Hybrid Chebyshev function bases for sparse spectral methods in parity-mixed PDEs on an infinite domain

Benjamin Miquel *, Keith Julien

Department of Applied Mathematics, University of Colorado, Boulder, CO 80309, USA

A B S T R A C T

We present a numerical spectral method to solve systems of differential equations on an infinite interval $y \in (-\infty, \infty)$ in presence of linear differential operators of the form $Q(y) \left(\partial / \partial_y \right)^b$ (where $Q(y)$ is a rational fraction and b a positive integer). Even when these operators are not parity-preserving, we demonstrate how a mixed expansion in interleaved Chebyshev rational functions $TB_n(y)$ and $SB_n(y)$ preserves the sparsity of their discretization. This paves the way for fast $O(N \ln N)$ and spectrally accurate mixed implicit-explicit time-marching of sets of linear and nonlinear equations in unbounded geometries.

© 2017 Elsevier Inc. All rights reserved.

1. Introduction

Different numerical strategies (see [1] for a review) are available to solve differential equations on an infinite interval $y \in (-\infty, \infty)$. The first subcategory of methods consists in artificially introducing a bounded interval, the width of which is large compared to the extension of the exact solution. Over this finite interval, the classic weaponry of finite differences methods or spectral methods [1–3] may be deployed. Other approaches rely on a representation of the solution using a basis of functions of infinite support. Usual candidates contain the Hermite functions, the Whittaker sinc function, or rational Chebyshev functions. In this paper, the focus is on the latter rational Chebyshev functions since fast Fourier transform algorithms can be advantageously utilized to compute the direct and inverse transforms from physical to spectral space.

Fast transform algorithms render amenable mixed implicit-explicit (IMEX) time-marching schemes for coupled systems of differential equations

$$\mathcal{M} \partial_t \boldsymbol{\psi} = \mathcal{L} \boldsymbol{\psi} + \mathcal{N}(\boldsymbol{\psi}) \tag{1}$$

with linear operators $\mathcal{M}, \mathcal{L}$ and where a popular choice is to treat pseudo-spectrally the nonlinear terms $\mathcal{N}$. Often, the system of differential equations are accompanied with a set of M_b linear boundary conditions, labeled by the integer $m \in [1, M_b]$:

$$\mathcal{B}_m \boldsymbol{\psi}(\boldsymbol{r}) = b_m(\boldsymbol{r}) \quad \text{for } \boldsymbol{r} \in \partial \Omega. \tag{2}$$

The stability and accuracy of implicit time-stepping schemes render them very desirable over explicit schemes, but comes at the cost of one or several linear algebraic solves of the form $\mathbf{Ax} = \mathbf{b}$ for each time-step. As the resolution N increases,

---

* Corresponding author.
  E-mail addresses: benjamin.miquel@colorado.edu (B. Miquel), julien@colorado.edu (K. Julien).

such solves become prohibitively costly and necessitate $O(N^3)$ operations if the chosen discretization results in a dense representation of $\mathscr{L}$. An implicit treatment of the linear terms, as in the popular IMEX schemes, hence advocates for a spectral decomposition that yields a sparse representation of both the operators $\mathscr{L}$ and $\mathscr{M}$. Additionally, solving the corresponding generalized eigenproblem $\lambda \mathscr{M} \boldsymbol{\psi} = \mathscr{L} \boldsymbol{\psi}$ and (2) in order to analyze the linear dynamics also benefit from sparse, and ideally well-conditioned, discretized operators.

In the case of a scalar differential equation, Boyd [4] has shown that the rational Chebyshev functions potentially provide a banded representation for operators $\mathscr{L}$ composed of multiplications by monomials and differential operators $y^a (\partial/\partial y)^b$. However, sparsity relies on the parity of $\mathscr{L}$, and is usually ruined if $\mathscr{L}$ is composed of operators of different parities.

In this paper, we introduce hybrid Chebyshev expansions based upon the definition of two families of function $TS_n(y)$ and $ST_n(y)$. These are obtained by interweaving Chebyshev rational functions $TB_n(y)$ and $SB_n(y)$. Upon leveraging these hybrid expansions, we expose a method that yields sparse representations for rational differential operators, even in cases where parity mixing would destroy the sparsity with regular Chebyshev expansions. Following the spirit of this approach for a single scalar equation, a recipe for obtaining a sparse discretized system for sets of coupled differential equations is also exposed.

For the sake of completeness and clarity, the authors have elected to use a pedagogical tone. As such the first sections will be familiar to seasoned users of spectral methods. The complexity of the examples developed is gradually ramped up throughout the paper. The role of these examples is hence to illustrate the mechanics of our expansions in different situations, so that to provide the reader with models that he can adapt to his own problem.

The paper is organized as follows. A brief summary of spectral methods and discretization is presented in section 2. The rational Chebyshev functions and hybrid expansions are presented in section 3. Their convergence properties are established and their utility for producing sparse discretized system is illustrated on a fundamental textbook example: the quantum harmonic oscillator. Section 4 is dedicated to the resolution of scalar equations with mixed-parity operator, through the example of the anharmonic oscillator. The method is extended to sets of coupled equations in Section 5, through the geophysical fluid dynamics example of the equatorial β-plane equations. Section 6 is dedicated to the generalization to multidimensional domains. The applicability of the method to time-stepping schemes for nonlinear equations is described in section 7 and illustrated on the example of the Kelvin–Helmholtz instability in shear flows. Finally, section 8 is composed of concluding remarks. Throughout, details of numerical approach are relegated to the appendices.

2. From equations to discrete representations

Three classes of spatial discretization methods are generally available for governing equations (1) and (2). The venerable finite differences method, first introduced in 1910 by Richardson [5] evaluates derivatives on a mesh and can be extended to the finite volume method on an irregular mesh [6]. Finite spectral element methods [7] decompose Ω in sub-domains where the unknowns are approximated by a small sample of basis functions. Finally, spectral methods, of interest here, express the unknowns as truncated expansion of N orthogonal basis functions $\{\phi_k(\boldsymbol{r})\}_{1 \le k \le N}$ on the whole domain Ω [1–3]:

$$\boldsymbol{\psi}_i(\boldsymbol{r}) = \sum_{k=1}^{N} \widetilde{\mathsf{X}}_{i,k} \phi_k(\boldsymbol{r}). \tag{3}$$

Upon substitution of the truncated expansion above into the governing equations (1), projecting each individual equation onto the $\{\phi_j\}$ basis yields the following algebraic system:

$$\mathbf{M} \partial_t \widetilde{\mathsf{X}} = \mathbf{L} \widetilde{\mathsf{X}} + \mathsf{N}[\widetilde{\mathsf{X}}]. \tag{4}$$

In the context of spectral methods boundary conditions can be enforced following three routes. *Galerkin* methods use basis functions $\phi_k(\boldsymbol{r})$, possibly tailored by basis recombination, that intrinsically obey the boundary conditions on $\partial \Omega$, so that no additional step is required for the whole expansion to comply with boundary conditions. For basis functions that violate the boundary conditions, two approaches are available. The *Collocation* method explicitly solves the governing equations on the physical space, i.e. on the Gauss-Lobatto or Gauss–Chebyshev grids, and enforces the boundary condition on the boundary locations $\partial \Omega$. The *tau method* evaluates the boundary conditions in coefficient space and incorporates the resulting tau lines in the matrices $\mathbf{M}$ and $\mathbf{L}$ at the expense of higher order projections of the governing equations. A more detailed discussion is dedicated to boundary conditions in section 6.

The linear limit of the algebraic system (4) is obtained by setting $\mathsf{N}[\widetilde{\mathsf{X}}] = 0$. In the case of autonomous sets of equations, a generalized eigenvalue problem where solutions are sought in the form of normal modes $\widetilde{\mathsf{X}}(t) = \widetilde{\mathsf{X}}_a \exp(\lambda t)$ with growth rate λ can be solved to characterize the linear dynamics:

$$\lambda \mathbf{M} \widetilde{\mathsf{X}}_a = \mathbf{L} \widetilde{\mathsf{X}}_a. \tag{5}$$

When the full spectrum is sought after, QR/QZ methods are available in linear algebra libraries (e.g. LAPACK). Regardless of the sparsity of the matrices, such methods are very costly with an $O(N^3)$ operations count. Moreover, it is often the case that the growthrate and frequency of the most unstable mode only is of primary interest. If this is so, using sparse iterative methods such as Arnoldi's method will prove much cheaper and faster, typically of complexity $O(N)$ for sparse matrices.

Iterative methods are therefore appealing for continuation in parameters. Granted, these methods suffer a certain lack of robustness: they usually require to be initialized with a good guess and commonly miss eigenvalues. However, a possible approach (previously described in [1]) for a practical problem such as computing marginal stability curves or surfaces, or computing the optimal growthrate in a parameter space, is to first sweep through the parameter space along a coarse mesh and to compute the full spectra at a modest resolution using a QR/QZ algorithm. Once a first intuition is gained, a much finer sweeping can be carried out using iterative methods and a higher resolution (potentially ruling out under resolved modes, for instance).

An alternate approach suitable to the study of instabilities consists in time-stepping the linearized equation: one may then observe the growth of the most unstable mode. A large variety of time-stepping schemes are now available. Having applications in fluid mechanics in mind, the focus in this paper is on developing, for systems on an infinite line, spectral methods that are suitable for IMEX time-steppers. IMEX schemes evaluate the linear term in (4) implicitly at the end of the time-step, as opposed to evaluation at the start of the time-step for explicit time-stepping schemes. This results in a greater numerical stability and accuracy in comparison with their explicit counterparts, particularly desirable when stiff operators are present. As a pedagogical example, the simplest implicit time-stepper is the first order accurate Backward Euler scheme, which obtains the value of the unknown $\widetilde{\mathsf{X}}^{(n+1)}$ after a time-step τ by solving:

$$(\mathbf{M} - \tau \mathbf{L}) \widetilde{\mathsf{X}}^{(n+1)} = \mathbf{M} \widetilde{\mathsf{X}}^{(n)} , \tag{6}$$

where the value at the beginning of the time-step $\widetilde{\mathsf{X}}^{(n)}$ is known. Higher-order L-stable implicit schemes [8,9], and more recently memory-efficient schemes [10] have been developed. The performance bottleneck of such schemes is the complexity of the linear solves of the form given in equation (6). As in the case of eigenproblems discussed above, such linear algebraic systems are costly to solve when the matrix is dense: the complexity of the solve algorithm is of order $O(N^3)$. In strong contrast, for a discretization that yields a banded operator $(\mathbf{M} - \tau \mathbf{L})$, this complexity is drastically reduced to $O(N)$. From this observation, a successful implementation of an implicit algorithm is conditional on the existence of sparse matrix representations $\mathbf{M}$ and $\mathbf{L}$ for operators $\mathscr{M}$ and $\mathscr{L}$.

Beyond the linear dynamics, one may investigate the fully nonlinear dynamics, for instance in view of understanding the saturation of an instability. To avoid a costly direct computation of the nonlinear term in spectral space by means of a convolution product, the pseudo-spectral approach is utilized where potential derivatives are first evaluated in spectral space and products are then computed in physical space. The complexity of this procedure is dominated by the complexity of the transforms required between the physical and spectral space. When implemented as matrix multiplications, these transforms usually necessitate $O(N^2)$ operations. This is largely the reason why spectral methods did not gain popularity until a fast Fourier transform (FFT) algorithm was published in 1965 by Cooley and Tukey [11], revisiting an idea of Gauss [12]. The FFT algorithm implements transforms in $O(N \ln N)$ operations only.

The great strength of spectral methods is their accuracy: the truncated expansion employed in spectral methods converges exponentially towards the true solution, when analytic, as opposed to algebraically with finite differences. Further, spectral methods have negligible spurious numerical dissipation, a well known feature of finite differences methods. A conservative numerical method is highly desirable for weakly dissipative equations, such as high Reynolds number fluid dynamics problems. The accuracy of spectral method comes at a price. First, they suffer from a poor flexibility concerning the geometry of the domain. Regular domains (e.g. rectangles, parallelepipeds, spheres, cylinders, *etc.*) are good candidates for a spectral treatment, but spectral methods are ill-suited for irregular geometries. Second, we already mentioned the necessity of transforms: this point strongly advocates for Chebyshev or Fourier basis, which both utilize fast Fourier transform algorithms. Third, the discretized representation of the operators are often dense.

As a conclusion of these observations, the emphasis of this paper is on choices of rational Chebyshev function expansions for functions on the infinite line. The central idea here is that sensibly chosen hybrid expansions obtained by interleaving different bases preserve the sparsity of discretized operators and exploit existing FFT algorithms.

3. The rational Chebyshev functions

In this paper, we introduce the basis of rational Chebyshev functions as remapped cosine and sine functions on the infinite line. Not only does this approach lead to an insightful definition of the collocation points on the infinite line, but it also provides an understanding of which operators will lend themselves to a sparse discrete representation when approximated under the basis of rational Chebyshev functions. We thus begin by focusing on functions on the interval $[0, \pi]$ prior to the remapping.

