

Efficient harmonic resolvent analysis via time-stepping

Ali Farghadan¹, Junoh Jung¹, Rutvij Bhagwat¹, Aaron Towne^{1*}

¹Department of Mechanical Engineering, University of Michigan, Ann Arbor, 48109, MI, USA.

*Corresponding author(s). E-mail(s): towne@umich.edu;

Contributing authors: alii@umich.edu; junohj@umich.edu; rutvij@umich.edu;

Abstract

We present an extension of the RSVD- Δt algorithm initially developed for resolvent analysis of statistically stationary flows to handle harmonic resolvent analysis of time-periodic flows. The harmonic resolvent operator, as proposed by Padovan et al. [1], characterizes the linearized dynamics of time-periodic flows in the frequency domain, and its singular value decomposition reveals forcing and response modes with optimal energetic gain. However, computing harmonic resolvent modes poses challenges due to **(i)** the coupling of all N_ω retained frequencies into a single harmonic resolvent operator and **(ii)** the singularity or near-singularity of the operator, making harmonic resolvent analysis considerably more computationally expensive than a standard resolvent analysis. To overcome these challenges, the RSVD- Δt algorithm leverages time stepping of the underlying time-periodic linearized Navier-Stokes operator, which is N_ω times smaller than the harmonic resolvent operator, to compute the action of the harmonic resolvent operator. We develop strategies to minimize the algorithm's CPU and memory consumption, and our results demonstrate that these costs scale linearly with the problem dimension. We validate the RSVD- Δt algorithm by computing modes for a periodically varying Ginzburg-Landau equation and demonstrate its performance using the flow over an airfoil.

Keywords: Turbulence, Harmonic resolvent analysis, Randomized singular value decomposition, Time stepping, Periodic flows

1 Introduction

Model reduction plays a critical role in the study of fluid mechanics due to the complex and high-dimensional nature of fluid flows, especially turbulent ones. In particular, modal decompositions, both data-driven and equation-based, have proven to be effective at identifying low-dimensional sets of modes associated with interpretable coherent structures that significantly influence quantities of engineering interest such as kinetic energy, heat and mass transfer, and noise emissions [2, 3]. In particular, resolvent analysis has been extensively employed to comprehend, model, and control statistically stationary turbulent flows such as wall-bounded turbulence [4], turbulent boundary layer [5], and jet flows [6]. However, real-world fluid systems often exhibit non-stationary behavior, including periodicity, rendering resolvent analysis less effective. Examples of flows exhibiting periodic behavior include internal waves in a stratified fluid [7], vortex shedding in the wake of a bluff body [8], and pulsatile blood flow [9]. Recently, an extension of resolvent analysis, known as harmonic resolvent analysis, has been developed by Padovan et al. [1]. This extension characterizes the perturbation forcing and responses around a periodic base flow, providing a more effective tool for analyzing, modeling, and controlling periodically time-varying flows. The harmonic resolvent operator can be considered a special case of the harmonic transfer function introduced by Wereley [10], wherein the operator exhibits a dominant frequency.

Investigating the harmonic resolvent operator is one way to identify the optimal spatio-temporal perturbations and their corresponding responses in fluid systems. Unlike the resolvent operator, the harmonic resolvent operator takes into consideration the periodicity of the base flow by representing it in the form of a Fourier series. This allows for characterization of the dynamic behavior of the system, particularly with respect to its spectral components and it provides a means to understand and explore the interactions between input and output modes through the linearized dynamics. Harmonic resolvent analysis determines the triadic interactions between the input frequency ω_1 , the base flow frequency $\omega_1 - \omega_2$, and the response frequency ω_2 . The number of Fourier modes of the base flow is determined by its power spectral density (PSD), which in turn determines the number of triads interacting with the forcing. Typically, periodic flows with a dominant periodicity exhibit a multimodal nature with their most energetic modes being harmonics of the base frequency. When the base flow is statistically stationary, linearization around the temporal average is sufficient, and traditional resolvent analysis arises as a special case.

Recent studies have expanded on this research area to model and control the dynamics of turbulent flows with periodic motions. Lin et al. [11] applied the harmonic resolvent framework for linear time-periodic systems with more than one dominant frequency, offering a tool to study the flow control of a plunging cylinder. In a distinct approach, Franceschini et al. [12] proposed a novel method aimed at capturing the phase dependency of small-scale turbulent phenomena in flows exhibiting a periodic limit cycle. This method builds upon classical spectral proper orthogonal decomposition (SPOD) and resolvent analyses but introduces a quasi-steady approximation that effectively separates the high-frequency turbulent component from the slow periodic oscillation. Moreover, Heidt and Colonius [13] have pioneered the theory and algorithm for cyclostationary SPOD (CS-SPOD). This innovative approach extends the conventional SPOD algorithm to effectively capture the characteristics of flows characterized by periodic motions. CS-SPOD is closely tied to harmonic resolvent analysis, akin to the relationship between SPOD and resolvent analysis [3].

Harmonic resolvent analysis offers a promising avenue for new discoveries in physics. For example, recent research conducted by Wu et al. [14] explored the response of a turbulent separation bubble to zero-net-mass-flux actuation, with a base flow characterized by periodicity. The noteworthy finding revealed an agreement between the reduction in separation bubble size identified through both harmonic resolvent optimal response and direct numerical simulation. Another contribution in this realm comes from the work of Padovan and Rowley [15], who investigated the driving mechanisms underlying vortex pairing in an incompressible forced axisymmetric jet. They showed that the optimal forcing of the subharmonic resolvent operator featured spatiotemporal structures leading to the nonlinear vortex pairing phenomenon. In general, harmonic resolvent analysis can facilitate the understanding of energy transfer mechanisms [16] and is a potentially powerful tool for secondary stability analyses involving periodic additions to the base flow, with applications such as boundary layer stability and transition [17], screeching jets [18], and the effect of mass injection on secondary instability [19].

From a computational standpoint, computing harmonic resolvent modes presents similar challenges to resolvent analysis, including the need to invert sparse matrices in Fourier space. This challenge is further exacerbated in harmonic resolvent analysis due to expanded system size, where all relevant frequencies are merged into a unified operator, in contrast to resolvent analysis, where each frequency has its own independent system. The second shared challenge involves performing the singular value decomposition (SVD). Turbulent and other complex flows characterized by periodic motions require high-resolution meshes and potentially large computational domains. This results in a sizable linearized operator, introducing formidable computational challenges when it comes to matrix inversion and performing SVD. A modified randomized SVD (RSVD) can be used to simultaneously overcome both computational challenges [1, 20, 21]. Nevertheless, the modified RSVD, referred to as “RSVD-LU” in this study, becomes impractical for tackling large-scale turbulent flows since it requires the lower-upper (LU) decomposition of the harmonic resolvent operator. Recently, RSVD- Δt , an improvement to RSVD-LU, addresses bottlenecks in resolvent analysis [22]. The approach of time-stepping, initially developed by Monokrousos et al. [23] and further optimized by Martini et al. [24], allows for computing the action of an operator on a vector without solving a linear system in Fourier space, making it a key component of RSVD- Δt .

The main objective of this study is to expand the applicability of the RSVD- Δt algorithm, initially crafted for resolvent analysis, to facilitate the computation of harmonic resolvent modes. A streaming approach is adopted for time integration, reducing memory usage while maintaining computational efficiency. Moreover, an efficient transient removal strategy is introduced, cutting the temporal length of the integration and

thereby minimizing the CPU cost of RSVD- Δt for achieving a desired accuracy. The RSVD- Δt algorithm for periodic flows maintains linear scalability in both CPU time and memory consumption. This scalability enables the efficient computation of harmonic resolvent modes for large-scale turbulent flows that were previously out of reach.

The structure of our paper unfolds as follows. We begin by explaining resolvent analysis in §2, followed by the formulation of harmonic resolvent analysis in §3, where the computation of harmonic resolvent modes using RSVD-LU is briefly described. The RSVD- Δt algorithm is introduced in §4. A comparison of the computational complexity of algorithms and suggested improvements appear in §5. Further details on performance and error analysis are presented in §6. In §7, we conduct two test cases to demonstrate the accuracy and cost efficiency of RSVD- Δt for two periodic systems, and we conclude with final remarks in §8.

2 Resolvent analysis

Before considering the time-periodic case, we briefly review standard resolvent analysis for stationary flows. The Navier-Stokes equations can be linearized around a base flow via Reynolds decomposition of the state, $\mathbf{q} = \bar{\mathbf{q}} + \mathbf{q}'$, where $\bar{\mathbf{q}}$ represents the temporal average of the flow, and \mathbf{q}' denotes the time-varying fluctuations. The dynamics of perturbations are then described by the linear time-invariant system

$$\frac{d\mathbf{q}'}{dt} = \mathbf{A}(\bar{\mathbf{q}})\mathbf{q}' + \mathbf{f}'(\bar{\mathbf{q}}, \mathbf{q}'), \quad (1)$$

where $\mathbf{A} \in \mathbb{C}^{N \times N}$ is the linearized Navier-Stokes (LNS) operator, N is the state dimension of the discretized system, and \mathbf{f}' represents an external forcing as well as nonlinear terms that are treated as an additional forcing, as suggested by McKeon and Sharma [25]. It should be noted that \mathbf{A} is time-independent because the base flow is time-invariant. By taking the Fourier transform

$$\mathcal{F}(\cdot) = \hat{(\cdot)}(\omega) = \int_{-\infty}^{+\infty} (\cdot) e^{-i\omega t} dt \quad (2)$$

in time, we obtain

$$\hat{\mathbf{q}}(\omega) = \mathbf{R}(\omega)\hat{\mathbf{f}}(\omega), \quad (3)$$

where ω represents the frequency and $\hat{(\cdot)}$ denotes the frequency-domain representation of the time-domain vector. The resolvent operator

$$\mathbf{R} = (i\omega\mathbf{I} - \mathbf{A})^{-1} \in \mathbb{C}^{N \times N} \quad (4)$$

is a transfer function that maps the input forcing to the output response. Here, $i = \sqrt{-1}$ and \mathbf{I} is the identity matrix of appropriate dimension.

To analyze the perturbation characteristics of a system, the transfer function can be examined across a range of frequencies. Typically, the most amplified responses are identified by computing the left singular vectors \mathbf{U}_R of the resolvent operator,

$$\mathbf{R} = \mathbf{U}_R \boldsymbol{\Sigma}_R \mathbf{V}_R^*, \quad (5)$$

where $(\cdot)^*$ denotes complex conjugate transpose. The optimal forcing inputs that lead to the highest amplifications comprise the right singular vectors \mathbf{V}_R , and the degree of amplifications are determined by the singular values $\boldsymbol{\Sigma}_R$.

The SVD in 5 implies that disturbance amplitudes are measured in a Euclidean norm. More generally, $\|\mathbf{x}\|_f^2 = \langle \mathbf{x}, \mathbf{x} \rangle_f = \mathbf{x}^* \mathbf{W}_f \mathbf{x}$ computes the f -norm of any vector \mathbf{x} , and input and output norms can be different, *i.e.*, $\|\cdot\|_q = \|\cdot\|_f$ is not necessary. Additionally, one can define an input matrix \mathbf{B} to restrict the forcing and an output matrix \mathbf{C} to extract the output of interest from the state. For the sake of notational brevity, we assumed the weight, input, and output matrices to be the identity, *i.e.*, $\mathbf{W} = \mathbf{B} = \mathbf{C} = \mathbf{I}$, throughout this paper. However, our algorithm, with minor adjustments, can accommodate non-identity weight, input, and output matrices by computing the modes of the modified resolvent operator

$$\tilde{\mathbf{R}} = \mathbf{W}_q^{1/2} \mathbf{C} \mathbf{R} \mathbf{B} \mathbf{W}_f^{-1/2}. \quad (6)$$

For more details, please refer to Farghadan et al. [22].

3 Harmonic resolvent analysis

In this section, we provide an overview of harmonic resolvent analysis. After reviewing its basic formulation, we discuss the singularity issue that can arise in some special flows, the truncation of the harmonic resolvent operator, and two variations of harmonic resolvent analysis, namely cross-frequency amplification and sub-harmonic resolvent. Lastly, a brief overview is provided for the state-of-the-art algorithm used to compute the harmonic resolvent modes and gains.

3.1 Formulation

Analogous to resolvent analysis, by inserting the generalized Reynolds decomposition $\mathbf{q}(t) = \bar{\mathbf{q}}(t) + \mathbf{q}'(t)$ into the Navier-Stokes equations, the linearization around a periodic base flow $\bar{\mathbf{q}}(t) = \bar{\mathbf{q}}(t+T)$ leads to the linear time-periodic system

$$\frac{d\mathbf{q}'}{dt} = \mathbf{A}_p \mathbf{q}' + \mathbf{f}, \quad (7)$$

where $\mathbf{A}_p(t) = \mathbf{A}_p(t+T) \in \mathbb{C}^{N \times N}$ is the periodic LNS operator, ω_f is the fundamental frequency, and $T = 2\pi/\omega_f$ is defined as the fundamental cycle length. As both \mathbf{A}_p and \mathbf{q}' are time-periodic, the Fourier transform of (7) incorporates interactions between all pairs of frequencies. For simplicity, we drop the prime notation for fluctuations from here on out.

