{A^{\prime\prime}}^{-1}(\mathbf{t}) \ker \pi_{A^{\prime\prime},A}.$$
(33)

The kernel thus consists of all discrete functions f that are reconstructed from linear combinations of harmonics with frequencies in $\Lambda'' \setminus \Lambda$, i.e. the *unresolved* frequencies.

=

Appendix B Choice of Numerical Optimization Algorithm

Two different optimization procedures were tested in this work. The first one was the OPT algorithm proposed by Guedeney et al. [13]. This algorithm is divided into two steps. In the first step, the objective function is evaluated for a very large number of equidistant time-samplings. In the second step, the best sampling is used to initialize an optimization based on the L-BFGS-B algorithm. The OPT algorithm thus benefits from both the strengths of a global search and a local optimization. The OPT algorithm available in SciPy [20]. The second optimization procedure that was employed was the Differential Evolution (DE) algorithm. DE is an evolutionary algorithm that uses mutation, recombination and tournament selection to evolve a large population (set of different time samplings) over several generations. In particular, DE benefits from a very efficient mutation step based on computing parameter differences between two (or more) individuals. In this work, the DE implementation available in SciPy was utilized.

In all optimizations the optimizer was allowed to modify all but the first time instance, which was set to zero. A constraint was also added in all optimizations to ensure that the time samples remained ordered, i.e. that $t_i > t_j$ if i > j. This was done to ensure that several optima, corresponding to different permutations of the time sampling, could not be found by the optimizer. When the norm of the alias operator, or the maximum value of the bounding coefficient α (Λ , Λ'' , **t**), were optimized, an additional constraint was imposed which ensured that the condition number of $E_A^{-1}(\mathbf{t})$ did not exceed the value of 10^3 . This guarantees that the inversion of $E_A^{-1}(\mathbf{t})$ remains a well-posed problem for all time samplings that came out of the optimization. All constraints were furthermore implemented with Lagrange multipliers.

Comparisons between the OPT and DE algorithms showed that they often give very similar results. The OPT algorithm was however found to be substantially faster than the DE algorithm. It is believed that one of the main reasons for this is that the time-samplings obtained from the mutation step in DE are not necessarily ordered. This will in turn lead to a large number of mutations that generate individuals that violate the ordering constraint, and thus not improve the population. If the time samplings on the other hand are not required to be ordered, the number of samplings that give an equivalent objective function will be M!. For large M, this implies that each individual can have the same objective function value, but different parameter vectors. In terms of DE, which is based on evolving the population using differences between parameter vectors, this should translate into bad performance since all individuals are far from each other in the design space. Given these difficulties with the DE algorithm for the particular problems considered in this work, it was decided to employ the OPT algorithm.