3.1. Trigonometric basis functions on a finite interval

We consider $\mathscr{V}$, the Hilbert space of real valued continuous functions on the interval $[0, \pi]$, that are also piecewise of class $\mathscr{C}^1$ (i.e. with a continuous first derivative) and that vanishes at the end of the interval $f(0) = f(\pi) = 0$. We define the inner product:

$$\langle f, g \rangle = \frac{1}{\pi} \int_0^\pi f(\theta) g(\theta) \mathrm{d}\theta , \tag{7}$$
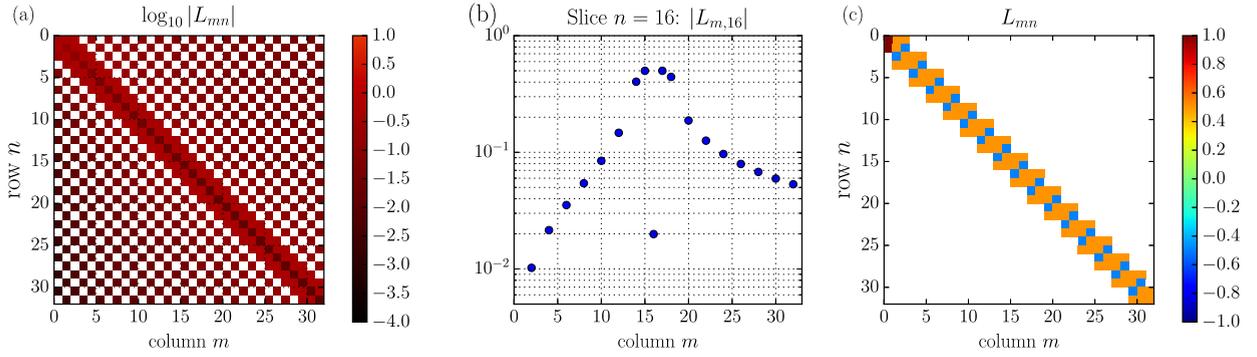
Fig. 1. Discretizations of $\mathscr{L} = \cos\theta + \sin\theta$ on a pure trigonometric series (figures a, b), or a mixed trigonometric series (figure c). Sparsity is illustrated by displaying the zero elements in white. Non-zero elements are color-coded. (a) Discretization on a sine basis $L_{mn} = \langle \sin(m\theta),\, \mathscr{L}\sin(n\theta)\rangle\, /\, \|\sin(m\theta)\|$ (see equation (13)). (b) Slice of matrix L_{mn} along row 16 (zero elements are ignored by the logarithmic scale), displaying a moderately rapid decay of the matrix elements away from the diagonal. (c) Matrix elements obtained with an expansion on $\breve{\beta}_n$ and a projection on $\mathring{\beta}_n$, namely $L_{mn} = \left\langle \mathring{\beta}_m(\theta),\, \mathscr{L}\breve{\beta}_n(\theta)\right\rangle / \|\mathring{\beta}_m(\theta)\|$.

and the corresponding 2-norm $\|f\| = \sqrt{\langle f,\, f\rangle}$. The families of functions $\{\cos(n\theta)\}_{n\in\mathbb{N}}$, indexed by nonnegative integers, and $\{\sin(n\theta)\}_{n\in\mathbb{N}^*}$, indexed by positive integers $\mathbb{N}^*$, each form an orthogonal and countable basis of $\mathscr{V}$. Hence a function $f \in \mathscr{V}$ can be expressed as a cosine series:

$$f(\theta) = \sum_{n \geq 0} \widetilde{f}_n \cos(n\theta)\,, \tag{8}$$

with

$$\text{for } n = 0: \quad \widetilde{f}_0 = \frac{1}{\pi}\int_0^\pi f(\theta)\mathrm{d}\theta\,, \tag{9a}$$

$$\text{for } n > 0: \quad \widetilde{f}_n = \frac{\langle \cos(n\theta),\, f(\theta)\rangle}{\|\cos(n\theta)\|^2} = \frac{2}{\pi}\int_0^\pi \cos(n\theta) f(\theta)\mathrm{d}\theta\,, \tag{9b}$$

or alternatively as a sine series:

$$f(\theta) = \sum_{n \geq 1} \widehat{f}_n \sin(n\theta)\,, \tag{10}$$

with

$$\widehat{f}_n = \frac{\langle \sin(n\theta),\, f(\theta)\rangle}{\|\sin(n\theta)\|^2} = \frac{2}{\pi}\int_0^\pi \sin(n\theta) f(\theta)\mathrm{d}\theta\,. \tag{11}$$

For functions in $\mathscr{V}$, both expansions converge uniformly to the function f. This does not imply that the expansions are equivalent: a discussion of the convergence properties of different representations, based on the analyticity of the possible extensions of $f(\theta)$ to the interval $[-\pi, \pi]$ (e.g., see [1–3]), can be found in section 3.5 below.

3.2. Hybrid trigonometric bases

Pure cosine or sine bases (presented above in equations (8), (10)) suffer the curse of dense discretization for symmetry breaking operators. For instance, consider $\mathscr{L} = \cos\theta + \sin\theta$ as example of an operator with no particular symmetry around $\pi/2$. To avoid confusion with the symmetry around $\theta = 0$ that stems in our discussion about convergence (section 3.5 and Appendix A), we emphasize that the symmetry that pertains to the present discussion is the one around $\pi/2$, the center of the segment $[0, \pi]$. Regardless of the selection of the pure basis (8) or (10) for expansion and projection, we would invariably obtain dense representations:

$$\frac{\langle \cos(m\theta),\, \mathscr{L}\cos(n\theta)\rangle}{\|\cos(m\theta)\|}\,, \quad \frac{\langle \sin(m\theta),\, \mathscr{L}\cos(n\theta)\rangle}{\|\sin(m\theta)\|}\,, \quad \frac{\langle \cos(m\theta),\, \mathscr{L}\sin(n\theta)\rangle}{\|\cos(m\theta)\|}\,, \text{ or } \frac{\langle \sin(m\theta),\, \mathscr{L}\sin(n\theta)\rangle}{\|\sin(m\theta)\|}\,, \tag{12}$$

all of which are band-unlimited. For example, the latter element of this list, represented on Fig. 1, yields:

$$\frac{\langle \sin(m\theta), \mathscr{L}\sin(n\theta)\rangle}{\|\sin(m\theta)\|} = \frac{1}{2}\left(\delta_{n+1}^m + \delta_{n-1}^m\right) + \frac{m\left(1 + (-1)^{m+n}\right)}{\pi(m-n+1)(m+n-1)} + \frac{m\left(1 + (-1)^{m+n}\right)}{\pi(m-n-1)(m+n+1)}. \tag{13}$$

To circumvent the apparition of dense matrices, the strategy is to introduce mixed trigonometric bases [13], described below. The core idea to this paper then consists in generating the corresponding mixed Chebyshev rational function bases when working on an unbounded interval.

The continuous functions $f \in \mathscr{V}$ may be uniquely decomposed into symmetric and antisymmetric components about the center of the segment $[0, \pi]$, namely $\theta = \pi/2$:

$$f(\theta) = f^{(s)}(\theta) + f^{(a)}(\theta), \tag{14a}$$

where

$$f^{(s)}(\theta) = f^{(s)}(\pi - \theta) \quad \text{and} \quad f^{(a)}(\theta) = -f^{(a)}(\pi - \theta). \tag{14b}$$

In contrast to the "pure" expansion in sine and cosine presented above in equations (8), (10), we define the two "hybrid" (or "mixed") cosine-sine expansions:

$$f(\theta) = f^{(s)}(\theta) + f^{(a)}(\theta) \approx \sum_{p=0}^{N} \widetilde{f}_{2p}\cos(2p\theta) + \sum_{p=0}^{N} \widehat{f}_{2p}\sin(2p\theta), \tag{15a}$$

or:

$$f(\theta) = f^{(s)}(\theta) + f^{(a)}(\theta) \approx \sum_{p=0}^{N} \widehat{f}_{2p+1}\sin([2p+1]\theta) + \sum_{p=0}^{N} \widetilde{f}_{2p+1}\cos([2p+1]\theta). \tag{15b}$$

As deduced from these mixed representations, each component $f^{(s)}$ and $f^{(a)}$ has two possible basis representations. For instance, $\{\cos(2p\theta)\}_{p\in\mathbb{N}}$ and $\{\sin([2p+1]\theta)\}_{p\in\mathbb{N}}$ each form an independent basis for symmetric functions about $\pi/2$, whereas $\{\cos([2p+1]\theta)\}_{p\in\mathbb{N}}$ and $\{\sin(2p\theta)\}_{p\in\mathbb{N}^*}$ form an independent basis for antisymmetric functions about $\pi/2$. Since this paper leverages the virtue of mixed expansions (15a), (15b) and makes a heavy use of them in producing sparse discretizations, each of these composite expansions is seen hereafter as a unique expansion by introducing the following compact notations:

$$f(\theta) \approx \sum_{n=1}^{N} \check{f}_n\check{\beta}_n(\theta) \quad \text{with} \quad \left\{\check{f}_n\right\}_{n\geq 1} = \left\{\widetilde{f}_0, \widehat{f}_2, \widetilde{f}_2, \widehat{f}_4, \widetilde{f}_4, \cdots\right\}, \tag{16a}$$

or:

$$f(\theta) \approx \sum_{n=1}^{N} \mathring{f}_n\mathring{\beta}_n(\theta) \quad \text{with} \quad \left\{\mathring{f}_n\right\}_{n\geq 1} = \left\{\widetilde{f}_1, \widehat{f}_1, \widetilde{f}_3, \widehat{f}_3, \widetilde{f}_5, \widehat{f}_5, \cdots\right\}. \tag{16b}$$

The two distinct hybrid bases obtained, $\{\check{\beta}_n\}$ and $\{\mathring{\beta}_n\}$, can be thought of as interleaved collections of sines and cosines:

$$\check{\beta}_n(\theta) = \begin{cases} \cos(2p\theta) & \text{for } n = 2p+1, \\ \sin(2p\theta) & \text{for } n = 2p, \end{cases} \tag{17a}$$

$$\mathring{\beta}_n(\theta) = \begin{cases} \cos([2p-1]\theta) & \text{for } n = 2p-1, \\ \sin([2p-1]\theta) & \text{for } n = 2p, \end{cases} \tag{17b}$$

or more explicitly:

$$\left\{\check{\beta}_n(\theta)\right\}_{n\geq 1} = \left\{1, \sin 2\theta, \cos 2\theta, \sin 4\theta, \cos 4\theta, \cdots\right\}, \tag{18a}$$

$$\left\{\mathring{\beta}_n(\theta)\right\}_{n\geq 1} = \left\{\cos\theta, \sin\theta, \cos 3\theta, \sin 3\theta, \cos 5\theta, \sin 5\theta, \cdots\right\}. \tag{18b}$$

Both bases are orthogonal sets, as proven in Appendix B. The benefits of using such functions become evident by revisiting the introductory example of this paragraph, $\mathscr{L} = \cos\theta + \sin\theta$. From the trigonometric identities for products, we deduce easily that the following matrix representations for $\mathscr{L}$ (one of which is represented on Fig. 1) are both sparse:

$$\frac{\left\langle \mathring{\beta}_m(\theta), \mathscr{L}\check{\beta}_n(\theta)\right\rangle}{\|\mathring{\beta}_m(\theta)\|}, \quad \frac{\left\langle \check{\beta}_m(\theta), \mathscr{L}\mathring{\beta}_n(\theta)\right\rangle}{\|\check{\beta}_m(\theta)\|}. \tag{19}$$

Therefore, a sparse discretization for $\mathscr{L}$ is obtained by considering an expansion in $\{\check{\beta}_n\}$ and a projection on the $\{\mathring{\beta}_n\}$ basis (or vice versa). This observation can be generalized in the following proposition.