Expanding $\mathbf{A}_p(t)$ and $\mathbf{q}(t)$ in terms of the Fourier series

$$\begin{aligned} \mathbf{A}_p(t) &= \sum_{j=-\infty}^{\infty} \hat{\mathbf{A}}_{p,j} e^{ij\omega_f t}, \\ \mathbf{q}(t) &= \sum_{j=-\infty}^{\infty} \hat{\mathbf{q}}_j e^{ij\omega_f t}, \end{aligned} \quad (8)$$

where $(\cdot)_j$ denotes j^{th} harmonic of ω_f , and substituting these expansions into (7), we obtain

$$[\mathbf{T}\hat{\mathbf{q}}]_k = ik\omega_f \hat{\mathbf{q}}_k - \sum_{j=-\infty}^{\infty} \hat{\mathbf{A}}_{k-j} \hat{\mathbf{q}}_j = \hat{\mathbf{f}}_k. \quad (9)$$

Here, \mathbf{T} is an infinite dimensional block matrix of the form

$$\mathbf{T} = \begin{bmatrix} \ddots & \vdots & \vdots & \vdots & \vdots & \ddots \\ \dots & \mathbf{L}_{-1} & -\hat{\mathbf{A}}_{-1} & -\hat{\mathbf{A}}_{-2} & -\hat{\mathbf{A}}_{-3} & \dots \\ \dots & -\hat{\mathbf{A}}_1 & \mathbf{L}_0 & -\hat{\mathbf{A}}_{-1} & -\hat{\mathbf{A}}_{-2} & \dots \\ \dots & -\hat{\mathbf{A}}_2 & -\hat{\mathbf{A}}_1 & \mathbf{L}_1 & -\hat{\mathbf{A}}_{-1} & \dots \\ \dots & -\hat{\mathbf{A}}_3 & -\hat{\mathbf{A}}_2 & -\hat{\mathbf{A}}_1 & \mathbf{L}_2 & \dots \\ \ddots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}, \quad (10)$$

where the diagonal entries are in fact the inverse of resolvent operators at various frequencies,

$$\mathbf{L}_j = \mathbf{R}_j^{-1} = ij\omega_f \mathbf{I} - \hat{\mathbf{A}}_{p,0}, \forall j \in \mathbb{Z}. \quad (11)$$

The harmonic resolvent operator

$$\mathbf{H} = \mathbf{T}^{-1} \quad (12)$$

transfers the harmonic inputs to the outputs for periodic base flows,

$$\hat{\mathbf{q}} = \mathbf{H}\hat{\mathbf{f}}. \quad (13)$$

The SVD of the harmonic resolvent operator

$$\mathbf{H} = \mathbf{U}_H \mathbf{\Sigma}_H \mathbf{V}_H^* \quad (14)$$

unveils the most amplified responses \mathbf{U}_H that correspond to optimal forcing \mathbf{V}_H , each accompanied by associated amplification magnitudes Σ_H . The matrix \mathbf{H} encompasses the base flow, the fundamental frequency, and higher harmonics. Whereas individual frequencies can be analyzed separately in a standard resolvent analysis, all frequencies are coupled within \mathbf{H} in harmonic resolvent analysis and must be considered simultaneously. Furthermore, the modes of the modified harmonic resolvent operator, defined as

$$\tilde{\mathbf{H}} = \mathbf{W}_q^{1/2} \mathbf{C} \mathbf{H} \mathbf{B} \mathbf{W}_f^{-1/2}, \quad (15)$$

can be computed using our algorithm with the same minor adjustments as detailed in Farghadan et al. [22]. In a special case where the base flow is time-independent, *i.e.*, $\hat{\mathbf{A}}_{p,j} = 0$ for $j \neq 0$, \mathbf{H} reduces to

$$\mathbf{H} = \begin{bmatrix} \ddots & \vdots & \vdots & \vdots & \ddots \\ \dots & \mathbf{R}_{-1} & 0 & 0 & \dots \\ \dots & 0 & \mathbf{R}_0 & 0 & \dots \\ \dots & 0 & 0 & \mathbf{R}_1 & \dots \\ \ddots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}, \quad (16)$$

so each frequency is decoupled. Therefore, resolvent analysis can be perceived as a special case of harmonic resolvent analysis.

3.2 Singular or near-singular nature of \mathbf{T}

The operator \mathbf{T} tends to be ill-conditioned for many periodic systems [1]. A system is considered ill-conditioned when the ratio of the largest and smallest singular values is large ($\kappa = \sigma_{\max}/\sigma_{\min} \sim O(10^4)$ or higher). The smallest singular value of \mathbf{T} can be shown to reach machine precision zero if the periodic base flow satisfies the Navier-Stokes equations [1, 26], *i.e.*, when

$$\frac{d\bar{\mathbf{q}}}{dt} = \mathcal{N}(\bar{\mathbf{q}}), \quad (17)$$

where \mathcal{N} is the nonlinear Navier-Stokes operator. Taking a time derivative of both sides,

$$\frac{d\left(\frac{d\bar{\mathbf{q}}}{dt}\right)}{dt} = \frac{d(\mathcal{N}(\bar{\mathbf{q}}))}{dt} = \frac{\partial \mathcal{N}}{\partial \bar{\mathbf{q}}} \frac{d\bar{\mathbf{q}}}{dt} = \mathbf{A}_p \frac{d\bar{\mathbf{q}}}{dt}. \quad (18)$$

Taking a Fourier transform of (18) and recalling that \mathbf{T} is obtained as shown in (9), $\widehat{d\bar{\mathbf{q}}/dt}$ is a non-trivial solution of

$$\mathbf{T} \mathbf{q} = 0, \quad (19)$$

residing within the null space of \mathbf{T} , thereby proving the singularity of \mathbf{T} . In systems where the base flow approximately satisfies (17), \mathbf{T} becomes nearly singular. Consequently, irrespective of the specific periodic system under consideration, \mathbf{T} might be poorly conditioned, presenting even more challenges when dealing with large systems.

Padovan et al. [1] demonstrated that defining the harmonic resolvent operator in a manner that eliminates the phase shift imparted by the Fourier coefficients of $d\bar{\mathbf{q}}/dt$ is effective, without adversely affecting the other dominant amplification mechanisms. By constraining \mathbf{T} to a subspace \mathbf{U}^\perp that is orthogonal to the direction of the phase shift $\widehat{d\bar{\mathbf{q}}/dt}$, the range of the harmonic resolvent operator is limited to \mathbf{U}^\perp . The key concept here is to eliminate the singular vectors associated with the smallest singular value of \mathbf{T} without affecting the other singular values and vectors of \mathbf{T} . Given that the phase shift resides in the null space of \mathbf{T} , the desired right singular vector $\mathbf{v} = \widehat{d\bar{\mathbf{q}}/dt} / \|\widehat{d\bar{\mathbf{q}}/dt}\|$ is readily available. We can also solve

$$\mathbf{T}^* \tilde{\mathbf{u}} = \mathbf{v} \quad (20)$$

and compute the corresponding left singular vector as $\mathbf{u} = \tilde{\mathbf{u}}/\|\tilde{\mathbf{u}}\|$. Upon obtaining \mathbf{v} and \mathbf{u} , the process becomes straightforward: projecting out these modes from the forcing and response terms effectively restricts \mathbf{T} to \mathbf{U}^\perp . These steps are elucidated in detail in Padovan et al. [1].

To compute the action of the harmonic resolvent operator on a given vector (or matrix), our goal is to determine the vector (or matrix) \mathbf{q} that satisfies $\mathbf{q} = \mathbf{H}\mathbf{f}$. In brief, to compute the action of a singular \mathbf{H} on a vector $\tilde{\mathbf{f}}$, the component along \mathbf{u} must be projected out as

$$\tilde{\mathbf{f}}_{in} = \tilde{\mathbf{f}} - \mathbf{u}(\mathbf{u}^* \tilde{\mathbf{f}}), \quad (21)$$

and this will provide a response

$$\mathbf{T}\tilde{\mathbf{q}} = \tilde{\mathbf{f}}_{in} \rightarrow \tilde{\mathbf{q}} = \mathbf{H}\tilde{\mathbf{f}}_{in}, \quad (22)$$

which is orthogonal to the direction of the phase shift. Nevertheless, to mitigate potential round-off errors, it is advisable to refine the output by projecting out the component along \mathbf{v} as

$$\tilde{\mathbf{q}}_{out} = \tilde{\mathbf{q}} - \mathbf{v}(\mathbf{v}^* \tilde{\mathbf{q}}). \quad (23)$$

Similarly, when computing the action of singular \mathbf{H}^* on a vector, the initial step involves projecting out the component along \mathbf{v} . After the response is obtained, to mitigate round-off errors, the component along \mathbf{u} is projected out.

3.3 Truncating the harmonic resolvent operator

Computing harmonic resolvent modes is contingent upon \mathbf{T} being of finite size. The size and block sparsity pattern of \mathbf{T} are determined by two key variables: the number of frequencies within the base flow, N_b , and within the response, N_ω . Here, we elucidate the impact of each variable on the operator.

The rows of \mathbf{T} consist of N_b -stencil blocks, dictated by the presence of non-zero elements in $\hat{\mathbf{A}}_{p,j}$. In the context of periodic flows, the time-dependent base flow can be expressed using a Fourier series expansion

$$\bar{\mathbf{q}}(t) = \sum_{j\omega_f \in \Omega_{\bar{\mathbf{q}}}} \hat{\bar{\mathbf{q}}}_j e^{ij\omega_f t}. \quad (24)$$

The set $\Omega_{\bar{\mathbf{q}}}$ encompasses the relevant frequencies associated with the dominant flow structures. Typically, $\Omega_{\bar{\mathbf{q}}}$ is a subset of $j\omega_f$ with $j \in \mathbb{Z}$ and encapsulates the majority of energy within the periodic flows. By retaining the base periodicity and a limited number of higher harmonics, it becomes possible to simplify the base flow representation, resulting in N_b frequencies within $\Omega_{\bar{\mathbf{q}}}$ and N_b non-zero elements within each row of \mathbf{T} .

On the other hand, we aim to resolve perturbations occurring at temporal frequencies $\omega \in \Omega_q$, where Ω_q represents a subset of $j\omega_f$ with $j \in \mathbb{Z}$, although the specific choice of harmonics for Ω_q may vary, as further elucidated in §3.5. Typically, $\Omega_{\bar{\mathbf{q}}}$ is a subset of Ω_q , *i.e.*, the perturbation frequency content spans a wider range of harmonics [1]. The number of frequencies N_ω expanding both \mathbf{f} and \mathbf{q} determines the dimensions of the block matrices, yielding a size of $(N_\omega N) \times (N_\omega N)$ for \mathbf{T} .

We provide a simple example to visually observe the harmonic resolvent operator. Suppose we use $\Omega_{\bar{\mathbf{q}}} = \{-2, -1, 0, 1, 2\}\omega_f$ to expand the LNS operator and $\Omega_q = \{-4, -3, \dots, 3, 4\}\omega_f$ to expand the perturbations. In this case, the stencil length would be $N_b = 5$, and the size of the block matrix would be $(9N) \times (9N)$. The harmonic resolvent operator is then

$$\mathbf{H} = \begin{bmatrix} \mathbf{L}_{-4} & -\hat{\mathbf{A}}_{-1} & -\hat{\mathbf{A}}_{-2} & 0 & 0 & 0 & 0 & 0 & 0 \\ -\hat{\mathbf{A}}_{-1} & \mathbf{L}_{-3} & -\hat{\mathbf{A}}_{-1} & -\hat{\mathbf{A}}_{-2} & 0 & 0 & 0 & 0 & 0 \\ -\hat{\mathbf{A}}_{-2} & -\hat{\mathbf{A}}_{-1} & \mathbf{L}_{-2} & -\hat{\mathbf{A}}_{-1} & -\hat{\mathbf{A}}_{-2} & 0 & 0 & 0 & 0 \\ 0 & -\hat{\mathbf{A}}_{-2} & -\hat{\mathbf{A}}_{-1} & \mathbf{L}_{-1} & -\hat{\mathbf{A}}_{-1} & -\hat{\mathbf{A}}_{-2} & 0 & 0 & 0 \\ 0 & 0 & -\hat{\mathbf{A}}_{-2} & -\hat{\mathbf{A}}_{-1} & \mathbf{L}_0 & -\hat{\mathbf{A}}_{-1} & -\hat{\mathbf{A}}_{-2} & 0 & 0 \\ 0 & 0 & 0 & -\hat{\mathbf{A}}_{-2} & -\hat{\mathbf{A}}_{-1} & \mathbf{L}_1 & -\hat{\mathbf{A}}_{-1} & -\hat{\mathbf{A}}_{-2} & 0 \\ 0 & 0 & 0 & 0 & -\hat{\mathbf{A}}_{-2} & -\hat{\mathbf{A}}_{-1} & \mathbf{L}_2 & -\hat{\mathbf{A}}_{-1} & -\hat{\mathbf{A}}_{-2} \\ 0 & 0 & 0 & 0 & 0 & -\hat{\mathbf{A}}_{-2} & -\hat{\mathbf{A}}_{-1} & \mathbf{L}_3 & -\hat{\mathbf{A}}_{-1} \\ 0 & 0 & 0 & 0 & 0 & 0 & -\hat{\mathbf{A}}_{-2} & -\hat{\mathbf{A}}_{-1} & \mathbf{L}_4 \end{bmatrix}^{-1}. \quad (25)$$