References

- 1. Besicovitch, A.S.: Almost Periodic Functions. Cambridge University Press, Cambridge (1932)
- Chua, L., Ushida, A.: Algorithms for computing almost periodic steady-state response of nonlinear systems to multiple input frequencies. IEEE Trans. Circuits Syst. 28(10), 953–971 (1981). https://doi.org/ 10.1109/TCS.1981.1084921
- Djeddi, R., Ekici, K.: Resolution of Gibbs phenomenon using a modified pseudo-spectral operator in harmonic balance CFD solvers. Int. J. Comput. Fluid Dyn. 30(7–10), 495–515 (2016). https://doi.org/10. 1080/10618562.2016.1242726
- Dutt, A., Rokhlin, V.: Fast Fourier transforms for nonequispaced data. SIAM J. Sci. Comput. 14(6), 1368–1393 (1993). https://doi.org/10.1137/0914081
- Ekici, K., Hall, K.C.: Nonlinear analysis of unsteady flows in multistage turbomachines using harmonic balance. AIAA J. 45(5), 1047–1057 (2007). https://doi.org/10.2514/1.22888
- Ekici, K., Hall, K.C.: Nonlinear frequency-domain analysis of unsteady flows in turbomachinery with multiple excitation frequencies. AIAA J. 46(8), 1912–1920 (2008). https://doi.org/10.2514/1.26006
- Frey, C., Ashcroft, G., Kersken, H.P.: Simulations of unsteady blade row interactions using linear and non-linear frequency domain methods. In: ASME Turbo Expo 2015: Turbine Technical Conference and Exposition, GT2015-43453 (2015). https://doi.org/10.1115/GT2015-43453
- Frey, C., Kersken, H.P., Ashcroft, G., Voigt, C.: A harmonic balance technique for multistage turbomachinery applications. In: ASME Turbo Expo 2014: Turbine Technical Conference and Exposition, GT2014-25230 (2014). https://doi.org/10.1115/GT2014-25230
- 9. Gelb, A., Gottlieb, S.: The resolution of the Gibbs phenomenon for Fourier spectral methods. In: Jerri, A. (ed.) Advances in the Gibbs Phenomenon. Sampling Publishing, Potsdam (2007)
- Gilmore, R.J., Steer, M.B.: Nonlinear circuit analysis using the method of harmonic balance—a review of the art. Part I. Introductory concepts. Int. J. Microw. Millim. Wave Comput. Aided Eng. 1(1), 22–37 (1991). https://doi.org/10.1002/mmce.4570010104
- Gilmore, R.J., Steer, M.B.: Nonlinear circuit analysis using the method of harmonic balance—a review of the art. Part II. Advanced concepts. Int. J. Microw. Millim. Wave Comput. Aided Eng. 1(2), 159–180 (1991). https://doi.org/10.1002/mmce.4570010205
- Gottlieb, D., Shu, C.W.: On the Gibbs phenomenon and its resolution. SIAM Rev. 39(4), 644–668 (1997). https://doi.org/10.1137/S0036144596301390
- Guédeney, T., Gomar, A., Gallard, F., Sicot, F., Dufour, G., Pugit, G.: Non-uniform time sampling for multiple-frequency harmonic balance computations. J. Comput. Phys. 236, 317–345 (2013). https://doi. org/10.1016/j.jcp.2012.11.010
- Guskov, M., Thouverez, F.: Harmonic balance-based approach for quasi-periodic motions and stability analysis. J. Vib. Acoust. 134(3), 031003-1-031003–11 (2012). https://doi.org/10.1115/1.4005823
- Hall, K.C., Crawely, E.F.: Calculation of unsteady flows in turbomachinery using the linearized Euler equations. AIAA J. 27(6), 777–787 (1989). https://doi.org/10.2514/3.10178
- Hall, K.C., Jeffrey, P.T., Clark, W.S.: Computation of unsteady nonlinear flows in cascades using a harmonic balance technique. AIAA J. 40(5), 879–886 (2002). https://doi.org/10.2514/2.1754
- He, L.: Fourier methods for turbomachinery applications. Prog. Aerosp. Sci. 46(8), 329–341 (2010). https://doi.org/10.1016/j.paerosci.2010.04.001
- Heners, J.P., Vogt, D.M., Frey, C., Ashcroft, G.: Investigation of the impact of unsteady turbulence effects on the aeroelastic analysis of a low-pressure turbine rotor blade. J. Turbomach. (2019). https://doi.org/ 10.1115/1.4043950
- Huang, H., Ekici, K.: Stabilization of high-dimensional harmonic balance solvers using time spectral viscosity. AIAA J. 52(8), 1784–1794 (2014). https://doi.org/10.2514/1.J052698
- Jones, E., Oliphant, T., Peterson, P., et al.: SciPy: open source scientific tools for Python (2001–). http:// www.scipy.org/ [Version 1.2.2]
- Junge, L., Ashcroft, G., Kersken, H.P., Frey, C.: On the development of harmonic balance methods for multiple fundamental frequencies. In: ASME Turbo Expo 2018: Turbomachinery Technical Conference and Exposition, GT2018-75495 (2018). https://doi.org/10.1115/GT2018-75495
- Junge, L., Frey, C., Ashcroft, G., Kügeler, E.: A new harmonic balance approach using multidimensional time. J. Eng. Gas Turbines Power 143(8), 081007-1–08100711 (2021). https://doi.org/10.1115/1.4049698
- Kotel'nikov, V.A.: On the carrying capacity of the "ether" and wire in telecommunications. In: Material for the First All-Union Conference on Questions of Communications (Russian), Izd., Red. Upr. Svyazi RKKA
- Krack, M., Gross, J.: Harmonic Balance for Nonlinear Vibration Problems. Mathematical Engineering. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-14023-6