Proposition 1. *Let $\mathscr{L}$ be a linear operator of the form*

$$\mathscr{L} = \sum_{i=1}^{N_{\mathscr{L}}} (\alpha_i \cos(n_i\theta) + \beta_i \sin(m_i\theta)) \left(\frac{\mathrm{d}}{\mathrm{d}\theta}\right)^{\gamma_i} \tag{20}$$

where $N_{\mathscr{L}} \in \mathbb{N}$, and for all integers $1 < i < N_{\mathscr{L}}$: $\alpha_i, \beta_i \in \mathbb{R}$, $\gamma_i \in \mathbb{N}$, and $(n_i, m_i) \in \mathbb{N}^2$ such that all integers n_i and m_i are of the same parity (even or odd). If all these integers are:

- *odd, then $\dfrac{\left\langle \mathring{\beta}_m(\theta), \mathscr{L}\check{\beta}_n(\theta) \right\rangle}{\left\| \mathring{\beta}_m(\theta) \right\|}$ and $\dfrac{\left\langle \check{\beta}_m(\theta), \mathscr{L}\mathring{\beta}_n(\theta) \right\rangle}{\left\| \check{\beta}_m(\theta) \right\|}$ are band-limited matrices,*

- *even, then $\dfrac{\left\langle \check{\beta}_m(\theta), \mathscr{L}\check{\beta}_n(\theta) \right\rangle}{\left\| \check{\beta}_m(\theta) \right\|}$ and $\dfrac{\left\langle \mathring{\beta}_m(\theta), \mathscr{L}\mathring{\beta}_n(\theta) \right\rangle}{\left\| \mathring{\beta}_m(\theta) \right\|}$ are band-limited matrices.*

Proof. These matrices are easily built as linear combinations of matrix product of the banded elementary matrices $\mathcal{S}$, $\mathcal{C}$, and $\mathcal{D}$ given in appendix C.1. □

The requirement that the n_i and m_i be of given parity will be naturally satisfied by operators originating from a remapping of rational operators on the infinite line, as discussed in the following sections.

For completeness, we close up this section by mentioning alternative expansions: analogous to the $\{\check{\beta}_n\}_{n\geq 1}$ basis is the $\{\exp(i[2p+1]\theta)\}_{p\in\mathbb{Z}}$ basis. Similarly, the $\{\exp(i[2p+1]\theta)\}_{p\in\mathbb{Z}}$ basis is analogous to the $\{\mathring{\beta}_n\}_{n\geq 1}$ basis. The changes of basis are straightforward, so that convergence properties (discussed below in section 3.5) will be unaffected and choosing between the trigonometric and the complex exponential bases is purely a matter of taste. In the rest of this paper, we discuss and utilize the hybrid trigonometric expansions $\{\check{\beta}\}$ and $\{\mathring{\beta}\}$ only.

3.3. Rational Chebyshev basis functions on the infinite line

Given the mixed representations on the $[0, \pi]$ interval, following Cain et al. [14] we now map this interval onto the infinite line using:

$$y = L \cot\theta, \tag{21a}$$

$$\theta = \mathrm{arccot}\,(y/L), \tag{21b}$$

where $L > 0$ is the mapping parameter. Derivatives on the two intervals are connected according to

$$\partial_y = -\frac{\sin^2\theta}{L^2}\partial_\theta. \tag{22}$$

By use of composite functions, a function on the infinite line $\psi(y)$ can be bijectively mapped on a function $f(\theta)$ on the segment $\theta \in [0, \pi]$:

$$\psi(y) = \psi(y(\theta)) = f(\theta). \tag{23}$$

Hence, we define the Chebyshev functions $TB_n(y)$ and $SB_n(y)$ as the mapped cosine and sine functions stretched on the infinite line:

$$TB_n(y) = \cos(n\,\mathrm{arccot}\,(y/L)), \tag{24a}$$

$$SB_n(y) = \sin(n\,\mathrm{arccot}\,(y/L)). \tag{24b}$$

From their definition and our discussion above, the Chebyshev functions inherit the properties of the trigonometric functions from which they are derived: they form two independent orthogonal bases on the infinite line and a truncated sum can approximate an arbitrary continuous bounded function:

$$\psi(y) \approx \sum_{n=0}^{N} \widetilde{\psi}_n TB_n(y), \tag{25a}$$

$$\psi(y) \approx \sum_{n=1}^{N} \widehat{\psi}_n SB_n(y). \tag{25b}$$

These expansions are the infinite line counterparts of the cosine and sine expansion given in equations (8), (10). The TB_n (respectively, SB_n) coefficients $\widetilde{\psi}_n$ (respectively, $\widehat{\psi}_n$) are equal to the cosine (respectively, sine) coefficients $\widetilde{f}_n$ (respectively, $\widehat{f}_n$). However, they can be computed directly on the infinite line, without mapping back to the interval $[0, \pi]$, by using

$$\int_0^\pi f(\theta)\mathrm{d}\theta = \frac{1}{L} \int_{-\infty}^{\infty} \frac{\psi(y)}{1+y^2}\mathrm{d}y \tag{26}$$

to generalize equations (9), (11) into:

$$\text{for } n = 0: \quad \widetilde{\psi}_0 = \frac{1}{\pi L} \int_{-\infty}^{\infty} \frac{\psi(y)}{1+y^2}\mathrm{d}y, \tag{27a}$$

$$\text{for } n > 0: \quad \widetilde{\psi}_n = \frac{2}{\pi L} \int_{-\infty}^{\infty} \frac{TB_n(y)\,\psi(y)}{1+y^2}\mathrm{d}y, \tag{27b}$$

$$\widehat{\psi}_n = \frac{2}{\pi L} \int_{-\infty}^{\infty} \frac{SB_n(y)\,\psi(y)}{1+y^2}\mathrm{d}y. \tag{27c}$$

3.4. Hybrid Chebyshev bases

By means of the mapping (21), our partition of symmetric and asymmetric functions around $\theta = \pi/2$ on the interval $[0,\pi]$ results in a partition of functions $\psi(y)$ into odd and even components $\psi^{(e)}(y)$ and $\psi^{(o)}(y)$ around $y = 0$. We emphasize here that in the following, symmetry properties are implicitly relative to $\theta = \pi/2$ for functions on the $[0,\pi]$ segment, and relative to $y = 0$ on the infinite line $(-\infty, \infty)$. The odd and even components of $\psi(y)$ can be expanded independently with either a TB_n or SB_n basis. Hence, the analogous extension of the hybrid sine-cosine expansions (eq. (15)) are the hybrid SB–TB expansions:

$$\psi(y) = \psi^{(e)}(y) + \psi^{(o)}(y) \approx \sum_{p=0}^{N/2-1} \widetilde{\psi}_{2p}\, TB_{2p-2}(y) + \sum_{p=1}^{N/2} \widehat{\psi}_{2p}\, SB_{2p}(y), \tag{28a}$$

$$\approx \sum_{p=0}^{N/2} \widehat{\psi}_{2p+1}\, SB_{2p+1}(y) + \sum_{p=0}^{N/2} \widetilde{\psi}_{2p+1}\, TB_{2p+1}(y). \tag{28b}$$

Following section 3.2, we introduce the following notations for the two bases $ST(y)$ and $TS(y)$:

$$TS_n(y) = \begin{cases} TB_{2p}(y) & \text{for } n = 2p+1, \\ SB_{2p}(y) & \text{for } n = 2p, \end{cases} \tag{29a}$$

$$ST_n(y) = \begin{cases} TB_{2p-1}(y) & \text{for } n = 2p-1, \\ SB_{2p-1}(y) & \text{for } n = 2p. \end{cases} \tag{29b}$$

More explicitly,

$$\left\{TS_n(y)\right\}_{n\geq 1} = \left\{1, SB_2(y), TB_2(y), SB_4(y), TB_4(y), \cdots\right\}, \tag{30a}$$

$$\left\{ST_n(\theta)\right\}_{n\geq 1} = \left\{TB_1(y), SB_1(y), TB_3(y), SB_3(y), TB_5(y), SB_5(y), \cdots\right\} \tag{30b}$$

such that the expansions (15a), (15b) become respectively:

$$\psi(y) \approx \sum_{n=1}^{N} \check{\psi}_n TS_n(y) \quad \text{with} \quad \left\{\check{\psi}_n\right\}_{n\geq 1} = \left\{\widetilde{\psi}_0, \widehat{\psi}_2, \widetilde{\psi}_2, \widehat{\psi}_4, \widetilde{\psi}_4, \cdots\right\}, \tag{31a}$$

$$\approx \sum_{n=1}^{N} \mathring{\psi}_n ST_n(y) \quad \text{with} \quad \left\{\mathring{\psi}_n\right\}_{n\geq 1} = \left\{\widetilde{\psi}_1, \widehat{\psi}_1, \widetilde{\psi}_3, \widehat{\psi}_3, \widetilde{\psi}_5, \widehat{\psi}_5, \cdots\right\}. \tag{31b}$$

In the remainder of the paper, we discuss how the sparsity of discretized systems is affected by the choice of this expansion. Particularly, we highlight cases when a mixed expansion is necessary to conserve sparsity, by extending our Proposition 1 that pertains to $\theta \in [0,\pi]$ onto the infinite line $y \in (-\infty, \infty)$ by means of the mapping $y = L\cot\theta$.

3.5. Choosing between $\{TB_n\}$, $\{SB_n\}$, $\{TS_n\}$, $\{ST_n\}$ expansions and alternatives

At this stage, it is necessary to discuss the global relevance of hybrid $\{TS_n\}$ and $\{ST_n\}$ expansions for a given function $\psi(y)$. Granted the fact that they prove useful in providing sparse discretization of parity-flipping operators, the question arises of under which conditions they inherit the convergence properties of the regular $\{TB_n\}$ and $\{SB_n\}$ expansions? The answer to this question is related to the analyticity of the continuations of the remapped function $f(\theta) = \psi(L \cot \theta)$ to the $[-\pi, \pi]$ interval. A detailed discussion is given in Appendix A, and boils down to the following conclusions.

Exponentially decaying functions that are dominated by $\exp(-b|y|^a)$ with $a, b > 0$ have subgeometrically converging coefficients for each of the four expansions. That is to say, hybrid expansions $\{TS_n\}$ and $\{ST_n\}$ will converge as fast as pure $\{TB_n\}$ and $\{SB_n\}$ expansions. For this reason, hybrid expansions should be given preference over the pure expansions in cases where they yield a sparse discretization, as discussed above and exemplified below. Therefore the question of using hybrid Chebyshev expansions instead of alternative methods such as expansions based on Hermite functions or Cardinal functions is really equivalent to the choice between regular Chebyshev expansions versus Hermite or Cardinal functions. Hence, we refer the reader to the abundant existing literature (a possible starting point is [1]). The bottom line is that in most cases these three options are equivalent and none really outperforms the others from the perspective of an optimal representation. In practice however, as stated in our introduction, it is our opinion that Chebyshev rational functions have an edge whenever a sparse representation is available. Indeed, Hermite functions prove to be inpractical in the context of nonlinear dynamics due to the absence of implementation of fast transforms in common linear algebra libraries–despite the existence of such algorithms [15,16]. Using the Cardinal functions collocation method does not require any transform but results in dense discretizations, which yield a much slower time-stepping or use of iterative eigensolvers.

3.6. The relevance of mapping before discretizing

As mentioned above, one could very well introduce an expansion of the forms given in equations (25) or (28) and discretize directly a system of the form:

$$\mathscr{L}_y \psi(y) = b(y). \tag{32}$$

However, more often than not, the path of least resistance to obtain sparse matrices consists in first mapping the equation to the $[0, \pi]$ segment before proceeding to the discretization. The reason behind this is that our intuition with the action of operators on the trigonometric functions is much more developed than our insight with the Chebyshev functions. As we exemplify in the next paragraph with the case of the quantum harmonic oscillator, tweaking the equations with a view to obtaining a sparse system is much easier with trigonometric functions, whereas dealing directly with Chebyshev functions remains obscure.