The forcing and response vectors can be represented as

$$\hat{\mathbf{f}} = \begin{bmatrix} \hat{\mathbf{f}}_{-4} \\ \hat{\mathbf{f}}_{-3} \\ \hat{\mathbf{f}}_{-2} \\ \hat{\mathbf{f}}_{-1} \\ \hat{\mathbf{f}}_0 \\ \hat{\mathbf{f}}_1 \\ \hat{\mathbf{f}}_2 \\ \hat{\mathbf{f}}_3 \\ \hat{\mathbf{f}}_4 \end{bmatrix}, \hat{\mathbf{q}} = \begin{bmatrix} \hat{\mathbf{q}}_{-4} \\ \hat{\mathbf{q}}_{-3} \\ \hat{\mathbf{q}}_{-2} \\ \hat{\mathbf{q}}_{-1} \\ \hat{\mathbf{q}}_0 \\ \hat{\mathbf{q}}_1 \\ \hat{\mathbf{q}}_2 \\ \hat{\mathbf{q}}_3 \\ \hat{\mathbf{q}}_4 \end{bmatrix}. \quad (26)$$

The accuracy of the obtained resolvent system in Fourier space depends on how well the chosen frequencies capture the relevant flow information, and it is affected by the truncation limit imposed on the frequency range. Although it is feasible to reduce the infinite-dimensional harmonic operator to a finite dimension based on perturbation frequencies, the choice of frequency range is consequential. To demonstrate the influence of truncation on response modes, consider a periodic flow with a dominant first harmonic, such that the base flow can be expressed as

$$\bar{\mathbf{q}}(t) = \hat{\mathbf{q}}_0 + \hat{\mathbf{q}}_1 e^{i\omega t} + \hat{\mathbf{q}}_{-1} e^{-i\omega t}. \quad (27)$$

Assume that the LNS equations are subjected to a constant forcing $\mathbf{f}(t) = \hat{\mathbf{f}}_0$. The LNS operator exhibits frequency content at 0 and $\pm\omega$, indicating that a constant forcing will directly induce a response at these frequencies. However, higher harmonics are also indirectly stimulated, and the triggering of responses at higher harmonics follows a cascade pattern, with norms gradually decaying towards zero, *i.e.*, $\lim_{j \rightarrow \infty} \|\hat{\mathbf{q}}_j\| = 0$. Hence, the truncation must be set at a sufficiently high level to preserve the response norms.

In practice, by defining Ω_q for both input perturbation and the response, we compute accurate harmonic resolvent modes as long as the norms of higher frequency modes become less relevant within the spectrum. In other words, if we set both forcing and response frequencies sufficiently high and, upon computing harmonic resolvent modes, find that the norms of input and output modes with higher frequency content are relatively small, it indicates convergence, and including higher frequencies has negligible impact on the results. In a case, for instance, where response modes with higher frequency content remain relatively important, extending the frequency range of the output modes is necessary without changing the input frequency content. In our algorithm, the input and output frequency ranges are adjustable as needed.

3.4 Cross-frequency harmonic resolvent analysis

Harmonic resolvent analysis captures the interaction between various frequencies within $\Omega_{\bar{q}}$ and Ω_q , and the typical objective is to identify the optimal forcing mode, potentially spanning various frequencies, that produces the optimal response across a range of frequencies. Alternatively, one can compute the unit-norm forcing at one frequency $\omega_1 \in \Omega_q$ that triggers the most amplified response at another frequency $\omega_2 \in \Omega_{\bar{q}}$. This is achieved by defining

$$\mathbf{H}_{\omega_2, \omega_1} = \mathbf{C}\mathbf{H}\mathbf{B} \quad (28)$$

and computing the modes and gains following the same procedure as before. Here, matrices \mathbf{B} and \mathbf{C} are constructed to extract the ω_1 and ω_2 frequencies from the forcing and response modes, respectively. This represents a special case of the modified harmonic resolvent operator in (15).

3.5 Subharmonic resolvent analysis

The frequency content of the forcing inputs is typically composed of harmonics of the base flow frequency. However, Padovan and Rowley [15] have shown that the harmonic resolvent system may exhibit sensitivity to subharmonic inputs. For instance, in the case where the fundamental base flow frequency is denoted by ω_f , the linearized system may be sensitive to forcing inputs with $\gamma < \omega_f$, such as $\omega_f/2$. More generally, if the frequency content of a periodic LNS operator is limited to the frequency range $\Omega_{\bar{q}}$, it may be desirable to compute input-output modes with frequencies outside of this range, *i.e.*, $\gamma \notin \Omega_{\bar{q}}$. In such cases, the subharmonic resolvent modes become relevant.

To identify unique subharmonic inputs, we can take advantage of the linear nature of the harmonic resolvent operator, which allows interaction with the fundamental frequency ω_f and its harmonics. It can be demonstrated that the interval $(-\omega_f/2, \omega_f/2]$ encompasses all sets of subharmonic frequencies where the input-output modes are unique. As discussed in §3, a forcing with a frequency $\gamma \in (-\omega_f/2, \omega_f/2]$ can only trigger responses at frequencies $\gamma + j\omega_f$, where j is an integer. Therefore, the input and output modes exist within the frequency range $\Omega_\gamma = \gamma \oplus \Omega_{\bar{q}}$, where \oplus denotes element-wise addition. Harmonics of γ trigger a unique set of outputs as long as $j\gamma \leq |\omega_f|$. Due to the linear relationship, we can show that all harmonics are isolated sets that need to be studied separately. Note that $-\omega_f/2$ is excluded since $\Omega_{-\omega_f/2} = \Omega_{\omega_f/2}$, resulting in redundancy. For $\gamma \in \Omega_\gamma$, one can derive the subharmonic resolvent system as [15]

$$\hat{\mathbf{q}} = (i\gamma\mathbf{I} - \mathbf{T})^{-1}\hat{\mathbf{f}}. \quad (29)$$

To further elucidate the point, consider an example where the subharmonic frequency of interest is $\omega_f/5$ along with its harmonics. Note that the valid harmonics are the ones that fulfill the condition $j\omega_f/5 \in (-\omega_f/2, \omega_f/2]$, which includes the set $\{-2/5, -1/5, 0, 1/5, 2/5\}$. While $\Omega_0 = 0 \oplus \{-N_b/2, \dots, N_b/2\}$, where N_b is the number of active frequencies within the base flow, is identical to harmonic inputs, the sets $\Omega_{-2/5}$, $\Omega_{-1/5}$, $\Omega_{1/5}$, and $\Omega_{2/5}$ represent distinct input-output modes. In general, for any two distinct subharmonic frequencies γ_1 and γ_2 within the interval $(-\omega_f/2, \omega_f/2]$, the corresponding subharmonic modes in Ω_{γ_1} are decoupled from those in Ω_{γ_2} as there exists no common frequency $\gamma_1 - \gamma_2 \in \Omega_{\bar{q}}$. While our algorithm described in §4 is described for harmonic resolvent analysis, it can be easily extended to compute subharmonic resolvent modes, as described in Appendix A.

The approach used to eliminate the singularity of \mathbf{T} described in §3.2 is not applicable for subharmonic resolvent analysis ($\gamma \neq 0$). Padovan and Rowley [15] formulated an alternative approach in which the domain of \mathbf{T} is constrained to a physically meaningful subspace for all choices of γ . In this approach, an oblique projection operator is defined as

$$\mathbf{P}(t) = \mathbf{I} - \mathbf{v}(t)\mathbf{u}^*(t), \quad (30)$$

where \mathbf{u} and \mathbf{v} are defined the same as in §3.2. The restricted operator can be written as

$$\mathbf{T}_p = \mathbf{T}\hat{\mathbf{P}}. \quad (31)$$

The resulting range of \mathbf{T}_p is deliberately designed to be invariant under T . This projection facilitates steps within the RSVD-LU algorithm, effectively addressing singularities. For more details on oblique projection properties and the efficient computation of \mathbf{u} and \mathbf{v} , we refer the reader to Padovan and Rowley [15].

3.6 Computing harmonic resolvent modes using RSVD-LU

The RSVD algorithm [27] is a randomized linear algebra technique developed to efficiently identify the singular vectors with the highest gains in a given matrix. In the context of harmonic resolvent analysis, the matrix of interest is the harmonic resolvent operator. This operator, similar to the resolvent operator, relies on the inverse of a sparse matrix, as shown in (12), making the computation of modes challenging. Modifying the original RSVD algorithm is necessary to address these challenges, enabling harmonic resolvent analysis. The modifications for resolvent analysis have been extensively documented in the literature by Ribeiro et al. [21] and Farghadan et al. [22], while the outline for the harmonic resolvent operator can be found in the appendix of Padovan et al. [1].

In brief, the computation involves computing the action of \mathbf{H} and \mathbf{H}^* for sketching the range and image of the harmonic resolvent operator, respectively, and approximating the leading harmonic modes. Computing actions of \mathbf{H} and \mathbf{H}^* , however, requires solving linear systems

$$\begin{aligned} \mathbf{T}\hat{\mathbf{q}} &= \hat{\mathbf{f}}, \\ \mathbf{T}^*\hat{\mathbf{q}} &= \hat{\mathbf{f}}, \end{aligned} \quad (32)$$

respectively, in Fourier space. These linear systems are solved via LU-decomposition of \mathbf{T} , but due to its large size and ill-conditioned nature, computing its LU decomposition can be a formidable computational

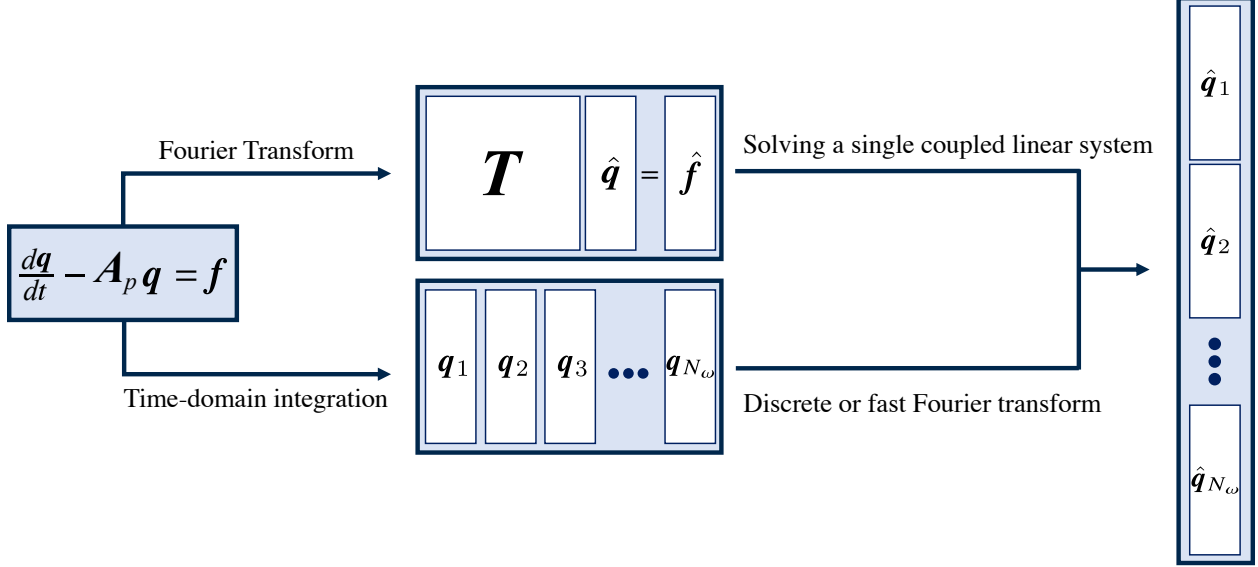


Fig. 1 Flowchart depicting the action of \mathbf{H} on a forcing comprised of N_ω frequencies within the RSVD-LU (top route) and the RSVD- Δt (bottom route) algorithms. Both routes produce the same result, but the bottom route is computationally advantageous for large systems.

obstacle, particularly for flows with three inhomogeneous directions. We propose an alternative way in the following section.

4 Computing harmonic resolvent modes using time stepping

In this section, we show how time stepping can be used to efficiently compute harmonic resolvent modes. The same concept was first used by Monokrousos et al. [23] in the context of the resolvent analysis and further optimized by Martini et al. [24] and Farghadan et al. [22]. Here, we introduce the extension of our algorithm to compute the harmonic resolvent modes and gains.

4.1 Computing the action of \mathbf{H} using time stepping

Computing the action of \mathbf{H} in the Fourier space poses a bottleneck within the RSVD-LU algorithm. To overcome this limitation, we propose to use time-stepping as an effective surrogate approach. Given that $\Omega_{\bar{q}}$ represents the frequency content of the base flow around which \mathbf{T} is constructed, we define the time-domain linearized operator as

$$\mathbf{A}_p(t) = \sum_{j\omega_f \in \Omega_{\bar{q}}} \hat{\mathbf{A}}_{p,j} e^{ij\omega_f t}. \quad (33)$$

The action of \mathbf{H} on $\hat{\mathbf{f}}$ in Fourier space is expressed in (13), where the frequency content of the forcing lies within Ω_q .