- Kundert, K.S., Sorkin, G.B., Sangiovanni-Vincentelli, A.: Applying harmonic balance to almost periodic circuits. IEEE Trans. Microw. Theory Tech. 36(2), 366–378 (1988). https://doi.org/10.1109/22.3525
- Kundert, K.S., White, J.K., Sangiovanni-Vincentelli, A.: Steady-State Methods for Simulating Analog and Microwave Circuits, The Kluwer International Series in Engineering and Computer Science, vol. 94. Kluwer Academic Publishers, Norwel (1990). https://doi.org/10.1007/978-1-4757-2081-5
- Kunisch, J., Wolff, I.: Determination of sampling points for nearly DFT-equivalent almost-periodic Fourier transforms. In: 1993 23rd European Microwave Conference (1993). https://doi.org/10.1109/EUMA.1993. 336680
- LaBryer, A., Attar, P.J.: High dimensional harmonic balance dealiasing techniques for a Duffing oscillator. J. Sound Vib. 324(3–5), 1016–1038 (2009). https://doi.org/10.1016/j.jsv.2009.03.005
- LaBryer, A., Attar, P.J.: A harmonic balance approach for large scale problems in nonlinear structural dynamics. Comput. Struct. 88(17–18), 1002–1014 (2010). https://doi.org/10.1016/j.compstruc.2010.06. 003
- Lee, J.Y., Greengard, L.: The type 3 nonuniform FFT and its applications. J. Comput. Phys. 206(1), 1–5 (2005). https://doi.org/10.1016/j.jcp.2004.12.004
- Liu, L., Thomas, J.P., Dowell, E.H., Attar, P., Hall, K.C.: A comparison of classical and high dimensional harmonic balance approaches for a Duffing oscillator. J. Comput. Phys. 215(1), 298–320 (2006). https:// doi.org/10.1016/j.jcp.2005.10.026
- McMullen, M., Jameson, A., Alonso, J.: Demonstration of nonlinear frequency domain methods. AIAA J. 44(7), 1428–1435 (2006). https://doi.org/10.2514/1.15127
- McMullen, M.S., Jameson, A.: The computational efficiency of non-linear frequency domain methods. J. Comput. Phys. 212(2), 637–661 (2006). https://doi.org/10.1016/j.jcp.2005.07.021
- Nakhla, M., Valch, J.: A piecewise harmonic balance technique for determination of periodic response of nonlinear systems. IEEE Trans. Circuits Syst. 23(2), 85–91 (1976). https://doi.org/10.1109/TCS.1976. 1084181
- Nimmagadda, S., Economon, T.D., Alonso, J.J., da Silva, C.R.I.: Robust uniform time sampling approach for the harmonic balance method. In: 46th AIAA Fluid Dynamics Conference, AIAA Paper 2016-3966 (2016). https://doi.org/10.2514/6.2016-3966
- 36. Orszag, S.A.: On the elimination of aliasing in finite-difference schemes by filtering high-wavenumber components. J. Atmos. Sci. 28(6), 1074 (1971). https://doi.org/10.1175/1520-0469(1971)028 < 1074: OTEOAI > 2.0.CO;2
- Rodrigues, P.J.C.: An orthogonal almost-periodic Fourier transform for use in nonlinear circuit simulation. IEEE Microw. Guided Wave Lett. 4(3), 74–76 (1994). https://doi.org/10.1109/75.275585
- 38. Shannon, C.E.: Communication in the presence of noise. Proc. IRE 37, 10-21 (1949)
- Steer, M.B., Chang, C.R., Rhyne, G.W.: Computer-aided analysis of nonlinear microwave circuits using frequency-domain nonlinear analysis techniques: the state of the art. Int. J. Microw. Millim. Wave Comput. Aided Eng. 1(2), 181–200 (1991). https://doi.org/10.1002/mmce.4570010206
- Whittaker, E.T.: On the functions which are represented by the expansion of interpolating theory. Proc. R. Soc. Edinb. 35, 181–194 (1915)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.