3.7. A fundamental example: the Quantum Harmonic Oscillator

We illustrate the mechanics of the Chebyshev decomposition on the classic example of the Quantum Harmonic Oscillator (QHO). The QHO is a fundamental example of a particle propagating in a potential $V(y)$. The wave function ψ and the eigenenergies E of the particle are the solutions to the Schrödinger equation:

$$E\psi(y) = \left(V(y) - \partial_{yy}\right)\psi(y). \tag{33}$$

The potential $V(y) = y^2$ is well-known to possess a point spectrum of exact analytic solutions, indexed by the integer $n \geq 0$:

$$E_n = 2n + 1, \quad \psi^{(n)}(y) = H_n(y) \exp\left(-y^2/2\right), \tag{34}$$

where H_n is the Hermite polynomial of degree n. As described above, we first solve this equation by mapping it to the segment $[0, \pi]$ using (21). Doing so yields the corresponding eigenproblem for $f(\theta) = \psi(y)$:

$$Ef(\theta) = -\frac{1}{L^2}\left(\sin^4\theta\,\partial_{\theta\theta} + 2\sin^3\theta\cos\theta\,\partial_\theta\right)f(\theta) + \frac{L^2\cos^2\theta}{\sin^2\theta}f(\theta). \tag{35}$$

In this form, the last term with a $\sin^2\theta$ denominator leads to a dense discretization regardless of which expansion is adopted for $f(\theta)$. Therefore, we regularize the equation by the multiplication $\sin^2\theta$ throughout. Doing so readily yields the generalized eigenproblem of the form

$$E\mathscr{M}f(\theta) = \mathscr{L}f(\theta), \tag{36a}$$

where
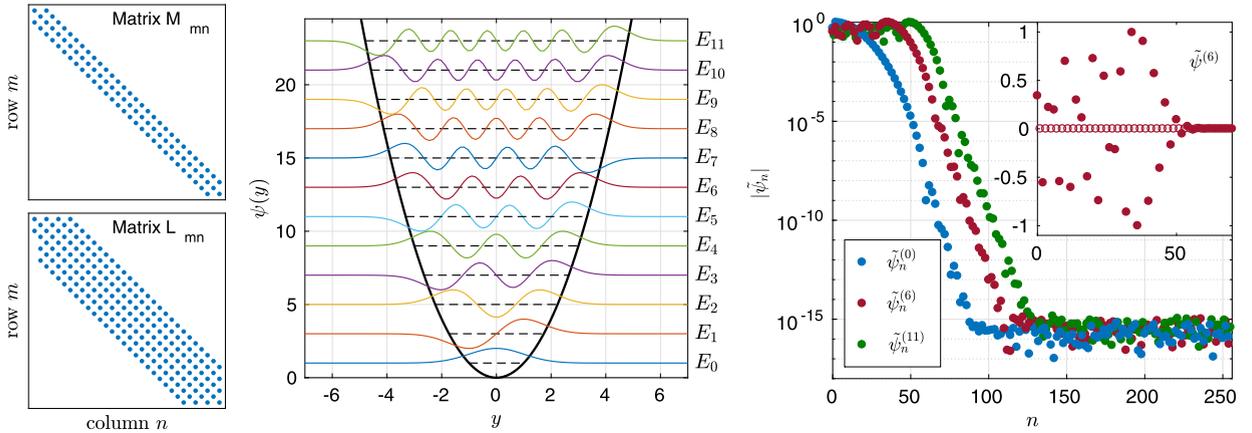
Fig. 2. Illustration of the Chebyshev spectral decomposition of the quantum harmonic oscillator using basis functions $TB_n(y)$. Leftmost panel: structure of the **M** (upper plot) and **L** (lower plot) matrices for 32 basis functions. Middle panel: the first 12 eigenmodes modes $\psi^n(y)$ of the parabolic well (continuous black parabola) obtained with 256 basis functions (solid color lines), shifted vertically by their energy level E_n (dashed horizontal lines). Right panel: Spectra $|\tilde{\psi}_n|$ for $n = 0, 6, 11$. Observe the exponential convergence of the expansion until machine precision is reached. Insert: Spectral coefficients of mode $\psi^{(6)}$. Solid circles are non-zero coefficients; open circles represent coefficients which are *rigorously* zero (*beyond* mere machine precision). The even symmetry of $\psi^{(6)}$ is preserved by our spectral method, as odd-degree coefficients $\tilde{\psi}_{2p+1}$ are exactly zero.

$$\mathcal{M} = \sin^2 \theta \,, \tag{36b}$$

$$\mathcal{L} = -\frac{1}{L^2} \left(\sin^6 \theta \partial_{\theta\theta} + 2 \sin^5 \theta \cos \theta \partial_\theta \right) + L^2 \cos^2 \theta \,. \tag{36c}$$

In this simple example, the discretization is a straightforward step: anticipating no particular difficulty, we adopt as a first guess a pure cosine expansion of N terms for $f(\theta)$. By projecting the equations onto the N basis functions $\{\cos(m\theta)\}_{0 \le m \le (N-1)}$, we obtain a $N \times N$ system of the form:

$$E \mathbf{M} \tilde{\boldsymbol{f}} = \mathbf{L} \tilde{\boldsymbol{f}} \quad \text{with} \quad \tilde{\boldsymbol{f}} = \left(\tilde{f}_0, \tilde{f}_1, \cdots, \tilde{f}_{N-1} \right)^{\mathrm{T}}, \tag{37}$$

where the matrix coefficients M_{mn} and L_{mn} are given by (for $1 \le m, n \le N$):

$$M_{mn} = \frac{\langle \cos([m-1]\theta), \, \mathcal{M} \cos([n-1]\theta) \rangle}{\| \cos([m-1]\theta) \|^2} \,, \tag{38a}$$

$$L_{mn} = \frac{\langle \cos([m-1]\theta), \, \mathcal{L} \cos([n-1]\theta) \rangle}{\| \cos([m-1]\theta) \|^2} \,. \tag{38b}$$

The exact expression of matrices **M** and **L** is omitted, but was obtained by using the procedure given in appendix C.1. Their banded form can be observed on Fig. 2. This sparse structure could have been easily anticipated, as the discretized operators **M** and **L** are composed exclusively by sparse representations on the cosines basis: i.e., multiplication by $\cos \theta$, $\cos^2(\theta)$, $\sin^2(\theta)$, second order derivation $\partial_{\theta\theta}$, and $\sin \theta \partial_\theta$. Note that an alternative choice consisting in expressing $f(\theta)$ as a sine expansion and projecting on the N functions $\{\sin(m\theta)\}_{1 \le m \le N}$ would yield a similarly sparse linear system $E \mathbf{M} \hat{\boldsymbol{f}} = \mathbf{L} \hat{\boldsymbol{f}}$ where:

$$M_{mn} = \frac{\langle \sin(m\theta), \, \mathcal{M} \sin(n\theta) \rangle}{\| \sin(m\theta) \|^2} \,, \tag{39a}$$

$$L_{mn} = \frac{\langle \sin(m\theta), \, \mathcal{L} \sin(n\theta) \rangle}{\| \sin(m\theta) \|^2} \,. \tag{39b}$$

In Fig. 3, we compare different methods used to solve the QHO equations, all of which contain a free length parameter L: the problem is solved on the infinite line using either a $TB_n(y)$ expansion (cosines), or a $SB_n(y)$ expansion. For comparison, we also solve our problem on a large segment $y \in [-L, L]$ by means of either a classic Chebyshev collocation method or finite differences. A Dirichlet boundary condition $\psi(L) = \psi(-L) = 0$ is imposed for the latter methods. In all cases, the number of grid points (finite differences and collocation method) or modes (spectral methods) is set to $N = 256$. The parameter L is increased between 0.25 and 20 by increments of 0.25. We display as a function of L the number $N_{<0.01}$ of eigenvalues correct within a relative error of 0.01 obtained by each method. For each method, an optimal value L_{opt} is found which corresponds to a optimal ratio $Q_{opt} = N_{<0.01}/256$ of accurate eigenvalues. The results are summarized in Table 1. From a sheer accuracy criteria, the SB_n and TB_n functions on the infinite line are sensibly preferable to Chebyshev polynomials. All three spectral methods are substantially more accurate than the method of finite differences. Finally, note on Fig. 3 that the poor conditioning of differentiation matrices on the Gauss–Lobatto grid (the collocation grid for T_n polynomials) is
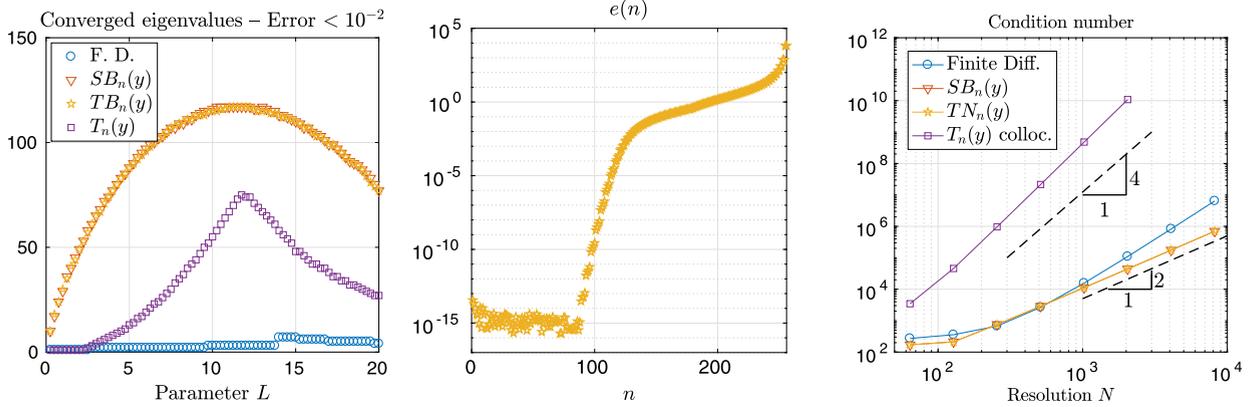
Fig. 3. Comparison of the different methods used to solved the QHO eigenproblem with $N = 256$ basis functions or grid points: spectral methods on the infinite line using Chebyshev functions $TB_n(y)$ (pentagrams) and $SB_n(y)$ (triangles) with the mapping parameter L; or methods on a finite segment $[-L, L]$: Finite differences (circles) or Chebyshev polynomial $T_n(y)$ collocation (squares). Left panel: number $N_{<0.01}$ of eigenvalues correct within a relative error of 0.01, as a function of L. In the remaining two panels, we set independently for each method $L = L_{opt}$, the optimal parameter that maximizes $N_{<0.01}$ in the left panel. Center panel: relative error $e(n) = \frac{E_n - (2n+1)}{2n+1}$ of the numerical eigenvalues E_n as a function of n. Right panel: Condition number of matrices **L**.

Table 1
Comparison of accuracy for the eigenmodes obtained by a spectral method on the infinite line (SB_n or TB_n), a Chebyshev collocation method, or first order finite differences with 256 modes or grid points. L_{opt} is the optimal mapping parameters for methods on the infinite line or the optimal size of the domain for the two other methods. $N_{<0.01}$ denotes the number of eigenvalues accurate within a relative error of 10^{-2}; $Q_{opt} = N_{<0.01}/256$ is the corresponding proportion of accurate eigenvalues.

Method	L_{opt}	$N_{<0.01}$	Q_{opt} (%)
$SB_n(y)$	11	132	51.6%
$TB_n(y)$	11	132	51.6%
Collocation $T_n(y)$	11.75	75	29.3%
Finite Differences	14	7	2.7%

visible, as the condition numbers increases as N^4. Discretizations based on rational Chebyshev functions do not suffer this phenomenon and their condition number scales quadratically with N.

We emphasize that in the present example as well as in the following examples, the computation of the full spectrum using a standard QR/QZ algorithm and its comparison to the analytic spectrum had been intended to provide a purely pedagogical illustration of the accuracy of the method. For real life applications, the sparsity of discretizations using Chebyshev functions should be leveraged through the use of iterative eigensolvers or sparse solvers.

4. Parity splitting operators in a scalar equation

4.1. Particle in a translated harmonic well: Chebyshev rational functions

We have illustrated how approximations using Chebyshev basis functions on the infinite line yields an efficient and accurate resolution of the QHO equation. However, the operators y^2 and ∂_{yy} are both even and hence parity-preserving. In this section, we first illustrate how the sparsity of the approximation method is ruined in presence of both even and odd parity operators. A resolution is then provided to restore sparsity for a large class of polynomials and rational differential operators.

Our introductory example is the case of a particle in a asymmetric parabolic potential $V_a(y) = y^2 - 6y$. The potential V_a is a mere translation of the harmonic potential accompanied by a renormalization of the ground state energy: applying the transformation $y = y' + 3$ and $E = E' - 9$, this problem can be mapped back to the QHO treated in the previous section. Hence we refer to this case as the "translated quantum harmonic oscillator" (TQHO) and we confront our approach against the known analytical solutions of this equation. Following a similar approach as above, we rewrite the Schrödinger equation for $\psi(y)$

$$E\psi(y) = \left(y^2 - 6y - \partial_{yy} \right) \psi(y), \tag{40}$$

in θ-space, and we regularize by multiplication by $\sin^2 \theta$, yielding:

$$E \sin^2 \theta f(\theta) = -\frac{1}{L^2} \left(\sin^6 \theta \partial_{\theta\theta} + 2 \sin^5 \theta \cos \theta \partial_\theta \right) f(\theta) + \left(L^2 \cos^2 \theta - 6L \sin \theta \cos \theta \right) f(\theta). \tag{41}$$
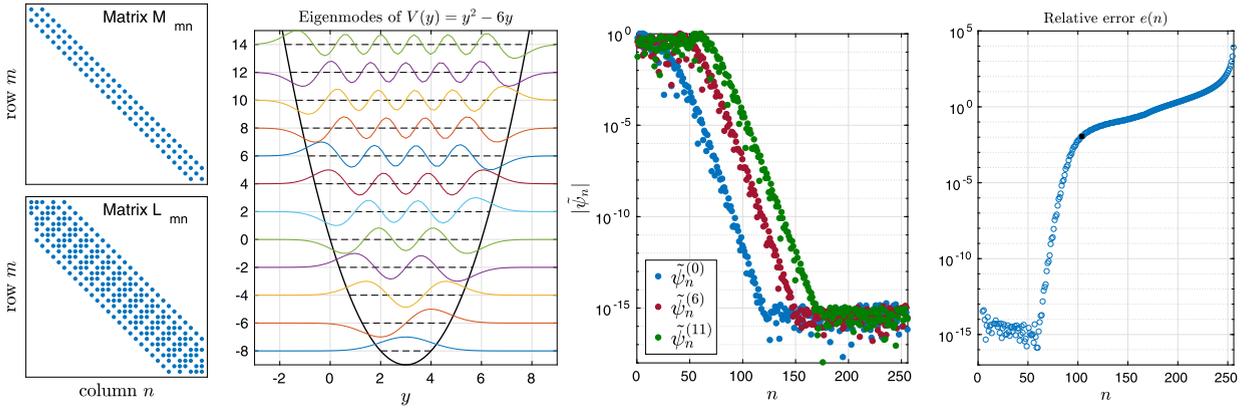
Fig. 4. From left to right: Sparse matrices **M** and **L** for the translated harmonic oscillator (equation (40)) for $N = 32$; Wave functions for $N = 256$, translated vertically by their energy (dashed horizontal lines); Spectra of modes 0, 6 and 11 illustrating exponential convergence; Relative error of the eigenenergies.