The steady-state solution of (7) is given by

$$\hat{\mathbf{q}}_s = \mathbf{H} \hat{\mathbf{f}}, \quad (34)$$

when subjected to the forcing represented as

$$\mathbf{f} = \sum_{j\omega_f \in \Omega_q} \hat{\mathbf{f}}_j e^{ij\omega_f t}. \quad (35)$$

The steady-state response can be obtained through time-stepping before taking a Fourier transform. The time-domain forcing is constructed to include frequencies identical to those present in $\hat{\mathbf{f}}$. Since the time-domain and frequency-domain frequencies are the same, both approaches yield identical results, *i.e.*, $\hat{\mathbf{q}}_s = \hat{\mathbf{q}}$.

In practice, the time-stepping method is inherently discrete, and the frequencies are retained up to the maximum limit determined by the chosen time step (below the Nyquist frequency [28]). By employing a suitably small time step, we can capture frequencies up to the desired limit.

The equivalence between computing the action of \mathbf{H} in both the RSVD-LU and RSVD- Δt algorithms is depicted in figure 1. In the upper route, within the RSVD-LU algorithm, the LNS equations are first transformed into Fourier space before solving a coupled linear system in the frequency domain for a given forcing. In the lower route, within the RSVD- Δt algorithm, the LNS equations are initially integrated in the time domain before being transformed into Fourier space. Both routes produce identical outputs so long as numerical artifacts are minimized.

4.2 RSVD- Δt for harmonic resolvent analysis

By recognizing the equivalence between time stepping and solving linear systems in the frequency domain, we can effectively address the limitations of the RSVD-LU algorithm, leading to the development of RSVD- Δt . The time-stepping aspect of RSVD- Δt relies on the linearized operator \mathbf{A}_p without explicitly constructing \mathbf{H} in Fourier space. The algorithmic steps for extending RSVD- Δt to harmonic systems are presented in Algorithm 1.

Algorithm 1 RSVD- Δt for harmonic resolvent analysis

- 1: **Inputs:** $\mathbf{A}_p, k, q, \Omega_q, \Omega_{\bar{q}}, \text{TSS}, dt, T_t$
- 2: $\hat{\Theta} \leftarrow \text{randn}(NN_\omega, k)$ ▷ Create random test matrices
- 3: $\hat{\mathbf{Y}} \leftarrow \text{DirectAction}(\mathbf{A}_p, \hat{\Theta}, \text{TSS}, dt, T_t)$ ▷ Sample the range of \mathbf{H}
- 4: **if** $q > 0$ **then** ▷ Optional power iteration
- 5: $\hat{\mathbf{Y}} \leftarrow \text{PI}(\mathbf{A}_p, \hat{\mathbf{Y}}, q, \text{TSS}, dt, T_t)$ ▷ Power iteration with time-stepping
- 6: **end if**
- 7: $\hat{\mathbf{Q}}_\Omega \leftarrow \text{qr}(\hat{\mathbf{Y}}_\Omega)$ ▷ Build the orthonormal subspace $\hat{\mathbf{Q}}_\Omega$
- 8: $\hat{\mathbf{S}} \leftarrow \text{AdjointAction}(\mathbf{A}_p^*, \hat{\mathbf{Q}}, \text{TSS}, dt, T_t)$ ▷ Sample the image of \mathbf{H}
- 9: $(\tilde{\mathbf{U}}_H, \tilde{\Sigma}_H, \tilde{\mathbf{V}}_H) \leftarrow \text{svd}(\hat{\mathbf{S}}_H)$ ▷ Obtain Σ_H, \mathbf{V}_H
- 10: $\mathbf{U}_H \leftarrow \hat{\mathbf{Q}}_\Omega \tilde{\mathbf{U}}_H$ ▷ Recover \mathbf{U}_H
- 11: **Outputs:** $\mathbf{U}_H, \Sigma_H, \mathbf{V}_H$

Algorithm 1. Inputs: linearized operator \mathbf{A}_p , number of modes k , number of power iterations q , frequency range of the perturbations and LNS operator Ω_q and $\Omega_{\bar{q}}$, respectively, time-stepping scheme abbreviated as TSS (*e.g.*, backward Euler), time step dt , and the transient length T_t . Outputs: k response modes \mathbf{U}_H , k forcing modes \mathbf{V}_H and the corresponding gains Σ_H . Here, $k, q, \Omega_q, \Omega_{\bar{q}}$ are common parameters with RSVD-LU. $(\cdot)_\Omega$ indicates all frequencies are merged into a single column. **DirectAction** and **AdjointAction** are functions that solve the direct and adjoint LNS equations, respectively, with a predefined forcing. **PI** is a function that performs the power iteration.

The algorithm proposed in this study is based on the RSVD- Δt algorithm introduced by Farghadan et al. [22], and we will provide a concise overview of the key steps. Initially, a random matrix $\hat{\Theta} \in \mathbb{C}^{NN_\omega \times k}$ is generated (line 2) to compute the sketch of \mathbf{H} through time stepping (line 3). The dimensions of the random matrix are determined by the state dimension N , the number of frequencies to resolve N_ω , and the desired number of modes to compute k . **DirectAction** is a function that computes the action of \mathbf{H} onto a given forcing using time-stepping and transforms the steady-state responses to Fourier space. As outlined in §4.1, **DirectAction** solves (7) with zero initial condition over a sufficiently long time interval, causing the initial transient response to dissipate. T_t represents the duration of this interval before the transient diminishes. The forcing term in (7) is constructed via inverse Fourier transform of $\hat{\Theta}$ and the steady-state response undergoes a Fourier transform to obtain $\hat{\mathbf{Y}}$. Next, an optional power iteration (PI) is performed (lines 4 and 5) via q successive applications of **DirectAction** and **AdjointAction**, enhancing the accuracy of harmonic resolvent modes. QR decomposition is performed to obtain $\hat{\mathbf{Q}}_\Omega$ (line 7) before constructing $\hat{\mathbf{S}}$ via time stepping (line 8). **AdjointAction** is the function that computes the action of \mathbf{H}^* onto a given forcing using time-stepping and transforms the steady-state responses to Fourier space. This function is similar to **DirectAction**, but it evolves the adjoint system. The forcing term in **AdjointAction** is constructed via the inverse Fourier

transform of $\hat{\mathbf{Q}}$. An SVD (line 9) is conducted to obtain the optimal forcing modes $\mathbf{V}_H \in \mathbb{C}^{NN_\omega \times k}$ and gains $\boldsymbol{\Sigma}_H \in \mathbb{R}^{k \times k}$ of \mathbf{H} . Lastly, the optimal response modes $\mathbf{U}_H \in \mathbb{C}^{NN_\omega \times k}$ are recovered in line 10.

Compared to the RSVD- Δt algorithm for resolvent analysis [22], two notable differences are present. In the extension of RSVD- Δt , the response modes at various frequencies are merged prior to both the QR decomposition and the SVD steps. This deviation from the original algorithm is motivated by the nature of the problem that requires interactions between frequencies, unlike in resolvent analysis where the QR decomposition and SVD are performed separately for each individual frequency of interest. The second difference is the generation of LNS operators over time, whereas in resolvent analysis, the LNS operator remains constant throughout the integration.

If \mathbf{T} approaches machine precision singularity, it is necessary to project out the forcing along the \mathbf{u} direction before the `DirectAction` and along \mathbf{v} before the `AdjointAction`, as described in §3.2. Moreover, to mitigate numerical artifacts, it is advisable to project out the response along \mathbf{v} and \mathbf{u} after the `DirectAction` and `AdjointAction`, respectively.

5 Computational complexity

A brief comparison is conducted between memory requirements and CPU cost of the RSVD-LU and RSVD- Δt algorithms. We also propose strategies to minimize the memory consumption. Throughout, we assume that the linearized operator \mathbf{A}_p is sparse, *i.e.*, that its number of non-zero entries scales as $\text{nnz}(\mathbf{A}_p) \sim O(N)$. Sparse operators are obtained when using sparse discretization schemes such as finite differences, finite volume, or finite elements. For more detailed information on the CPU and memory scaling of the LU decomposition of the resolvent operator, we refer readers to Farghadan et al. [22]. To offer empirical insights and confirmation of the theoretical scalings discussed below, we present a two-dimensional test case in §7.

5.1 Memory benefits of RSVD- Δt

In the case of resolvent analysis, the memory scaling of the LU decomposition is empirically $O(N^{1.2})$ and $O(N^{1.6})$ for two- and three-dimensional systems, respectively [22, 29]. The memory scaling of LU decomposition is expected to be worse for the harmonic resolvent operator, as the lower- and upper-triangular matrices are denser. When \mathbf{T} is singular or nearly singular, the computational task becomes even more challenging, as discussed in Padovan et al. [1]. An alternative approach for solving (32) involves leveraging iterative solvers, such as Krylov subspace methods [14, 30]. This approach eliminates the need for computing the LU decomposition and becomes more memory-efficient, especially for large systems. However, finding a suitable preconditioner remains a challenging task.

On the other hand, the memory usage of RSVD- Δt for harmonic resolvent analysis is contingent on the sizes of various matrices, specifically the sparse LNS operators $\hat{\mathbf{A}}_p \in \mathbb{C}^{N \times N \times N_b}$, and the dense forcing matrix $\hat{\mathbf{F}} \in \mathbb{C}^{NN_\omega \times k}$ and response matrix $\hat{\mathbf{Q}} \in \mathbb{C}^{NN_\omega \times k}$, all of which are stored in Fourier space. The scaling for resolvent analysis is $O(NN_\omega)$, and the sole extra space is for the storage of LNS operators.

One approach to generating LNS operators for all time points involves retaining N_b coefficients in Fourier space and subsequently constructing $\mathbf{A}_\Omega = \{\mathbf{A}_{p,1}, \mathbf{A}_{p,2}, \mathbf{A}_{p,3}, \dots, \mathbf{A}_{p,N_t}\}$ once for all time steps within a period. However, this approach can be memory-intensive for large systems, especially when $N_t \sim O(10^3 - 10^5)$. An alternative strategy to reduce memory consumption is to create LNS operators on the fly, as shown in figure 2. As a result, no time-domain LNS operator is permanently stored in memory. Each LNS operator is created on the fly using the discrete Fourier transform (DFT) of N_b Fourier coefficients,

$$\mathbf{A}_p(t) = \sum_{j=-N_b/2}^{N_b/2} \hat{\mathbf{A}}_{p,j} e^{ij\omega_f t}. \quad (36)$$

Since \mathbf{A}_p is very sparse, $\text{nnz}(\mathbf{A}_p) \sim O(N)$, resulting in an overall memory consumption of $O(NN_b)$. The methods employed for generating LNS operators, both in the streaming and memory-intensive approaches, mirror those used for creating forcing terms in resolvent analysis, as elaborated in Farghadan et al. [22]. Additionally, the CPU cost of this procedure scales as $O(\text{nnz}(\mathbf{A}_p)N_b)$ or $O(NN_b)$.

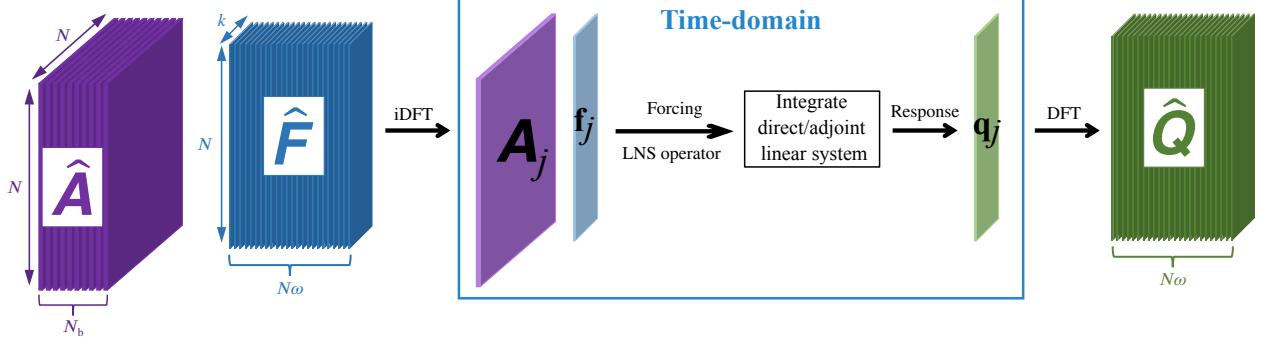


Fig. 2 Schematic of the action of \mathbf{H} with streaming DFT/iDFT to transform between the Fourier and time domains.

Streaming DFT is indeed effective in reducing memory consumption; nevertheless, additional memory savings can be specifically achieved for real-valued LNS operators. In many cases, \mathbf{A}_p is real-valued, *i.e.*, confined to $\mathbb{R}^{N \times N}$, yielding $\hat{\mathbf{A}}_{p,j} = \tilde{\hat{\mathbf{A}}}_{p,-j}$, where $\tilde{(\cdot)}$ denotes the complex conjugate. Hence, $\hat{\mathbf{A}}_{p,j}$ with $j \geq 0$ suffices to construct \mathbf{A}_p as

$$\mathbf{A}_p(t) = \hat{\mathbf{A}}_{p,0} + 2\mathcal{R}e \left(\sum_{j=1}^{N_b/2} \hat{\mathbf{A}}_{p,j} e^{ij\omega_f t} \right), \quad (37)$$

resulting in halving the number of Fourier coefficients of the LNS operator to $\lfloor N_b/2 \rfloor + 1$. Here, $\mathcal{R}e$ represents the real part.

Another opportunity to reduce memory usage for real-valued LNS matrices arises from the symmetry between positive and negative frequencies. Rewriting (9) for $-k\omega_f$ yields

$$i(-k\omega_f)\hat{\mathbf{q}}_{-k} - \sum_{j=-\infty}^{\infty} \hat{\mathbf{A}}_{-k+j}\hat{\mathbf{q}}_{-j} = \hat{\mathbf{f}}_{-k}, \quad (38)$$

and the complex conjugate version of (9) becomes

$$-ik\omega_f\bar{\hat{\mathbf{q}}}_k - \sum_{j=-\infty}^{\infty} \tilde{\hat{\mathbf{A}}}_{k-j}\bar{\hat{\mathbf{q}}}_j = \bar{\hat{\mathbf{f}}}_k. \quad (39)$$

Both equations (38) and (39) yield the same output for a given $\bar{\hat{\mathbf{f}}}_k = \hat{\mathbf{f}}_{-k}$ since $\tilde{\hat{\mathbf{A}}}_{k-j} = \hat{\mathbf{A}}_{-k+j}$, inducing the harmonic resolvent response and forcing modes containing $\mathbf{U}_{H,j} = \bar{\mathbf{U}}_{H,-j}$ and $\mathbf{V}_{H,j} = \bar{\mathbf{V}}_{H,-j}$ for $j \neq 0$, respectively. Hence, we only keep $\lfloor N_\omega/2 \rfloor + 1$ Fourier coefficients of the forcing and response modes, reducing memory consumption by half when storing these dense matrices.

5.2 CPU benefits of RSVD- Δt

The computational cost of computing harmonic resolvent modes using RSVD-LU is predominantly determined by the LU decomposition of \mathbf{T} . In the case of resolvent analysis, the CPU scaling of the LU decomposition is $O(N^{1.5})$ and $O(N^2)$ for two- and three-dimensional systems, respectively [22, 29]. When, in general, block matrix \mathbf{T} contains $\mathbf{A}_{p,i}$ terms, and yields a denser matrix than block diagonal with a complex sparsity pattern, the scaling is expected to be $O((N_\omega N)^{1.5})$ or $O((N_\omega N)^2)$ or worse. Alternatively, iterative solvers may outperform LU decomposition in terms of CPU scalability, contingent on improving the condition number of \mathbf{T} via applying an appropriate preconditioner [14, 30], which itself may or may not be inexpensive to obtain.

The cost of using RSVD- Δt is directly tied to the cost of computing the actions of \mathbf{H} and \mathbf{H}^* . In resolvent analysis, the total CPU cost is proportional to the state dimension N . This cost can be broken down into three main components: (i) time-stepping, *e.g.*, Adams-Bashforth methods, which scales as $O(N)$ when \mathbf{A} is sparse, (ii) creating forcing from Fourier space to the time domain, which scales as $O(NN_tN_\omega)$, (iii) taking

the response from the time domain to Fourier space, which scales as $O(NN_\omega^2)$. Among these components, the time-stepping part is the most significant contributor to the overall cost. The CPU cost of harmonic resolvent analysis can be assessed similarly to resolvent analysis, with one exception. For periodic flows, the LNS operator varies over time, whereas it remains constant for statistically stationary flows. We showed in §5.1 that the cost of creating LNS operators over time scales as $O(NN_b)$ for sparse matrices.

In summary, the CPU and memory costs of RSVD- Δt exhibit linear scaling with respect to the state dimension N , regardless of whether the base flow is steady or periodic. Moreover, the dominant terms in the CPU cost are independent of the number of retained frequencies N_ω . Both of these properties make the RSVD- Δt algorithm considerably more scalable than RSVD-LU for harmonic resolvent analysis.

6 Minimizing the CPU cost of RSVD- Δt

Similar to time stepping in the context of resolvent analysis, periodic systems may also experience a slow decay of transient response, which is undesirable, since we require the steady-state response in isolation to compute the action of the harmonic resolvent operator using time stepping. This phenomenon might be linked to instances when \mathbf{T} is near singularity. Regardless, in cases where the LNS equations are absolutely stable but exhibit slowly decaying modes, we propose a strategy that can effectively reduce the CPU cost of time stepping.

6.1 Stability analysis: Floquet theorem and transient response

The steady-state response of the LNS equations, subject to forcing as described in equation (7), is of interest for our analysis. In the case of a linear time-invariant system where \mathbf{A} is independent of time, the decay rate is controlled by the least-damped eigenvalue of \mathbf{A} . This decay rate can be estimated by integrating a homogeneous system over a sufficiently long period. However, for a linear time-periodic system, the decay rate is determined by the least-damped Floquet exponent.

Assuming $\Phi(t)$ is the fundamental solution of (7), the monodromy matrix is constructed as

$$\mathbf{M} = \Phi(0)^{-1}\Phi(T), \quad (40)$$

where T is the period of the flow. When $\Phi(0) = \mathbf{I}$, $\Phi(t)$ is referred to as the principal fundamental matrix, and the monodromy matrix simplifies to $\Phi(T)$. In other words, we evaluate the principal fundamental matrix at the end of the first period. The eigenvalues of \mathbf{M} , known as Floquet multipliers μ_j , allow us to determine the Floquet exponents $\lambda_j = \log(\mu_j)/T$. The least-damped Floquet exponent has the smallest real part and it determines the decay rate of the transient response. If any mode is located in the unstable half-plane, *i.e.*, $\text{Re}(\lambda_j) > 0$, the system is globally unstable.

It is worth noting that the transient length remains independent of the steady-state period. This characteristic is particularly important as it ensures that the length of integration does not rely on T , thereby avoiding potential deterioration in the performance of RSVD- Δt for flows with low periodicity. In practice, both finding the principal fundamental matrix and determining the Floquet exponents can be computationally expensive, especially for large-scale systems. As with autonomous cases, it is possible to run a homogeneous simulation to analyze the long-term behavior of the transient response for periodic flows. This equivalence has been illustrated on a test case in §7.

If \mathbf{T} is singular but all other modes are stable, then the Floquet exponent with the smallest real part will be zero with $d\bar{\mathbf{q}}/dt$ as the corresponding Floquet mode. This is a well-known fact from Floquet theory [10] and can also be seen from (18), where the phase shift direction is a non-trivial solution of the homogenous system. Leclercq and Sipp [26] have also shown that $d\bar{\mathbf{q}}/dt$ is associated with a zero Floquet exponent under the assumption that the base flow satisfies the Navier-Stokes equations as in (17).

6.2 Efficient transient removal

Our strategy leverages the periodic nature of the steady-state part and the linear evolution of the transient part of the solution to directly compute and eliminate the undesired transient portion. For a given pair of solutions \mathbf{q}_1 and \mathbf{q}_2 at times t_1 and $t_2 = t_1 + T$, respectively, they can be expressed as a sum of their

steady-state and transient components as

$$\begin{aligned}\mathbf{q}_1 &= \mathbf{q}_{s,1} + \mathbf{q}_{t,1}, \\ \mathbf{q}_2 &= \mathbf{q}_{s,2} + \mathbf{q}_{t,2},\end{aligned}\tag{41}$$

where $\mathbf{q}_{t,1}$ and $\mathbf{q}_{t,2}$ represent the transient part which decays to zero as $t \rightarrow \infty$ and $\mathbf{q}_{s,1}$ and $\mathbf{q}_{s,2}$ denote the steady-state part which evolves periodically. The time distance between two snapshots is one period, hence

$$\mathbf{q}_{s,2} = \mathbf{q}_{s,1}.\tag{42}$$

Also, the evolution of the transient part can be expressed as

$$\mathbf{q}_{t,2} = \Phi(T)\mathbf{q}_{t,1},\tag{43}$$

where Φ is the principal fundamental matrix of (7). Simplifying (41), (42), and (43) leads to

$$(\Phi(T) - \mathbf{I})\mathbf{q}_{t,1} = \mathbf{b},\tag{44}$$

where $\mathbf{b} = \mathbf{q}_2 - \mathbf{q}_1$ is known.

By obtaining $\mathbf{q}_{t,1}$ from solving (44), the steady-state solution can be recovered as $\mathbf{q}_{s,1} = \mathbf{q}_1 - \mathbf{q}_{t,1}$. The central challenge in removing the transient is the computational cost associated with solving the linear system in (44). Instead, we employ Petrov-Galerkin (or Galerkin) projection to obtain an approximate solution to (44) in a more cost-effective manner.

A low-dimensional representation of the transient part can be expressed as

$$\mathbf{q}_{t,1} = \phi\beta_1,\tag{45}$$

where $\phi \in \mathbb{C}^{N \times r}$ is an orthonormal low-dimensional trial basis with $r \ll N$, and $\beta_1 \in \mathbb{C}^r$ represents the coefficients describing the transient response in this basis. Substituting (45) into (44), the linear system

$$(\Phi(T) - \mathbf{I})\phi\beta_1 = \mathbf{b}\tag{46}$$

is overdetermined. We use Petrov-Galerkin projection with a low-dimensional test basis $\psi \in \mathbb{C}^{N \times r}$ to close (46), yielding

$$\tilde{M}\beta_1 = \psi^*\mathbf{b},\tag{47}$$

where

$$\tilde{M} = \psi^*(\Phi(T) - \mathbf{I})\phi\tag{48}$$

is the map between coefficients. Solving (47) for β_1 is inexpensive due to its reduced dimension, and from (45) the transient estimate is obtained as

$$\mathbf{q}_{t,1} = \phi(\psi^*\phi - \tilde{M})^{-1}\psi^*\mathbf{b}.\tag{49}$$

The reduced operator \tilde{M} is obtained by integrating the columns of ϕ for one period, giving $\Phi(T)\phi$, and then projecting $(\Phi(T) - \mathbf{I})\phi$ against the columns of ψ . This integration happens once and its CPU cost is equivalent to integrating the LNS equations for r periods. The analogous process occurs for the adjoint equations.

How should the trial and test bases be selected? We have found a trial basis spanning the least-damped modes to be effective. Rather than determining these modes through an exhaustive Floquet analysis, we employ a homogenous simulation as illustrated in the top row of the flowchart in figure 3. By initiating a sufficiently long simulation from a normalized random initial condition, the system asymptotically converges to the least-damped modes. The longer the integration length, the lower-dimensional the space becomes. As demonstrated in a test case in §7, the dimension of the trial bases can, in some cases, be as small as a single column. Our test cases employ Galerkin projection with identical trial and test bases.

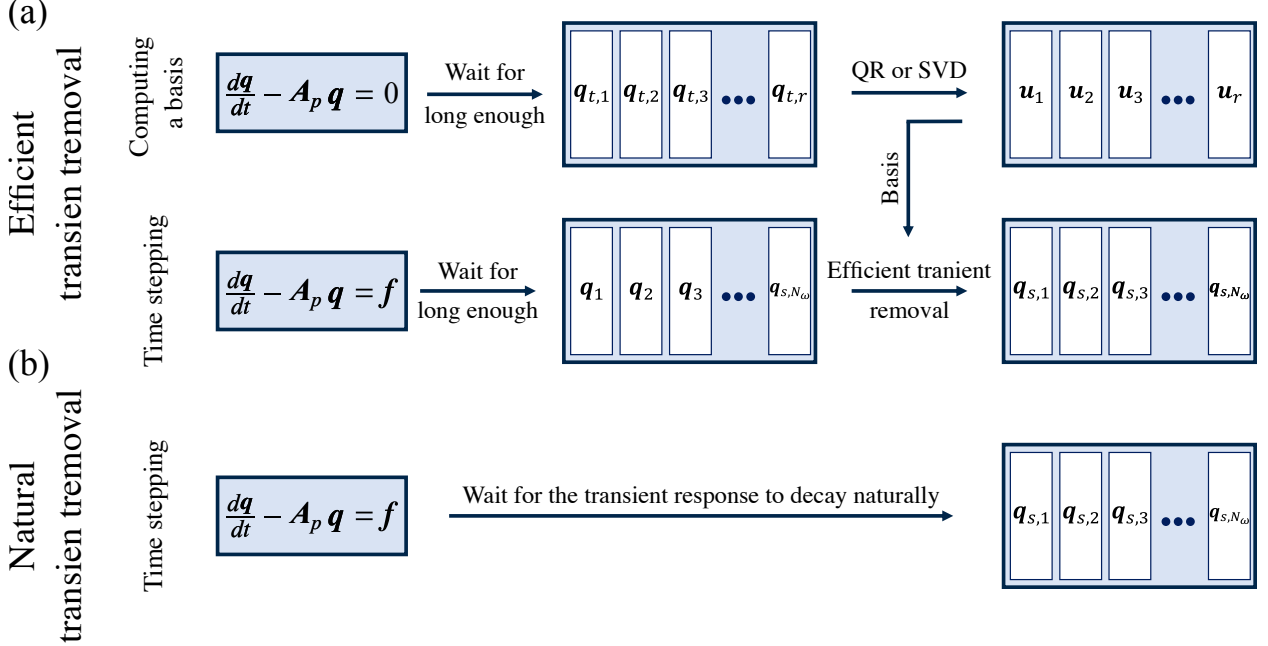


Fig. 3 Flowchart depicting the action of \mathbf{H} using time stepping (a) with efficient transient removal strategy and (b) with natural transient decay.