We identify on this simple example the trouble that accompanies parity flipping operators. Equation (40) contains two classes of operators: the identity operator on the left-hand-side and the operator $(y^2 - \partial_{yy})$ preserve the parity of the function they act on. By contrast, the operator $-6y$ reverses the parity. As a consequence, were ψ to be written as a "pure" TB or SB expansion, at least one term will be band-unlimited and the coupling matrix would be dense. Suppose we assume a pure cosine or sine expansion (8), (10) for f, as we did for the QHO. A discrete set of equations of the form $E\,\mathbf{M}\,\widetilde{f} = \mathbf{L}\,\widetilde{f}$ could be obtained by projection onto the cosine or sine basis. However, by contrast with the QHO and regardless of which one of the four options we elect, this system would never yield a sparse representation. Indeed, the two possible representations of $\mathscr{L}$:

$$\frac{\langle \cos(m\theta),\ \mathscr{L}\cos(n\theta)\rangle}{\|\cos(m\theta)\|^2},\ \frac{\langle \sin(m\theta),\ \mathscr{L}\sin(n\theta)\rangle}{\|\sin(m\theta)\|^2}, \tag{42}$$

are all dense matrices due the existence the parity-break term that results in a band-unlimited representation.

We note, however, that the differential operators present in equation (41) obey the requirement of our Proposition 1. Therefore, the sparsity can be restored if the odd and even components of ψ are expanded by a different basis.

A solution consists in assuming an hybrid expansion in sines and cosines of the form (15) for the symmetric and antisymmetric components $f^{(s)}$ and $f^{(a)}$. For instance, if we choose the basis $\{\breve{\beta}_n\}$ (i.e. a cosine and sine expansion (see equation (15a)) for $f^{(s)}$ and $f^{(a)}$, respectively), then we can compute the matrices representing the action of the operators that compose $\mathscr{L}$ on this basis:

$$\mathbf{A} = \frac{\left\langle \breve{\beta}_m,\ \sin^6\theta\,\partial_{\theta\theta}\breve{\beta}_n\right\rangle}{\left\|\breve{\beta}_m\right\|^2}, \tag{43a}$$

$$\mathbf{B} = \frac{\left\langle \breve{\beta}_m,\ \cos\theta\sin^5\theta\,\partial_\theta\breve{\beta}_n\right\rangle}{\left\|\breve{\beta}_m\right\|^2}, \tag{43b}$$

$$\mathbf{C} = \frac{\left\langle \breve{\beta}_m,\ \cos^2\theta\,\breve{\beta}_n\right\rangle}{\left\|\breve{\beta}_m\right\|^2}, \tag{43c}$$

$$\mathbf{D} = \frac{\left\langle \breve{\beta}_m,\ \sin^2\theta\,\breve{\beta}_n\right\rangle}{\left\|\breve{\beta}_m\right\|^2}, \tag{43d}$$

$$\mathbf{E} = \frac{\left\langle \breve{\beta}_m,\ \cos\theta\sin\theta\,\breve{\beta}_n\right\rangle}{\left\|\breve{\beta}_m\right\|^2}, \tag{43e}$$

all of which are sparse. With these building blocks in hand, the expression of which is given in appendix C.1, the discretized system:

$$E\,\mathbf{M}\,\breve{f} = \mathbf{L}\,\breve{f} \tag{44}$$

is readily obtained and the sparse structure of matrices **L** and **M** is illustrated in Fig. 4:

$$\mathbf{M} = \mathbf{D}, \tag{45a}$$

$$\mathbf{L} = -\frac{1}{L^2}\left(\mathbf{A} + 2\mathbf{B}\right) + L^2\mathbf{C} - 6L\mathbf{E}. \tag{45b}$$

As can also be observed in Fig. 4 when compared to Fig. 2, neither the exponential convergence of the series nor the excellent accuracy of the method are adversely affected by our using of a mixed basis. This illustrates the discussion of convergence given in section 3.5. For exponentially decaying solutions, expansions based upon the $\{\mathring{\beta}_n\}$ basis or the $\{\check{\beta}_n\}$ basis yield a similarly sparse representations and converge equally rapidly.

4.2. The generalized case

Having motivated, illustrated and benchmarked our method on a simple example, we now summarize and formalize the procedure to tackle a broader class of eigenproblems:

$$E\psi(y) = \mathcal{H}\psi(y), \tag{46}$$

where the operator $\mathcal{H}$ is a sum of rational fractions and derivatives of maximal order $N_{\mathcal{H}}$:

$$\mathcal{H} = \sum_{i=0}^{N_{\mathcal{H}}} \frac{P_i(y)}{Q_i(y)} \left(\frac{\mathrm{d}}{\mathrm{d}y}\right)^i, \tag{47}$$

and P_i and Q_i are polynomials. Note that even more general equations such as $E\mathcal{H}_1\psi = \mathcal{H}_2\psi$ (with $\mathcal{H}_1$ and $\mathcal{H}_2$ of the same form as given above in (47)) can be considered. However, for the sake of brevity and clarity, we present the method in the case of equation (46) only.

The key idea is that eigenproblems of this form will necessarily fall back into the domain of validity of Proposition 1 after remapping (21). In contrast, we point out the following limitation to our technique: transcendental operators will not lend themselves to a sparse discretization. The method to obtain sparse representations for rational differential operators proceeds as follows.

As a first step, multiply the equation by $C(y)$, the least common multiple polynomial of the $Q_i(y)$ polynomials:

$$E\, C(y)\psi = \sum_{i=0}^{N_{\mathcal{H}}} P_i'(y) \left(\frac{\mathrm{d}}{\mathrm{d}y}\right)^i, \tag{48}$$

where the polynomials $P_i'(y) = C(y)P_i(y)/Q_i(y)$. As a second step, map the equation to θ-space, using the mapping (21) and the chain rule (22). Determine the highest power h_p of $\sin\theta$ present as a denominator: multiply the equation by $(\sin\theta)^{h_p}$ to cancel out all denominators, yielding to an equation of the form:

$$E\mathcal{M}f(\theta) = \mathcal{L}f(\theta). \tag{49}$$

Finally, obtain a sparse discretization of this system using a mixed expansion (15).

4.3. Particle in an arbitrary asymmetric well

The general method described above is quickly illustrated on an example of greater complexity than the translated parabolic well considered in section 4.1. We consider a potential of the form:

$$V(y) = V_0 \frac{a_4 y^4 + a_2 y^2}{y^2 + y + 1}, \tag{50}$$

which corresponds to the Hamiltonian:

$$\mathcal{H} = -\partial_{yy} + V_0 \frac{a_4 y^4 + a_2 y^2}{y^2 + y + 1}, \tag{51}$$

of the form given in equation (47). We first eliminate the rational fractions:

$$E\left(y^2 + y + 1\right)\psi(y) = -\left(y^2 + y + 1\right)\partial_{yy}\psi(y) + V_0\left(a_4 y^4 + a_2 y^2\right)\psi(y). \tag{52}$$

Mapping this equation to θ-space and eliminating the $\sin\theta$ denominators by multiplying by $\sin^4\theta$ gives, in this case:
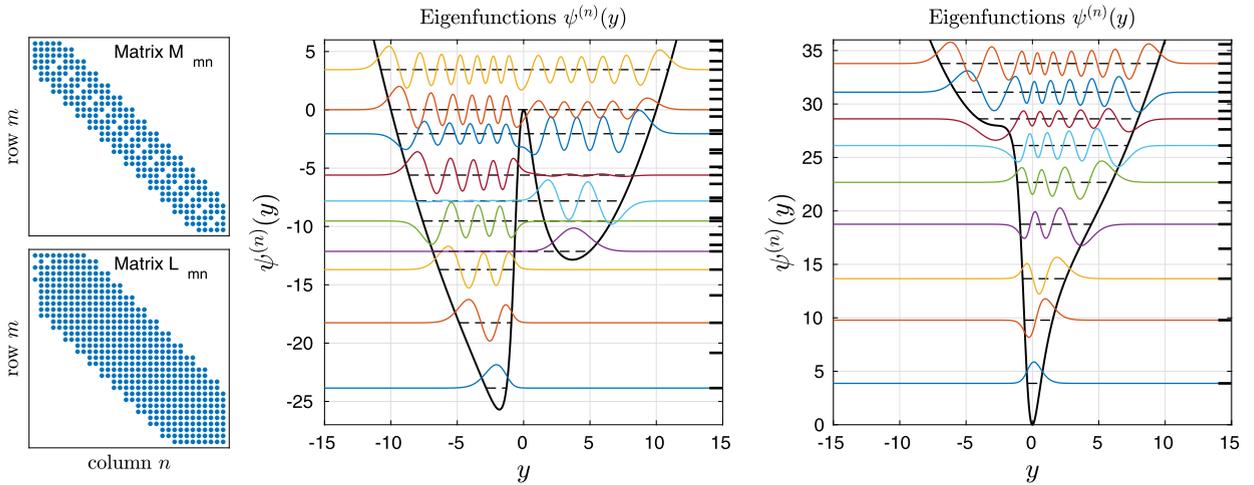
Fig. 5. A quantum particle in an asymmetric potential well (50). Left: Discretization of the Schrödinger equation. Eigenfunctions of potential (50) with $V_0 = 20$, $a_4 = 0.01$ and $a_2 = 1$ (center) or $a_2 = -1$ (right). Only a subset of the eigenfunctions (shifted vertically by the corresponding eigenenergy) are shown for legibility. All the eigenenergies are marked by thick ticks on the right of each figure.

$$E \sin^2 \theta \left(L^2 \cos^2 \theta + L \sin \theta \cos \theta + \sin^2 \theta \right) f(\theta) = V_0 \left(a_4 L^4 \cos^4 \theta + a_2 L^2 \cos^2 \theta \sin^2 \theta \right) f(\theta)$$

$$- \sin^4 \theta \left(L^2 \cos^2 \theta + L \sin \theta \cos \theta + \sin^2 \theta \right) \left(\frac{\sin^2 \theta}{L^2} \partial_{\theta\theta} + \frac{2 \cos \theta \sin \theta}{L^2} \partial_\theta \right) f(\theta). \tag{53}$$

This equation readily yields a sparse discretization when a mixed expansion is used for f. The sparse matrices involved are displayed in Fig. 5. The eigenfunctions of two different wells are computed to illustrate the variety of potentials amenable with our method.

5. Coupled systems: the equatorial β-plane

5.1. The equatorial β-plane

We now motivate and illustrate the application of the method for coupled systems of equations, namely, the linearized equatorial β-plane model [17] relevant for tropical atmospheric or oceanic dynamics. On the surface of a sphere covered with a shallow layer of fluid, the azimuthal velocity u, meridional velocity v and height of the fluid layer h obey the shallow water governing equations around the equator:

$$\partial_t u = -au + yv - \partial_x h, \tag{54a}$$

$$\partial_t v = -yu - av - \partial_y h, \tag{54b}$$

$$\partial_t h = -\partial_x u - \partial_y v, \tag{54c}$$

where x is the azimuthal coordinate, y is the meridional coordinate, and a is the Rayleigh friction coefficient.