The obtained trial basis effectively captures the transient response at a specific phase θ . Utilizing ϕ and performing Galerkin projection, we obtain $\tilde{\mathbf{M}}$. The time stepping is then carried out for a sufficiently long duration to allow the initial transients to vanish before applying (47) to determine the transient at phase θ as shown in the middle row of the flowchart in figure 3. Once the steady-state response at the same phase is updated as $q_{s,1} = q_1 - q_{t,1}$, it is used as a new initial condition, which is synchronized with the forcing and will not excite a transient response (within the span of the bases used to construct $\tilde{\mathbf{M}}$). Then, integration for one period is sufficient, and N_ω steady-state snapshots are collected as desired. The resulting steady-state snapshots are identical to those obtained by a prolonged wait in the case of natural decay, as illustrated in the bottom row of the flowchart in figure 3. This procedure is also applicable to the adjoint LNS equations. Given that the adjoint equations differ from the LNS equations, a new basis needs to be constructed for them. Implementing this strategy can reduce the integration length by a factor of 10 or more, depending on the desired accuracy.

7 Test cases

We assess the effectiveness of RSVD- Δt in periodic systems using two test cases. First, we use a modified Ginzburg-Landau equation to validate the RSVD- Δt algorithm and to illustrate the transient removal strategy. Second, we consider a periodic flow around an airfoil, similar to a test case from Padovan et al. [1], to evaluate the accuracy and cost-savings of the RSVD- Δt algorithm compared to RSVD-LU.

7.1 Periodically varying complex Ginzburg-Landau equation

The one-dimensional complex Ginzburg-Landau equations is a widely used model for understanding and controlling the non-modal growth of instabilities in transitional and turbulent shear flows [31–35]. Here, a modified system is used as an inexpensive test case to validate our algorithm. The original complex

Ginzburg-Landau system follows the form of (1) with

$$\begin{aligned}\mathcal{A} &= -\nu \frac{\partial}{\partial x} + \gamma \frac{\partial^2}{\partial x^2} + \mu(x), \\ \mu(x) &= (\mu_0 - c_\mu^2) + \frac{\mu_2}{2} x^2, \\ \mathbf{B} &= \mathbf{C} = \mathbf{I}.\end{aligned}\tag{50}$$

Following Bagheri et al. [33], we set $c_\mu = 0.2$, $\mu_2 = -0.01$, $\gamma = 1 - i$, and $\nu = 2 + 0.2i$. This system is globally stable when $\mu_0 < \mu_{0,cr} \approx 0.3977$ [33]. By substituting $\mu_0(t) = \bar{\mu}_0 + \mu_p \sin(\omega_f t - \pi/2)$ in place of μ_0 , we transform the Ginzburg-Landau equations to a periodically varying system with fundamental frequency ω_f , following the form of (7). We set $\mu_0 = 0.395$, $\mu_p = 0.1$, and $\omega_f = 0.1$; the mean value of $\mu(t)$ is equal to μ_0 , whose value is close to the critical value $\mu_{0,cr}$ to promote significant growth. The periodic linearized operators $\mathbf{A}_{p,j}$ are constructed using a central finite-difference method for $x \in [-50, 50]$ with $N = 1000$ grid points, and we set $\mathbf{W} = \mathbf{I}$ on account of the uniform grid.

Nine operators are built within the interval $[0, T)$, with $T = 2\pi/\omega_0 \approx 62.83$, such that $\mathbf{A}_p(t) = \mathbf{A}_p(t + T)$, before conducting a DFT to obtain $\hat{\mathbf{A}}_p$. Figure 4(a) depicts the Frobenius norm of $\hat{\mathbf{A}}_p$ up to the fourth harmonic, normalized relative to the maximum norm. The spectrum reveals that only the mean and the first harmonic are active, as expected, since $\mu_0(t)$ oscillates at a single frequency. Accordingly, \mathbf{A}_p can be reconstructed at any given time using $\hat{\mathbf{A}}_{p,0}$, and $\hat{\mathbf{A}}_{p,\pm 1}$, and the harmonic resolvent operator is constructed around $\Omega_{\bar{q}} = \{-1, 0, 1\}\omega_f$, *i.e.*, $N_b = 3$. Notably, the first harmonic's norm is several orders of magnitude smaller than that of the base flow. Nonetheless, it should not be disregarded, as we will elucidate its significance in subsequent discussions.

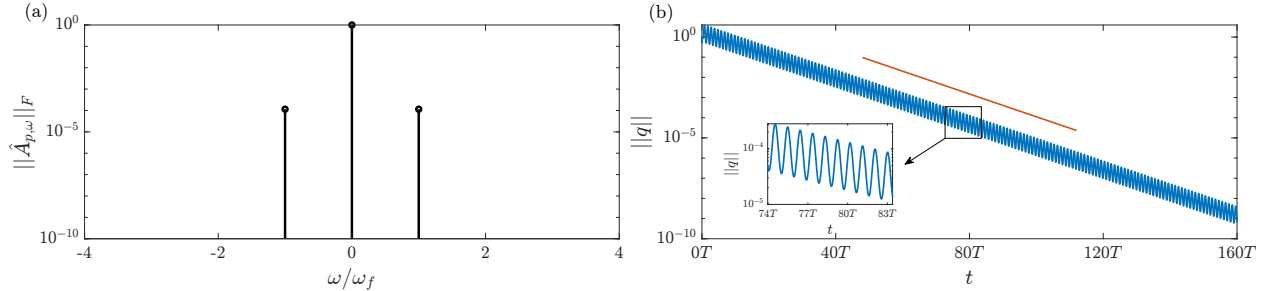


Fig. 4 Ginzburg-Landau test case: (a) the normalized Frobenius norm of $\hat{\mathbf{A}}_{p,\omega}$ up to the fourth harmonic. (b) The norm of the transient response over time (blue) along with the expected decay rate from Floquet analysis (red).

An essential initial step involves ensuring the stability of \mathbf{A}_p . Typically, this is achieved by integrating the homogeneous system, *i.e.*, integrating (7) with zero forcing, starting from a normalized random initial condition. However, the compact size of this system allows us to validate the time-stepping approach for computing the least-damped decaying mode against a Floquet stability analysis. The least-damped Floquet exponent, computed as $\lambda_l = -0.0021 - 0.0477i$, is expected to govern the decay rate of the transient response. Figure 4(b) displays the norm of snapshots of the homogeneous system response over time, along with a reference line representing $e^{\lambda_l t}$. The natural decay rate aligns with the reference line, as expected, ensuring that the transient run will ultimately converge to the least-damped Floquet mode. Additionally, the system is not singular, as the transient dynamics never reach a naturally stable state, and none of the Floquet exponents have a zero real value.

Before proceeding to the computation of harmonic resolvent modes and gains using RSVD- Δt and RSVD-LU, the transient removal strategy is demonstrated, and computing the actions of \mathbf{H} and \mathbf{H}^* are validated for a given random forcing containing $\Omega_q \in \{-10, -9, \dots, 9, 10\}\omega_f$, *i.e.*, $\omega \in [-1, 1]$ with $\Delta\omega = 0.1$. The responses have been computed up to the 20th harmonic, *i.e.*, $\omega \in [-2, 2]$, using both time stepping and directly in Fourier space as outlined in (13). We used a 4th order Runge-Kutta (RK4) integration scheme and a time step of $dt = 0.001$ and $T_t = 15000$ for this purpose.

Figure 5(a) depicts the spectrum of response norms. In order to highlight the substantial difference in output generated by the harmonic resolvent operator around the base flow ($\Omega_{\bar{q}} = \{0\}\omega_f, N_b = 1$), the norm of the responses for the same input forcing is also presented. Figure 5(b) displays the relative errors. The harmonic resolvent with $N_b = 1$ produces inaccurate results across the entire spectrum and has zero response for $|\omega| > 1$ since the modes at different frequencies are entirely decoupled. Therefore, neglecting $\hat{\mathbf{A}}_{p,\pm 1}$ is impractical, as mentioned earlier, and induces significant alterations in the harmonic resolvent modes and gains. On the other hand, the time-stepping output exhibits almost machine-precision accuracy within $|\omega| \leq 1$, and its accuracy diminishes as response norms decrease within $|\omega| > 1$. The validation of \mathbf{H}^* follows a similar procedure but is omitted here for brevity.




Fig. 5 Ginzburg-Landau test case: (a) the norms of the Fourier space response for $N_b = 3$ (black) and $N_b = 1$ (blue), along with the norm of the steady-state response obtained via time-stepping (red). (b) The relative errors of the norms, with the norms from Fourier space with $N_b = 3$ (black line in (a)) serving as the reference.

The transient norm decay rate, as shown in figure 4(b), follows $e^{-0.0021t}$. This indicates that approximately 40 periods are necessary for the transient norm to decay to 1%. However, we can significantly reduce this duration by employing our transient removal strategy. To begin, we precompute the trial basis by conducting a sufficiently long homogeneous simulation, and we find that a single mode suffices. Next, we integrate the orthogonalized basis for one period to obtain \tilde{M} and complete the preprocessing step.

Employing the same forcing as the validation case, we obtain the responses in time at the identical phase (initiated at $t = 0$ or $\theta = 0$ with a time interval of $\Delta t = T$), and compare them with the solutions after the removal of transients. Figure 6(a) presents the norm of the responses, which illustrates the effectiveness of our strategy commencing from the second period. This plot suggests that after two periods using our transient removal strategy with Galerkin projection, we obtain an accurate steady-state snapshot which can be used as the synchronized initial condition for the forced equations. Figure 6(b) demonstrates that the snapshots ultimately align perfectly with the updated snapshot. Hence, using our strategy, a total of three periods are required – two periods before transient removal and one period after the initial condition is obtained – to acquire all N_ω steady-state snapshots, a 10-fold speed-up compared to the natural decay. A similar observation was made regarding the adjoint equations. In brief, our strategy can significantly reduce computational costs, potentially by one or more orders of magnitude, depending on the desired level of accuracy.

Finally, RSVD- Δt employing an RK4 integration scheme with $dt = 0.003$ and $T_t = 2T$ is carried out along with our efficient transient removal strategy to compute the harmonic resolvent modes for $k = 5$ and $q = 1$. To establish a reference, the harmonic resolvent modes are also computed using RSVD-LU with an equivalent number of test vectors and power iterations. Figure 7(a) illustrates that RSVD- Δt computes the five leading gains identically to RSVD-LU. The gain relative error in figure 7(b) confirms that relative errors remain below 10^{-9} across all modes. The inner products of response and forcing modes, computed via RSVD- Δt and RSVD-LU, demonstrate parallel directions as the relative errors remain below 10^{-8} across all modes. The




Fig. 6 Ginzburg-Landau test case: (a) the response norms before (black) and after (red) transient removal at the same phase, *i.e.*, snapshots with $\Delta t = T$ interval. (b) Comparison of the normalized response at the second period in blue, the 50th period in black, and the second period after transient removal in red.

inner-product error between two unit-norm vectors \mathbf{v}_1 and \mathbf{v}_2 is defined as

$$e_{ip} = 1 - \langle \mathbf{v}_1, \mathbf{v}_2 \rangle. \quad (51)$$

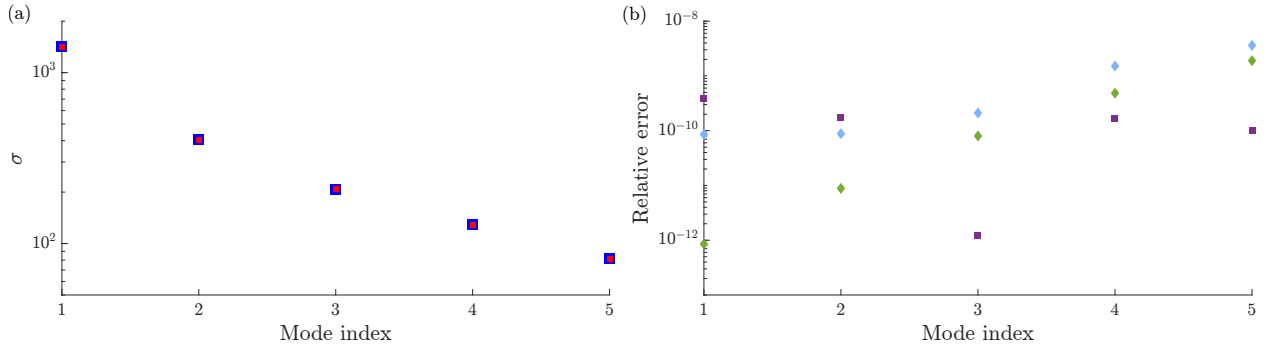


Fig. 7 Ginzburg-Landau test case: (a) the five leading gains using RSVD- Δt (blue) and RSVD-LU (red). (b) The corresponding gain relative error (purple) and relative inner product errors (51) between the response modes (cyan) and the forcing modes (green) computed via RSVD- Δt and RSVD-LU.