As with the harmonic oscillator, we first present the analytical results against which the numerical method will be validated. In absence of friction, i.e. $a = 0$, the solutions are waves and these equations can be solved analytically by considering normal modes $(U(x, y), V(x, y), H(x, y)) = (u(y), v(y), h(y)) \exp(i[kx - \omega t])$. The variables $u(y)$ and $h(y)$ can both be expressed as functions of $v(y)$:

$$\left(\omega^2 - k^2 \right) h = i \left(ky - \omega \partial_y \right) v, \tag{55a}$$

$$\left(\omega^2 - k^2 \right) u = i \left(\omega y - k \partial_y \right) v. \tag{55b}$$

The variable v itself is determined as the solution of a Weber equation, similarly to the quantum harmonic oscillator:

$$\left(y^2 - \partial_{yy} \right) v = \Omega v \quad \text{where} \quad \Omega = \left(\omega^2 - k^2 - \frac{k}{\omega} \right). \tag{56}$$

As above, the solutions are parabolic cylinder functions and Ω is quantized:
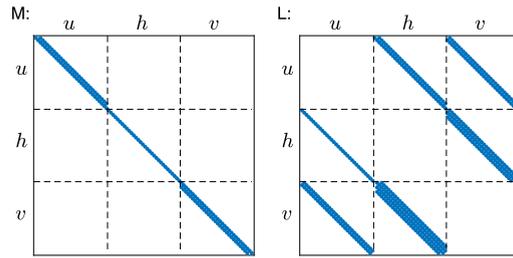
Fig. 6. Discretization of system of equations (60) using expansions (61) and ignoring friction, i.e. $a = 0$.

$$\Omega_n = 2n + 1, \quad n \in \mathbb{N}, \tag{57a}$$

$$v_n(y) = v_0 \, H_n(y) e^{-y^2/2}. \tag{57b}$$

Each value Ω_n corresponds to three different modes of frequency $\omega_n^{(1)}$, $\omega_n^{(2)}$, and $\omega_n^{(3)}$ found as the roots of equation (56). The wave functions of each mode $u_n^{(1)}(y)$, $h_n^{(1)}(y)$, $u_n^{(2)}(y)$, $h_n^{(2)}(y)$, $u_n^{(3)}(y)$, and $h_n^{(3)}(y)$, are obtained by using formulas (55). For $n \neq 0, -1$ these modes represent (westward and eastward propagating) inertio-gravity waves and a (westward propagating) Rossby wave. Two exceptions appear to this picture: for $n = 0$, one of the three roots, $\omega = -k$, must be discarded. Indeed, this root corresponds to an exponentially growing solution $v(y) = \exp\left(y^2/2\right)$. The remaining root represents the mixed Rossby-gravity (Yanai) wave. Further, the Kelvin waves are not captured by expression (57). Sometimes dubbed the "$n = -1$ mode", these waves correspond to v identically zero, yielding:

$$\omega_{-1} = k, \tag{58a}$$

$$v_{-1}(y) = 0, \tag{58b}$$

$$u_{-1}(y) = h_{-1}(h) \propto \exp\left(-y^2/2\right). \tag{58c}$$

5.2. Numerical treatment

When Rayleigh friction is present, the complex growth rate λ of a normal mode of azimuthal wave number k obeys the eigenproblem:

$$\lambda \begin{pmatrix} u \\ h \\ v \end{pmatrix} = \mathscr{L} \begin{pmatrix} u \\ h \\ v \end{pmatrix} \quad \text{where} \quad \mathscr{L} = \begin{pmatrix} \boxed{\begin{matrix} -a & -ik \\ -ik & -a \end{matrix}} & \begin{matrix} y \\ -\partial_y \end{matrix} \\ \begin{matrix} -y & -\partial_y \end{matrix} & \boxed{-a} \end{pmatrix}. \tag{59}$$

Observe that variables form two groups: u and h form the first group and v forms the second. The reason behind this segregation becomes clear when observing the blocks which form $\mathscr{L}$. In equation (59), the diagonal blocks (marked with rectangles) are composed of parity-preserving operators, whereas the off-diagonal blocks contain exclusively parity-flipping operators. Hence, in the spirit of the mixed expansions we have been using before in scalar equations, here we will use expansions of opposed parities in y for different variables. As emphasized before, this remedy becomes clear after remapping the equation to θ-space ($u^\dagger(\theta) = u(L \cot\theta)$, etc.):

$$\lambda \mathscr{M} \begin{pmatrix} u^\dagger \\ h^\dagger \\ v^\dagger \end{pmatrix} = \mathscr{L} \begin{pmatrix} u^\dagger \\ h^\dagger \\ v^\dagger \end{pmatrix}, \tag{60a}$$

$$\text{where} \quad \mathscr{L} = \begin{pmatrix} \boxed{\begin{matrix} -a\sin\theta & -ik\sin\theta \\ -i & -a \end{matrix}} & \begin{matrix} L\cos\theta \\ L^{-1}\sin^2\theta\,\partial_\theta \end{matrix} \\ \begin{matrix} -L\cos\theta & L^{-1}\sin^3\partial_\theta \end{matrix} & \boxed{-a\sin\theta} \end{pmatrix}, \quad \mathscr{M} = \begin{pmatrix} \boxed{\begin{matrix} \sin\theta & 0 \\ 0 & 1 \end{matrix}} & \begin{matrix} 0 \\ 0 \end{matrix} \\ \begin{matrix} 0 & 0 \end{matrix} & \boxed{\sin\theta} \end{pmatrix}. \tag{60b}$$

It becomes immediately clear that the sparsity of the discretization of the system is preserved if expansions of opposite parities are used for $(u^\dagger, h^\dagger)$ and $v^\dagger$. In the following, we present results obtained with a cosine expansion for $u^\dagger$ and $h^\dagger$ and a sine expansion for $v^\dagger$ (which yield to the structure depicted in Fig. 6 for matrices **M** and **L**):

$$u^\dagger(\theta) = \sum_m \widetilde{u}_m \cos(m\theta), \quad h^\dagger(\theta) = \sum_m \widetilde{h}_m \cos(m\theta), \quad v^\dagger(\theta) = \sum_m \widetilde{v}_m \sin(m\theta). \tag{61}$$

We compare in Fig. 7 the frequencies obtained by solving the discretized eigenproblem to Rossby and inertio-gravity waves (57a), and Kelvin waves (58a). The accuracy of a spectral method based on Chebyshev functions yields an excellent agreement. We also display on Fig. 7 some eigenmodes. We illustrate how for a given Ω_n, Rossby and inertio-gravity
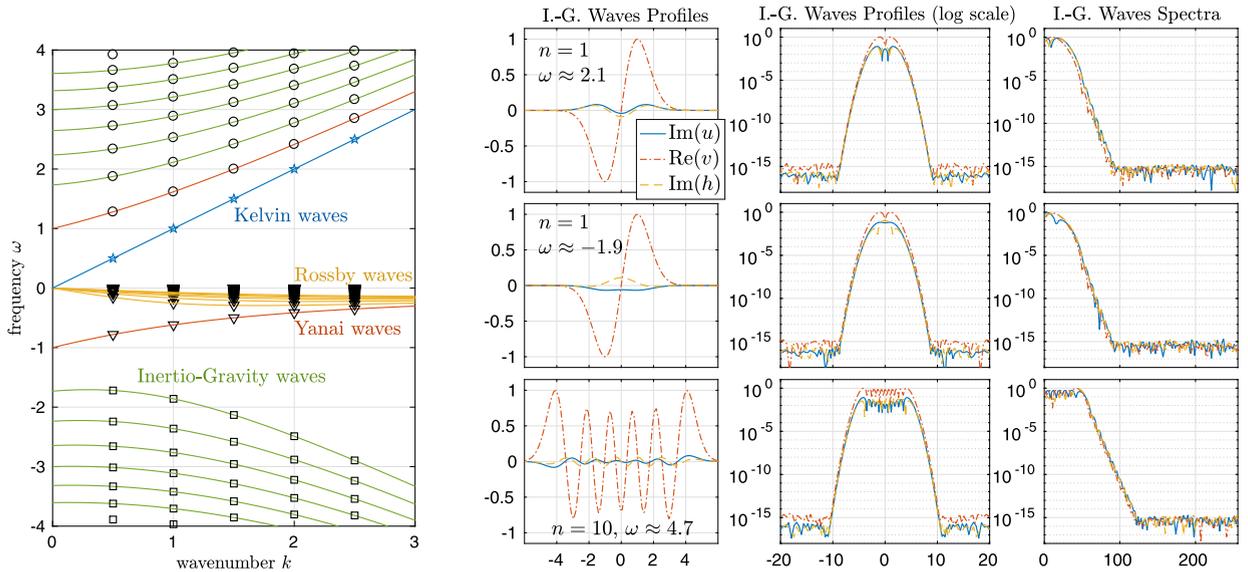
Fig. 7. Eigenmodes of the frictionless beta-plane. Left panel: the plain lines indicate the different branches of the dispersion relation; the blue line correspond to Kelvin waves ($n = -1$ mode), the orange lines to the Yanai wave ($n = 0$ mode), and the yellow and black lines the inertio-gravity and Rossby waves respectively ($\Omega_n = 2n + 1$). The symbols correspond to the frequency determined numerically, for $k = 1/2$, 1, 3/2, 2, and 5/2. Right panels concern Inertio-Gravity waves: the first column displays meridional profiles of eigenmodes u, v and h as a function of y; the second column uses a logarithmic scale to illustrate the decay of the modes to zero within machine precision on the collocation grid; the third column illustrates the exponential convergence of the expansion. First two rows: two different modes for $n = 1$ and $\Omega_1 = 3$: note the common value for v (given by (57b)) but the different u and h profiles. The third row corresponds to $n = 10$ and $\Omega_{10} = 21$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
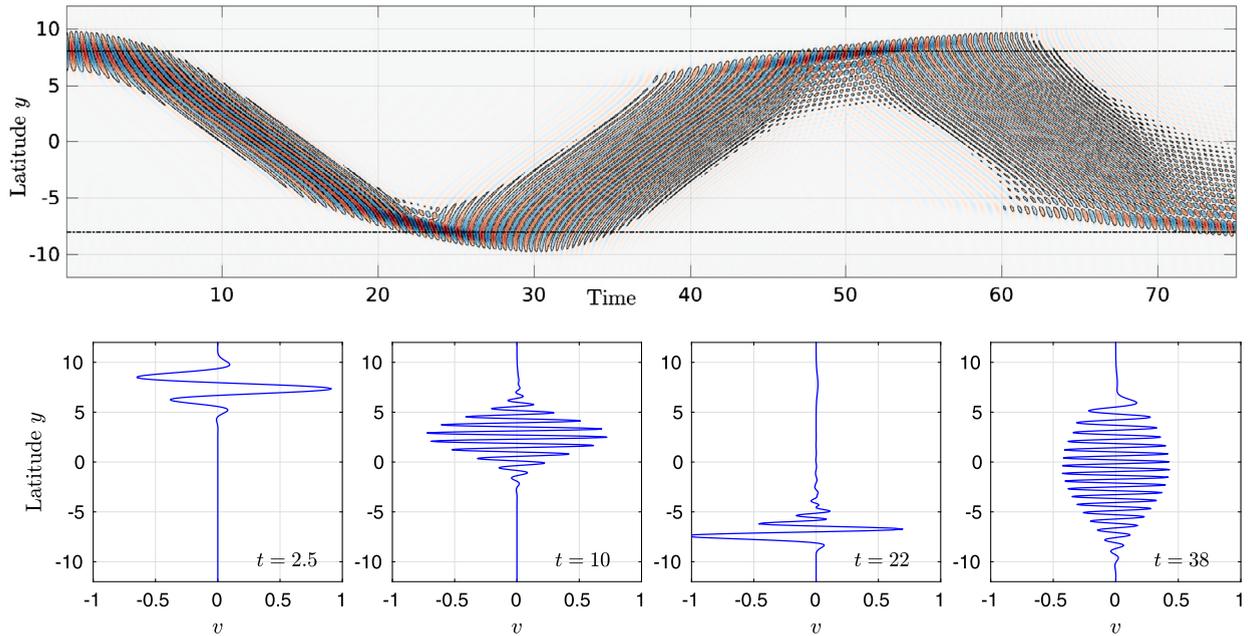
Fig. 8. Upper panel: Space–time diagram of a latitudinally trapped wave packet, with the initial condition: $u = 0$, $h = 0$, $v = \cos(x)e^{-\frac{(y-8)^2}{2}}$. We display $v(y, t)$ at $x = 0$ as a function of y and t. Lower panels: corresponding velocity profiles $v(y)$ at $x = 0$ for different times.

waves with different frequencies correspond to the roots of equation (56). These modes have a common azimuthal velocity v but distinct meridional velocity u and height of fluid h. We show plots of the amplitude of the profiles $|u(y)|$, $|v(y)|$, and $|h(y)|$ using a semi-logarithmic scale to illustrate the exponential decay of both variables as $|y| \to \infty$. Finally, the exponential convergence of the three spectra $|\tilde{u}|$, $|\tilde{v}|$, and $|\tilde{h}|$ is shown on the last panel of Fig. 7.

We readily obtain a time-stepping scheme for the dynamics of equatorial waves by utilizing the matrices **M** and **L** with a implicit third-order Runge–Kutta scheme [8]. We show temporal dynamics in Fig. 8. A initial perturbation localized around

a latitude y_0 travels across the equator and is reflected when reaching the opposite latitude $-y_0$. In absence of dissipation, the wavepacket then continues bouncing back and forth between the bounding latitudes $-y_0$ and y_0, illustrating a well known feature (see for instance [17]) of the equatorial region: this region behaves as a trapping wave guide.

6. Multidimensional domains

6.1. The spirit of the method

The machinery of the Chebyshev functions on the infinite line is readily generalized for separable operators to bi- and tridimensional domains with at least one unbounded direction. For the sake of brevity, this paper describes only bidimensional domains $(x, y) \in \Omega$, the generalization to 3D being a mere bookkeeping exercise. As described in [18], the philosophy consists in decomposing variables as a truncated expansion of a well-suited family of basis functions $\alpha_n(x)$ and $\beta_m(y)$ along each direction in the following fashion:

$$f(x, y) \approx \sum_{n=1}^{N_x} \sum_{m=1}^{N_y} \widetilde{f}_{nm} \alpha_n(x) \beta_m(y). \tag{62}$$

The procedure is as follows. First, select the bases α_n and β_m: Fourier modes should be used along periodic directions; Chebyshev polynomials are almost always the best choice along bounded directions; unbounded directions can be remapped to $[0, \pi]$, where cosines, sines, or the hybrid bases $\breve{\beta}$ or $\check{\beta}$ can be used, depending on the parity properties of the operators involved as discussed above.

Second, regularize the equations so that they contain only operators with a sparse representation. We have indicated a procedure for unbounded directions in section 4.2; for bounded directions where Chebyshev polynomials are used, follow the quasi-inverse method of [19] equivalent to repeated integrations to remove all derivatives along these directions.

Third, decompose the separable operators into a sum of the form $\mathscr{L} = \sum_i \mathscr{P}_i \mathscr{Q}_i$, where the $\mathscr{P}_i$ (resp. $\mathscr{Q}_i$) operators act on the variable x (resp. y) only. Compute the matrices that represent the action of these operators on the selected bases $\alpha_n(x)$ and $\beta_m(y)$:

$$\mathbf{P}_i = \frac{\langle \alpha_m, \mathscr{P}_i \alpha_n \rangle}{\|\alpha_m\|^2}, \qquad \mathbf{Q}_i = \frac{\langle \beta_m, \mathscr{Q}_i \beta_n \rangle}{\|\beta_m\|^2}. \tag{63}$$

These sparse matrices form the building blocks for the discretization of $\mathscr{L}$ which may now be assembled using the Kronecker product to obtain the matrix representing $\mathscr{L}$:

$$\mathbf{L} = \sum_i \mathbf{P}_i \otimes \mathbf{Q}_i. \tag{64}$$

A similar treatment should be applied to the operator $\mathscr{M}$. The Kronecker product offers a simple way to compute the representation of operators $\mathscr{M}$ and $\mathscr{L}$ on the basis $(\alpha_0\beta_0, \alpha_0\beta_1, \ldots, \alpha_0\beta_{N_y}, \alpha_1\beta_0, \ldots \alpha_{N_x}\beta_{N_y})$. The matrix of spectral coefficients $\widetilde{f}_{nm}$ can be vectorized into a long column vector of state

$$\boldsymbol{f} = \text{vec}(f_{nm}), \tag{65}$$

which represents a stack of the columns of $\widetilde{f}_{nm}$. Finally, we obtain the familiar form $E\mathbf{M}\boldsymbol{f} = \mathbf{M}\boldsymbol{f}$.

We illustrate this method on the three cases that occur with at least one unbounded direction. In this section, we consider a 2D quantum harmonic oscillator on the infinite plane $\Omega = \mathbb{R}^2$ and on the infinite strip with finite width $\Omega = [x_1, x_2] \times \mathbb{R}$. In the next section, the case of a domain with an infinite direction and a periodic direction (infinite cylinder) is considered through the example of the Kelvin–Helmholtz instability. This is investigated through a time-stepping method in contrast with the eigenvalue problems considered in the present section.

6.2. Bidimensional benchmark: the 2D QHO

We consider a particle in the two-dimensional potential

$$V(x, y) = \frac{1}{2}x^2 - x + \frac{5}{6}y^2 - \frac{10}{3}y + \frac{\sqrt{3}}{3}(xy - 2x - y) + \frac{13}{3}, \tag{66}$$

and obeying the 2D Schrödinger equation

$$E\psi = \left(-\partial_{xx} - \partial_{yy} + V(x, y)\right)\psi. \tag{67}$$

The motivation behind the choice of potential $V(x, y)$ is that this potential is nothing more than the simple potential $\bar{V}(X, Y) = X^2 + Y^2/3$ rotated by an angle $\pi/3$ and translated by the quantity $(1, 2)$. The resulting potential thus possesses no parity in x nor y and features a coupling term xy. Hence, equation (67) is a test problem of reasonable complexity

for our method. Despite this apparent complexity, the results are easily checked against analytic solutions. Indeed, since translations and rotations are unitary transforms, the eigenvalues of V are equal to the eigenvalues of $\bar{V}$, namely:

$$E_{nm} = (2n+1) + \frac{1}{\sqrt{3}}(2m+1), \quad (n,m) \in \mathbb{N}^2. \tag{68}$$

Further, the eigenfunctions of V are obtained by applying the same rotation and translation to those of $\bar{V}$:

$$\psi^{(nm)}(X,Y) = H_n(X) \exp\left(-X^2/2\right) \times H_m(Y/3^{1/4})\left(-Y^2/\left[2\sqrt{3}\right]\right). \tag{69}$$

6.3. The infinite plane $\Omega = \mathbb{R}^2$

On the infinite plane, since we have demonstrated the efficiency and accuracy of the Chebyshev approximation, we simply decompose our variable on two sets of Chebyshev functions along x and y. As in the 1D case, more intuition is generally available when dealing with trigonometric functions instead of the full-fledged Chebyshev functions: hence we remap both the coordinates x and y using:

$$x = L_x \cot(\chi), \tag{70a}$$
$$y = L_y \cot(\theta). \tag{70b}$$

where L_x and L_y are two real mapping parameters that can be optimized for accuracy, depending on the problem. Our Schrödinger equation (67) for $\psi(x,y)$ is mapped into the following equation for $f(\chi,\theta) = \psi(x(\chi),y(\theta))$:

$$E\mathscr{M}f(\chi,\theta) = \mathscr{L}f(\chi,\theta), \tag{71a}$$

where

$$\mathscr{M} = \sin^2\chi\sin^2\theta, \tag{71b}$$

$$\begin{aligned}
\mathscr{L} = & -\frac{1}{L_x^2}\left(\sin^6\chi\sin^2\theta\,\partial_{\chi\chi} + 2\cos\chi\sin^5\chi\sin^2\theta\,\partial_\chi\right) - \frac{1}{L_y^2}\left(\sin^6\theta\sin^2\chi\,\partial_{\theta\theta} + 2\cos\theta\sin^5\theta\sin^2\chi\,\partial_\theta\right) \\
& + \frac{L_x^2}{2}\cos^2\chi\sin^2\theta - \left(1 + \frac{2\sqrt{3}}{3}\right)L_x\cos\chi\sin\chi\sin^2\theta \\
& + \frac{5L_y^2}{6}\sin^2\chi\cos^2\theta - \left(\frac{10}{3} + \frac{\sqrt{3}}{3}\right)L_y\cos\theta\sin\theta\sin^2\chi \\
& + \frac{\sqrt{3}L_xL_y}{3}\cos\chi\sin\chi\cos\theta\sin\theta + \frac{13}{3}\sin^2\chi\sin^2\theta,
\end{aligned} \tag{71c}$$

after multiplication (regularization) by $\sin^2\chi\sin^2\theta$ to preserve sparsity. Due to the presence of parity mixing in both directions θ and χ, we discretize this equation using mixed expansion in χ and θ: for instance $\breve{\beta}(\chi)$ and $\breve{\beta}(\theta)$:

$$f(\chi,\theta) \approx \sum_{n,m=0}^{N}\widetilde{f}_{nm}\breve{\beta}_n(\chi)\breve{\beta}_m(\theta). \tag{72}$$

Note that nothing would prevent us from using different expansions in each direction, for instance, $\breve{\beta}(\chi)$ and $\mathring{\beta}(\theta)$. Doing so, however, would unnecessarily demand to compute the action of the operators on two different bases, yielding twice as much matrix building compared to using a common basis for both χ and θ. The discretized version of system (71), with $\boldsymbol{f} = \text{vec}(\widetilde{f}_{nm})$:

$$E\,\mathbf{M}\,\boldsymbol{f} = \mathbf{L}\,\boldsymbol{f} \tag{73a}$$

can be expressed with a set of only five matrices:

$$\mathbf{M} = \mathbf{D} \otimes \mathbf{D}, \tag{73b}$$

$$\begin{aligned}
\mathbf{L} = & -\frac{1}{L_x^2}\left(\mathbf{A}\otimes\mathbf{D} + 2\mathbf{B}\otimes\mathbf{D}\right) - \frac{1}{L_y^2}\left(\mathbf{D}\otimes\mathbf{A} + 2\mathbf{D}\otimes\mathbf{B}\right) + \frac{L_x^2}{2}\mathbf{C}\otimes\mathbf{D} - \left(1 + \frac{2\sqrt{3}}{3}\right)L_x\mathbf{E}\otimes\mathbf{D} \\
& + \frac{5L_y^2}{6}\mathbf{D}\otimes\mathbf{C} - \left(\frac{10}{3} + \frac{\sqrt{3}}{3}\right)L_y\mathbf{D}\otimes\mathbf{E} + \frac{\sqrt{3}L_xL_y}{3}\mathbf{E}\otimes\mathbf{E} + \frac{13}{3}\mathbf{D}\otimes\mathbf{D}.
\end{aligned} \tag{73c}$$

where matrices $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, $\mathbf{D}$, and $\mathbf{E}$ are defined in equations (43) and computed in appendix C.1.

6.4. The infinite strip

We now turn to solving equation (67) on an infinite strip $\Omega = [x_1, x_2] \times \mathbb{R}$ with Dirichlet boundary conditions $\psi(x_1) = \psi(x_2) = 0$. As usual, the infinite y-direction is remapped to $\theta \in [0, \pi]$ and the bounded x-direction is remapped to $\chi \in [-1, 1]$ via:

$$x = x_0 + \chi \, \Delta x, \quad \text{with} \quad x_0 = \frac{x_1 + x_2}{2} \quad \text{and} \quad \Delta x = \frac{x_2 - x_1}{2}, \tag{74a}$$

$$\theta = L \cot(\theta). \tag{74b}$$

Since the spirit of this paper is to provide recipes for memory efficient discretization methods along the infinite line, it is desirable to couple them with memory efficient methods along the other directions too. When a truncated Chebyshev polynomial expansion is considered along a bounded direction, a sparse representation is obtained for operators containing only polynomials and derivatives. The quasi-inverse method [3] is employed that consists of integrating the governing equation or set of equations as many times as the highest derivative present. In our example, we will integrate twice to cancel out the second order x derivative, so that equation (67) becomes:

$$E \sin^2 \theta \iint d^2x f(x, \theta) = -\sin^2 \theta f(x, \theta) - \frac{1}{L^2} \iint d^2x \left[\sin^6 \theta \partial_{\theta\theta} + 2 \cos \theta \sin^5 \theta \partial_\theta \right] f(x, \theta)$$

$$+ \iint d^2x \left[\frac{1}{2} x^2 \sin^2 \theta - \left(1 + \frac{2\sqrt{3}}{3} \right) x \sin^2 \theta + \frac{5}{6} \cos^2 \theta - \left(\frac{10}{3} + \frac{\sqrt{3}}{3} \right) \cos \theta \sin \theta + \frac{\sqrt{3}}{3} x \cos \theta \sin \theta \right] f(x, \theta)$$

$$+ ax + b. \tag{75}$$

Observe that by integrating twice we have introduced two arbitrary integration constants a and b. The value of these yet floating integration constants will be determined by imposing boundary conditions as the last step of our discretization. For the time being, we define the discrete representation of three useful operators in Chebyshev space, denoted $\mathbf{Q}$:

$$\mathbf{Q}^{(0)} = \frac{\langle T_m, \iint T_n d^2x \rangle}{\| T_m \|^2}, \tag{76a}$$

$$\mathbf{Q}^{(1)} = \frac{\langle T_m, \iint x T_n d^2x \rangle}{\| T_m \|^2}, \tag{76b}$$

$$\mathbf{Q}^{(2)} = \frac{\langle T_m, \iint x^2 T_n d^2x \rangle}{\| T_m \|^2}. \tag{76c}$$