7.2 Flow over an airfoil

Our second test case consists of the flow over a NACA0012 airfoil at a low Reynolds number of $Re = 200$ and a steep angle of attack of $\alpha = 20^\circ$, which is dominated by periodic motions. This problem serves as a benchmark for computational cost and accuracy comparisons. A direct numerical simulation is conducted using the “CharLES” compressible flow solver. To mimic the incompressible simulations by Padovan et al. [1], we set the Mach number to 0.05. The simulation employs a C-shaped mesh with a total of 62,000 cells. The leading edge of the airfoil is located at the origin $(x/L_c, y/L_c) = (0, 0)$, where $L_c = 1$ is the chord length. The computational domain spans $x/L_c \times y/L_c \in [-49, 50] \times [-50, 50]$. A constant time step of $\Delta t U_\infty / L_c = 6.88 \times 10^{-5}$ is used, corresponding to a Courant-Friedrichs-Lewy number of 0.91, where U_∞ is the inflow streamwise velocity. Integration continues for a sufficient duration until the flow reaches a periodic state. Figure 8(a) displays the power spectral density computed from the transverse velocity at $(x/L_c, y/L_c) = (3.0, -0.43)$, where a pronounced vortex shedding occurs behind the trailing edge. The shedding frequency is $St_f = 0.114$, where the Strouhal number is defined as $St = \frac{\omega L_c \sin(\alpha)}{2\pi U_\infty}$. The consistency between the dominant frequency identified in the power spectral density and obtained from our linearized code has been confirmed in our previous work [36].




Fig. 8 Airfoil test case: (a) the PSD spectrum based on transverse velocity at $(x/L_c, y/L_c) = (3.0, -0.43)$. (b) The normalized Frobenius norm of $\hat{\mathbf{A}}_{p,St}$ up to the eighth harmonic.

The linearized operators are constructed at 100 time points within $T = 2\pi/St_f \approx 55$. Performing a DFT on the \mathbf{A}_i , we generate 100 LNS operators $\hat{\mathbf{A}}_j$ in Fourier space to be used for harmonic resolvent analysis. The Frobenius norms $\|\hat{\mathbf{A}}_{p,St}\|_F$, depicted in figure 8, indicate that $\mathbf{A}_p(t)$ can be accurately reconstructed using up to the 5th harmonic. Given that the linearized operator $\mathbf{A}_p(t)$ is real-valued for this problem, we retained only $\lfloor N_b/2 \rfloor + 1 = 6$ coefficients, *i.e.*, the zeroth and 5 positive harmonics. For time stepping, $dt = 0.0045$ is chosen to ensure the stability and accuracy of the RK4 integration scheme. The input and output perturbation frequency range spans up to the 7th harmonic, *i.e.*, $N_\omega = 15$. The domain of interest is $x/L_c \times y/L_c \in [-4, 12] \times [-2.5, 2.5]$, identical to that in Padovan et al. [1]. We seek optimal forcing and response under the Chu energy norm [37], which has been used in several previous studies [3, 13, 38]. The number of test vectors is set to $k = 5$, and $q = 2$ power iteration proves sufficient for the convergence of both gains and modes.

A homogeneous simulation of the time-periodic system exhibits a slow decay rate, so we employ the transient removal strategy. Three snapshots constitute the trial basis, obtained after a sufficiently long time interval where most initial transients are eliminated. A random forcing, encompassing the set $\Omega_q \in \{-7, -6, \dots, 6, 7\}St_f$ frequencies, is applied to the LNS equations for over 30 periods. The snapshot norms slowly approach towards the steady-state norm but remain far from converging, as illustrated in figure 9. Utilizing our strategy, the relative norm error falls below the 1% threshold after 20-30 periods in most cases with different random forcings. The efficacy of our strategy is further emphasized in figure 9(b), where the 1% relative error is achieved before 10 periods for the optimal forcing (from RSVD- Δt output). Ultimately, we set $T_t = 29T$ and conducted our simulations for a total duration of 30 periods to compute the actions of \mathbf{H} and \mathbf{H}^* . Extrapolating the natural decay rate suggests that, in the absence of our transient-removal strategy, the norms reach the 1% threshold after around 2000 periods, demonstrating over a 60-fold acceleration in time stepping and overall algorithm performance with our strategy. This observation holds true for the adjoint equations as well.




Fig. 9 Airfoil test case: (a) the response norms before (black) and after (red) transient removal, and the true steady-state norm (blue), at the same phase, *i.e.*, snapshots with $\Delta t = T$ interval, for a random forcing. (b) The relative error between the steady-state response norm and the response norms before (black) and after (red) transient removal.

The optimal forcing and response modes (apart from the phase shift mode) are forcing and response modes are shown in figure 10. The vorticity patterns of the first output mode of the harmonic resolvent closely predict vortical structures observed in simulations that are induced by sinusoidal perturbations to the periodic base flow across all frequencies[1]. This close match is attributed to the low-rank nature of the harmonic resolvent operator of the airfoil. This agreement also implies that regardless of the forcing type, resultant flow perturbations resemble the suboptimal output modes. Moreover, the proximity of forcing modes to the airfoil suggests sensitivity to perturbations near it. The significance of harmonic resolvent analysis lies in its ability to uncover flow structures that would remain hidden when linearizing around the temporal mean. It reveals the intricate interplay between different frequencies, the base frequency and its harmonics, which collectively contribute to the observed flow patterns during real perturbation analysis.

The results from RSVD- Δt closely align with the existing data obtained from RSVD-LU by Padovan et al. [1], despite variations in the CFD solver, boundary conditions, domain setup, and energy norm. The gain plot in figure 11(a) exhibits a similar pattern to Padovan et al. [1], with the exception that we have retained the optimal gain associated with the phase shift (the exact values differ due to differences in the problem setup). We noted that the leading mode and gain, *i.e.*, the mode associated with the phase shift, converged without power iterations due to a substantial two-orders-of-magnitude gap between the optimal and suboptimal gains.

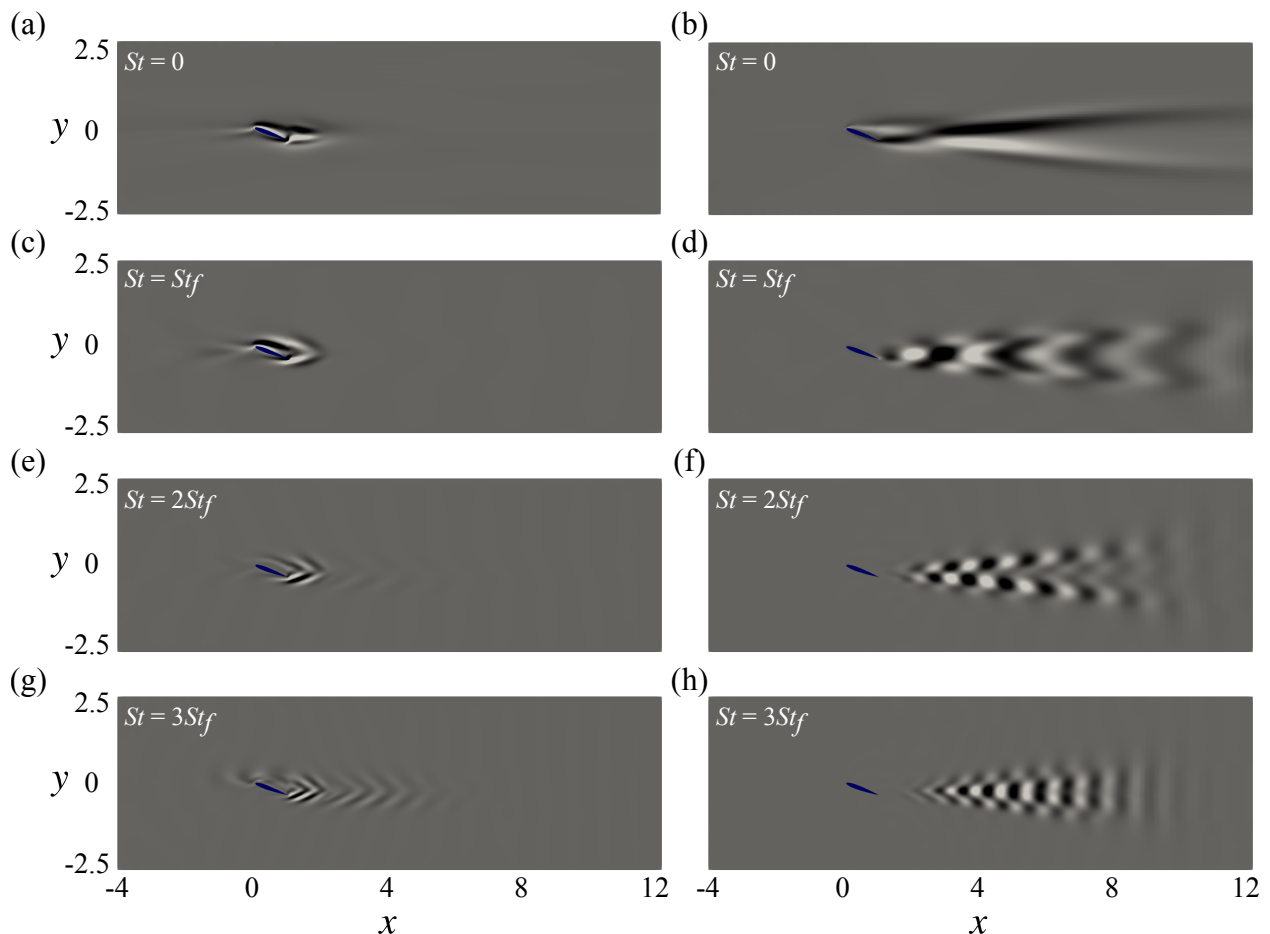


Fig. 10 Airfoil test case: real part of the vorticity field computed from (a, c, e, g) the input mode and (b, d, f, h) the output mode associated with the first suboptimal gain.

From a computational standpoint, a fair comparison between the RSVD-LU and RSVD- Δt algorithms is feasible when both utilize the same parameters. However, the use of $N_b = 11$ proves to be excessively

memory-intensive for the RSVD-LU algorithm, exceeding the available 3.5 TB of memory on our cluster. The CPU and memory scaling plots in figure 11(b) and (c), respectively, depict the LU decomposition of \mathbf{T} with $N_b = 3$, maintaining a constant N while varying the number of blocks N_ω from 5 to 15 with $\Delta N_\omega = 2$. On the same plots, we present the cost of computing the action of \mathbf{H} on a vector using time-stepping. The total duration is set to 30 periods to obtain an accurate solution. We only vary N_ω while keeping $N_b = 3$, the overall cost of RSVD- Δt remains constant and unaffected by changes in N_ω . This implies that the creation of LNS operators and the time-stepping process incur significantly higher costs compared to the transformations between Fourier space and the time domain for forcing and response. The memory of RSVD- Δt to store LNS operators remains independent of N_ω , as expected (see §5.1), while storing \mathbf{Q} and \mathbf{F} matrices grows linearly with the number of frequencies N_ω .

RSVD-LU is estimated to require 1543 CPU-hours for the LU decomposition of \mathbf{T} , and an additional 15 CPU-hours for solving the LU-decomposed system for each test vector, yielding a total cost of $\text{CPU}_{\text{RSVD-LU}} \approx 1543 + 15 \times 5 \times 2 \times 2 = 1843$ CPU-hours for $k = 5, q = 1$. On the other hand, employing RSVD- Δt with the aforementioned time-stepping parameters and $N_b = 11$ and $N_\omega = 15$ incurs a cost of approximately 7 CPU-hours per period for a test vector, resulting in a total cost of $\text{CPU}_{\text{RSVD-}\Delta t} \approx 30 \times 7 \times 5 \times 2 \times 2 = 4200$ CPU-hours. The cost of QR decomposition and the final SVD are ignored as they are orders of magnitude faster. In terms of memory consumption, RSVD-LU peaks at $\text{RAM}_{\text{RSVD-LU}} \approx 2036\text{GB}$ for the LU decomposition of \mathbf{T} . RSVD- Δt , on the other hand, stores $\lfloor N_b/2 \rfloor + 1 = 6$ sparse matrices, each of size 0.4 GB, and three dense matrices of size $Nk(\lfloor N_\omega/2 \rfloor + 1) = 313630 \times 5 \times 8$, totaling $\text{RAM}_{\text{RSVD-}\Delta t} \approx 0.4 \times 6 + 0.185 \times 3 \approx 3\text{GB}$. This translates to a memory saving of approximately three orders of magnitude, even with an unbalanced number of base frequencies. Overcoming the memory consumption hurdle, which is typically the limiting factor in practice, RSVD- Δt emerges as a viable tool for analyzing larger-scale flows compared to RSVD-LU.




Fig. 11 Airfoil test case: (a) the five leading gains using RSVD- Δt . (b) The CPU-hour, and (c) memory usage scaling of RSVD-LU (red) and RSVD- Δt (blue) to compute the action of \mathbf{H} onto a vector, *i.e.*, $k = 1$. The memory usage of RSVD- Δt is decomposed into memory required to store LNS operators (solid) and forcing and response matrices (dashed).

8 Conclusions

This paper presents an extension of the RSVD- Δt algorithm that was originally developed for resolvent analysis of steady base flows to handle harmonic resolvent analysis of periodic flows. Specifically, we demonstrate that the time-stepping technique employed within RSVD- Δt can effectively replace the actions of \mathbf{H} and its complex conjugate \mathbf{H}^* in Fourier space. In terms of CPU cost, the RSVD- Δt algorithm exhibits a scaling of $O(N)$ for both resolvent and harmonic resolvent analyses. This offers a significant advantage, considering that the LU decomposition of \mathbf{T} scales as $O((N_\omega N)^c)$ with $c \geq 1.5$. Regarding memory usage, our algorithm only stores relevant matrices in Fourier space and exhibits $O(NN_\omega)$ scaling; in contrast, the memory peak of LU decomposition of \mathbf{T} empirically scales with $O((N_\omega N)^{1.5})$ or worse. One difference between applying RSVD- Δt to resolvent and harmonic resolvent analyses lies in the necessity to update the LNS operators during the time-stepping process for the latter. These operators are efficiently generated on the fly from their N_b Fourier components.

The error sources in our algorithm align with those in resolvent analysis, extensively investigated in Farghadan et al. [22]. A unique contribution of this paper is the introduction of a novel transient removal

strategy for harmonic resolvent analysis. While reminiscent of our approach for resolvent operators, the challenge in harmonic resolvent analysis lies in the intertwining of all retained frequencies. Our strategy takes advantage of the differing evolution of the steady-state and transient components of the response and Petrov-Galerkin or Galerkin projections.

A validation of RSVD- Δt against RSVD-LU is conducted using a periodic Ginzburg-Landau system. Additionally, we verify the governing role of the Floquet exponent in determining the decay rate of the least-damped mode of the system. Extending the application of our algorithm, we analyze a two-dimensional flow passing an airfoil, providing insights into the CPU and memory complexities. Despite dealing with a mid-sized flow scenario, a substantial memory gap persists between RSVD-LU and RSVD- Δt algorithms. The computed forcing and response modes corresponding to the first suboptimal gain closely resemble those obtained by Padovan et al. [1]. The adaptation of the transient removal strategy for these test cases significantly enhanced the performance of RSVD- Δt , resulting in a 10-fold speed-up for the Ginzburg-Landau problem and a 60-fold speed-up for the airfoil to reach a 1% relative error. The speed-up is even greater for lower error tolerances.

Moreover, our algorithm exhibits versatility, accommodating non-identity weight, input, and output matrices. This capability extends to computing modes for the modified harmonic resolvent operator, and in particular, the analysis of cross-frequency amplification mechanisms via $\mathbf{H}_{\omega_2, \omega_1}$. Our algorithm tackles the computation of subharmonic resolvent modes if desired. Implementation of the RSVD- Δt algorithm within Petsc [39] and SlepC [40] environments leverage parallel computations for enhanced efficiency. Our time-stepping approach allows for a matrix-free application. This can be implemented using any code equipped with linear direct and adjoint capabilities, bypassing the explicit formation of \mathbf{T} [24, 41]. Finally, the time stepping within our algorithm can also be used along with other SVD algorithms, *e.g.*, Arnoldi and power iteration, if desired.

In turbulent flows, where the state dimension N can become very large due to higher resolution requirements, CPU cost and memory requirements can swiftly impose constraints on the applicability of RSVD-LU. This limitation applies even to steady problems and is further exacerbated for harmonic resolvent analysis due to the inflated frequency-domain operators necessitated by triadic frequency coupling. Our algorithm circumvents this issue by strategically operating in the time domain, avoiding the need for limiting operations like LU decomposition, leading to linear cost scaling. Because of this key distinction, the RSVD- Δt could enable harmonic resolvent analysis of previous intractable turbulent flows.

Declarations

- Funding: This work was funded in part by Air Force Office of Scientific Research (AFOSR) grant no. FA9550-20-1-0214 and by Catalyst Grant from the Michigan Institute for Computational Discovery and Engineering (MICDE).
- Acknowledgements: We acknowledge the University of Michigan’s Great Lakes cluster for providing the essential computational resources to conduct all the simulations in this study.
- Competing interests: The authors have no relevant financial or non-financial interests to disclose.
- Availability of data and materials: A subset of the data is available from the author upon reasonable request
- Authors’ contributions: Conceptualization: Aaron Towne, Ali Farghadan; Methodology: Aaron Towne, Ali Farghadan; Formal analysis and investigation: Ali Farghadan; Writing - original draft preparation: Ali Farghadan; Writing - review and editing: All authors; Software: Ali Farghadan, Rutvij Bhagwat; Validation: Junoh Jung; Funding acquisition: Aaron Towne; Supervision: Aaron Towne
- Consent for publication: Not applicable
- Ethics approval and consent to participate: Not applicable

Appendix A RSVD- Δt for the subharmonic resolvent operator

In §3.5, we provided an overview of subharmonic resolvent analysis. In this appendix, we briefly outline the application of RSVD- Δt in computing subharmonic resolvent modes and gains. Consider the frequency of interest as $\gamma \in \Omega_\gamma$. This set specifies the perturbation frequency, while the base flow frequency content is confined to $\Omega_{\bar{q}}$.

To compute the actions of \mathbf{H} and \mathbf{H}^* using time-stepping, we must adhere to both the base flow frequency, enforcing a duration of $T = 2\pi/\omega_f$, and the perturbation frequency, enforcing a duration of $T_p = 2\pi/\gamma$, in order to obtain N_ω steady-state solutions to (7). Thus, we need to integrate for T_{sub} such that $T_{sub}/T \in \mathbb{N}$ and $T_{sub}/T_p \in \mathbb{N}$, during which N_ω equidistant snapshots are saved. For instance, if we consider $\gamma = \omega_f/5$, requiring $T_p = 2\pi/(\omega_f/5) = 5T$, integrating for $T_{sub} = 5T$ is sufficient to obtain the steady-state solutions. All the other steps remain the same as introduced in Algorithm 1.

References

- [1] A. Padovan, S. E. Otto, and C. W. Rowley. Analysis of amplification mechanisms and cross-frequency interactions in nonlinear flows via the harmonic resolvent. *Journal of Fluid Mechanics*, 900, 2020.
- [2] K. Taira, S. L. Brunton, S. T. M. Dawson, C. W. Rowley, T. Colonius, B. J. McKeon, O. T. Schmidt, S. Gordeyev, V. Theofilis, and L. S. Ukeiley. Modal analysis of fluid flows: An overview. *AIAA Journal*, 55(12):4013–4041, 2017.
- [3] A. Towne, O. T. Schmidt, and T. Colonius. Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis. *Journal of Fluid Mechanics*, 847:821–867, 2018.
- [4] S. T. M. Dawson and B. J. McKeon. On the shape of resolvent modes in wall-bounded turbulence. *Journal of Fluid Mechanics*, 877:682–716, 2019.
- [5] D. Sipp and O. Marquet. Characterization of noise amplifiers with global singular modes: the case of the leading-edge flat-plate boundary layer. *Theoretical and Computational Fluid Dynamics*, 27(5):617–635, 2013.
- [6] E. Pickering, A. Towne, P. Jordan, and T. Colonius. Resolvent-based modeling of turbulent jet noise. *The Journal of the Acoustical Society of America*, 150(4):2421–2433, 2021.
- [7] C. D. Troy and J. R. Koseff. The instability and breaking of long internal waves. *Journal of Fluid Mechanics*, 543:107–136, 2005.
- [8] A. E. Giannenas, S. Laizet, and G. Rigas. Harmonic forcing of a laminar bluff body wake with rear pitching flaps. *Journal of Fluid Mechanics*, 945:A5, 2022.
- [9] A. Farghadan and A. Arzani. The combined effect of wall shear stress topology and magnitude on cardiovascular mass transport. *International Journal of Heat and Mass Transfer*, 131:252–260, 2019.
- [10] N. M. Wereley. *Analysis and control of linear periodically time varying systems*. PhD thesis, Massachusetts Institute of Technology, 1990.
- [11] C.-T. Lin, M.-L. Tsai, and H.-C. Tsai. Flow control of a plunging cylinder based on resolvent analysis. *Journal of Fluid Mechanics*, 967:A41, 2023.
- [12] L. Franceschini, D. Sipp, O. Marquet, J. Moulin, and J. Dandois. Identification and reconstruction of high-frequency fluctuations evolving on a low-frequency periodic limit cycle: application to turbulent cylinder flow. *Journal of Fluid Mechanics*, 942:A28, 2022.
- [13] L. Heidt and T. Colonius. Spectral proper orthogonal decomposition of harmonically forced turbulent flows. *arXiv preprint arXiv:2305.05628*, 2023.
- [14] W. Wu, C. Meneveau, R. Mittal, A. Padovan, C. W. Rowley, and L. Cattafesta. Response of a turbulent separation bubble to zero-net-mass-flux jet perturbations. *Physical Review Fluids*, 7(8):084601, 2022.

- [15] A. Padovan and C. W. Rowley. Analysis of the dynamics of subharmonic flow structures via the harmonic resolvent: Application to vortex pairing in an axisymmetric jet. *Physical Review Fluids*, 7(7):073903, 2022.
- [16] B. Jin, S. Symon, and S. J. Illingworth. Energy transfer mechanisms and resolvent analysis in the cylinder wake. *Physical Review Fluids*, 6(2):024702, 2021.
- [17] T. Herbert. Analysis of the subharmonic route to transition in boundary layers. In *22nd Aerospace Sciences Meeting*, page 9, 1984.
- [18] D. Edgington-Mitchell, T. Wang, P. Nogueira, O. Schmidt, V. Jaunet, D. Duke, P. Jordan, and A. Towne. Waves in screeching jets. *Journal of Fluid Mechanics*, 913:A7, 2021.
- [19] C. Kumar and A. Prakash. Effect of mass injection on secondary instability of hypersonic boundary layer over a blunt cone. *Physics of Fluids*, 34(6), 2022.
- [20] R. Moarref, A. S. Sharma, J. A. Tropp, and B. J. McKeon. Model-based scaling of the streamwise energy density in high-Reynolds-number turbulent channels. *Journal of Fluid Mechanics*, 734:275–316, 2013.
- [21] J. H. M. Ribeiro, C.-A. Yeh, and K. Taira. Randomized resolvent analysis. *Physical Review Fluids*, 5(3):033902, 2020.
- [22] A. Farghadan, E. Martini, and A. Towne. Scalable resolvent analysis for three-dimensional flows. *arXiv preprint arXiv:2309.04617*, 2023.
- [23] A. Monokrousos, E. Åkervik, L. Brandt, and D. S. Henningson. Global three-dimensional optimal disturbances in the blasius boundary-layer flow using time-steppers. *Journal of Fluid Mechanics*, 650:181–214, 2010.
- [24] E. Martini, D. Rodríguez, A. Towne, and A. V. G. Cavalieri. Efficient computation of global resolvent modes. *Journal of Fluid Mechanics*, 919, 2021.
- [25] B. J. McKeon and A. S. Sharma. A critical-layer framework for turbulent pipe flow. *Journal of Fluid Mechanics*, 658:336–382, 2010.
- [26] C. Leclercq and D. Sipp. Mean resolvent operator of a statistically steady flow. *Journal of Fluid Mechanics*, 968:A13, 2023.
- [27] N. Halko, P. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.
- [28] H. Nyquist. Certain topics in telegraph transmission theory. *Transactions of the American Institute of Electrical Engineers*, 47(2):617–644, 1928.
- [29] A. Towne, G. Rigas, O. Kamal, E. Pickering, and T. Colonius. Efficient global resolvent analysis via the one-way Navier-Stokes equations. *Journal of Fluid Mechanics*, 948:A9, 2022.
- [30] G. Rigas, D. Sipp, and T. Colonius. Nonlinear input/output analysis: application to boundary layer transition. *Journal of Fluid Mechanics*, 911:A15, 2021.
- [31] J. M. Chomaz, P. Huerre, and L. G. Redekopp. Bifurcations to local and global modes in spatially developing flows. *Physical Review Letters*, 60(1):25, 1988.
- [32] R. E. Hunt and D. G. Crighton. Instability of flows in spatially developing media. *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences*, 435(1893):109–128, 1991.

- [33] S. Bagheri, D. S. Henningson, J. Hoepffner, and P. J. Schmid. Input-output analysis and control design applied to a linear model of spatially developing flows. *Applied Mechanics Reviews*, 62(2), 2009.
- [34] K. K. Chen and C. W. Rowley. H2 optimal actuator and sensor placement in the linearised complex Ginzburg–Landau system. *Journal of Fluid Mechanics*, 681:241–260, 2011.
- [35] A. V. G. Cavalieri, P. Jordan, and L. Lesshafft. Wave-packet models for jet dynamics and sound radiation. *Applied Mechanics Reviews*, 71(2), 2019.
- [36] J. Jung, R. Bhagwat, and A. Towne. Resolvent-based estimation of laminar flow around an airfoil. *AIAA Paper #2023-0077*, 2023.
- [37] B.-T. Chu. On the energy transfer to small disturbances in fluid flow (part i). *Acta Mechanica*, 1(3): 215–234, 1965.
- [38] O. T. Schmidt, A. Towne, G. Rigas, T. Colonius, and G. A. Brès. Spectral analysis of jet turbulence. *Journal of Fluid Mechanics*, 855:953–982, 2018.
- [39] S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. Gropp, et al. *PETSc users manual*. Argonne National Laboratory, 2019.
- [40] V. Hernandez, J. E. Roman, and V. Vidal. Slepnc: A scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Transactions on Mathematical Software (TOMS)*, 31(3):351–362, 2005.
- [41] M. F. de Pando, D. Sipp, and P. J. Schmid. Efficient evaluation of the direct and adjoint linearized dynamics from compressible flow solvers. *Journal of Computational Physics*, 231(23):7739–7755, 2012.