These matrices are banded with respectively two-, three- and four-bands. We give their full expressions in the appendix C.2 and show that given the discretized matrix $\mathbf{x}$ for x we arrive at $\mathbf{Q}^{(1)} = \mathbf{Q}^{(0)}\mathbf{x}$ and $\mathbf{Q}^{(2)} = \mathbf{Q}^{(0)}\mathbf{x}^2$ with simply matrix multiplications. Similarly to the previous section, we easily obtain our matrices $\mathbf{M}$ and $\mathbf{L}$ by combining the sparse representations of operators along different directions with the Kronecker product:

$$\mathbf{M} = \mathbf{Q}^{(0)} \otimes \mathbf{D}, \tag{77}$$

$$\mathbf{L} = -\mathbf{1}_x \otimes \mathbf{D} - \frac{1}{L^2} \mathbf{Q}^{(0)} \otimes [\mathbf{A} + 2\mathbf{B}] + \frac{1}{2} \mathbf{Q}^{(2)} \otimes \mathbf{D}$$

$$- \left(1 + \frac{2\sqrt{3}}{3} \right) \mathbf{Q}^{(1)} \otimes \mathbf{D} + \frac{5}{6} \mathbf{Q}^{(0)} \otimes \mathbf{C} - \left(\frac{10 + \sqrt{3}}{3} \right) \mathbf{Q}^{(0)} \otimes \mathbf{E} + \frac{\sqrt{3}}{3} \mathbf{Q}^{(1)} \otimes \mathbf{E}. \tag{78}$$

The last step consists in implementing boundary conditions: as mentioned above the first two coefficients of the expansion in x are polluted by integration factors $a(\theta)$ and $b(\theta)$. As a consequence, the $2N_y$ lines in $\mathbf{M}$ and $\mathbf{L}$ that correspond to coefficients f_{0m} and f_{1m} (with $m \in [1, N_y]$) have to enforce Dirichlet boundary conditions, namely:

$$\forall m: \quad \sum_n f_{nm} = 0, \quad \sum_n (-1)^n f_{nm} = 0, \tag{79}$$

which we write

$$\forall m: \quad \sum_n b_{0n} f_{nm} = 0, \quad \sum_n b_{1n} f_{nm} = 0, \tag{80}$$

for a greater generality. Any boundary conditions can be implemented easily by using only matrix products in the following manner. We define a boundary-condition matrix $\mathbf{b}$ which is zero everywhere, except for the first two lines b_{0n} and b_{1n}, and a diagonal matrix $\mathbf{1}^{(2)}$ equal to the identity everywhere except on the first two lines, which are identically zero. Define:
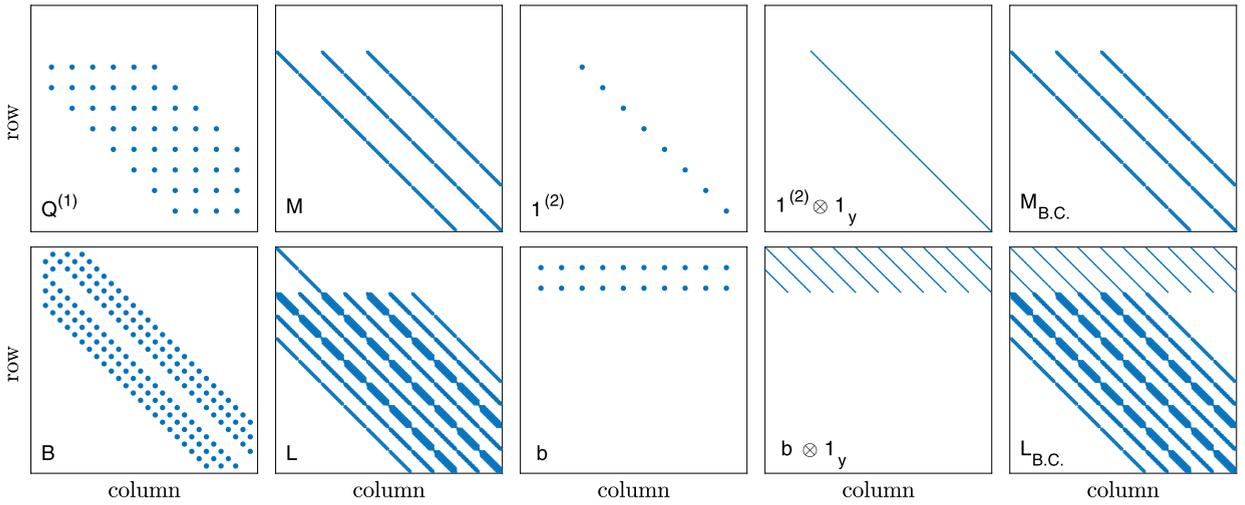
Fig. 9. Discretization of equation (71a) for the infinite strip geometry. Matrices **M** and **L** are obtained as Kronecker products of sparse matrices acting on the unbounded direction (for instance **B**) or on the bounded direction (for instance **Q**), and therefore inherit their sparsity. Along the bounded direction, boundary conditions are expressed with matrices **b** and $\mathbf{1}^{(2)}$. Finally, the discretized system including boundary conditions is represented by matrices $\mathbf{M}_{\text{B.C.}}$ and $\mathbf{L}_{\text{B.C.}}$ following equations (81).
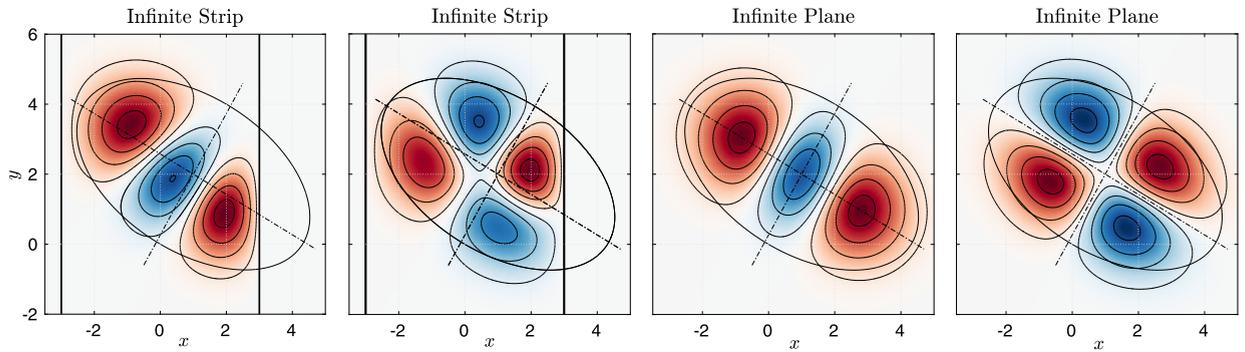
Fig. 10. Solution $\psi(x, y)$. Modes $(1, 3)$ and $(2, 2)$ for a particle in the 2D quadratic potential given in equation (66), in presence of walls at $x = \pm 3$ (thick vertical lines, two left panels), or on an infinite plane (two right panels). To guide the eye, the dashed ellipses represent an equipotential line; the dashed lines indicate the two main directions of the potential.

$$\mathbf{M}_{\text{B.C.}} = \left(\mathbf{1}^{(2)} \otimes \mathbf{1}_y \right) \mathbf{M} \,, \tag{81a}$$

$$\mathbf{L}_{\text{B.C.}} = \left(\mathbf{1}^{(2)} \otimes \mathbf{1}_y \right) \mathbf{L} + \left(\mathbf{b} \otimes \mathbf{1}_y \right) \,. \tag{81b}$$

The mechanics of this simple algebra, illustrated on Fig. 9, is the following: multiplying by $\mathbf{1}^{(2)} \otimes \mathbf{1}_y$ zeroes out the $2N_y$ lines of matrices **M** and **L** that correspond to projection onto the two lowest order Chebyshev modes that are polluted by arbitrary integration constants, as seen in equation (75). In $\mathbf{L}_{\text{B.C.}}$, these lines are replaced with the boundary condition by adding $\mathbf{b} \otimes \mathbf{1}_y$. Eigenmodes for both the infinite strip and the infinite plane are shown on Fig. 10. The accuracy of this bidimensional example is omitted but has been checked to be as good as in the 1D case.

7. Nonlinear time-stepping: the Kelvin–Helmholtz instability

7.1. Background parallel flow and formulation

We consider a streamwise invariant parallel flow $U_0(y)\boldsymbol{e}_x$ imposed by external forces in a periodic domain along the streamwise direction $x \in [0, 2\pi[$ with an unbounded spanwise direction $y \in \mathbb{R}$. Perturbations $\boldsymbol{u}(x, y, t) = -\partial_z \psi \, \boldsymbol{e}_x + \partial_x \psi \, \boldsymbol{e}_z$ to the background flow obey the incompressible governing equation:

$$\partial_t \nabla^2 \psi + U_0 \, \partial_x \nabla^2 \psi - \partial_{zz} U_0 \, \partial_x \psi + \mathcal{J}\left(\psi, \nabla^2 \psi \right) = \frac{1}{\text{Re}} \nabla^4 \psi \,. \tag{82}$$

Such flows are known to be potentially linearly unstable if they present an inflection point $\partial_{zz} U_0 \neq 0$ (see for example [20]). The Reynolds number, measuring the strength of the viscous dissipation force, controls the stability of the shear flow, i.e., the flow stabilizes as Re is lowered. In this section, the stability of the flow

$$U_0(y) = \frac{y}{1 + y^4} \tag{83}$$

is analyzed by means of a time-stepping integration of the full nonlinear governing equation (82).

7.2. Discretization

Following the usual prescription, we remap the equations to θ-space and multiply out with the factor $\left(L^4 \cos^4 \theta + \sin^4 \theta\right)^3$ to obtain:

$$\mathcal{M} \partial_t f = \mathcal{L} f + \mathcal{N}(f), \tag{84}$$

where $\mathcal{M}$ and $\mathcal{L}$ are linear operators and $\mathcal{N}(f)$ is the nonlinear advection term:

$$\mathcal{M} = \left(L^4 \cos^4 \theta + \sin^4 \theta\right)^3 \nabla^2 \tag{85a}$$

$$\mathcal{L} = -L \cos \theta \sin^3 \theta \left(L^4 \cos^4 \theta + \sin^4 \theta\right)^2 \partial_x \nabla^2 + 4L^3 \cos^3 \theta \left(3 \cos^4 \theta L^4 - 5 \sin^4 \theta\right) \sin^5 \theta \partial_x$$
$$+ \frac{1}{\text{Re}} \left(L^4 \cos^4 \theta + \sin^4 \theta\right)^3 \nabla^4 \tag{85b}$$

$$\nabla^2 = \left(\frac{\sin^4 \theta}{L^2} \partial_{\theta\theta} + \frac{2 \sin^3 \theta \cos \theta}{L^2} \partial_\theta + \partial_{xx}\right) \tag{85c}$$

$$\mathcal{N}(f) = \left(L^4 \cos^4 \theta + \sin^4 \theta\right)^3 \frac{\sin^2 \theta}{L} \left(-\partial_x f \partial_\theta \nabla^2 f + \partial_\theta f \partial_x \nabla^2 f\right) \tag{85d}$$

We expand our variable f as a Fourier series along the periodic direction x. Observing that parity mixing is present amongst the operators, we chose to represent f as a hybrid cosine and sine series in θ:

$$f(x, \theta) \approx \sum_{k=0}^{N_x - 1} \sum_{n=1}^{N_y} \left(\widetilde{f}_{k+1,n} \breve{\beta}_n(\theta) e^{ikx}\right) + \text{c.c.} \tag{86}$$

As the operators $\mathcal{M}$ and $\mathcal{L}$ are separable, we easily obtain their discrete representation **M** and **L** by using Kronecker products. We compute the nonlinear term pseudospectrally by first evaluating derivatives $\partial_x f$, $\partial_\theta \nabla^2 f$, *etc.*, in spectral space. By means of FFT, we then transform these variables to physical space where the products are computed, and transform back to spectral space. A 2/3 dealiasing rule, appropriate for quadratic nonlinearities, is used.

7.3. Semi-implicit time-stepping schemes

IMEX schemes are a popular class of time-stepping schemes in computational fluid mechanics: these schemes combine an implicit treatment of the linear terms, whereas nonlinearities are treated explicitly. This popularity stems from the salient features of the Navier–Stokes equations which typically contains stiff linear terms, advocating for the use of stable and accurate implicit schemes. An implicit treatment of the nonlinear term would be costly and ultimately requires iterative methods compared to direct solves. Fortunately an explicit treatment yields satisfactory results. The simplest IMEX scheme, used in the following, is the order one Backward Euler which obtains $f^{(i+1)}$ from $f^{(i)}$ after a timestep τ by solving the linear system:

$$\left(\mathbf{M} - \tau \mathbf{L}\right) f^{(i+1)} = \mathbf{M} f^{(i)} + \tau \mathbf{N}\left(f^{(i)}\right). \tag{87}$$

7.4. Stability of the flow: numerical results

Equation (82) is investigated by time-integration starting from an initial condition consisting of small amplitude noise uniformly distributed over the resolved spatial scales. The time series of the kinetic energy E_K is displayed on Fig. 11 for seven values of the Reynolds number Re. Three phases are distinguished.

During the initial phase, as most of the modes that compose the initial condition are stable, the total energy first plummets (for $t \lesssim 20$). During the second phase, an optimal mode emerges from the sea of decaying modes: this mode is either the most unstable (fastest growing) mode or the least damped mode, depending on the value of Re. The planform of
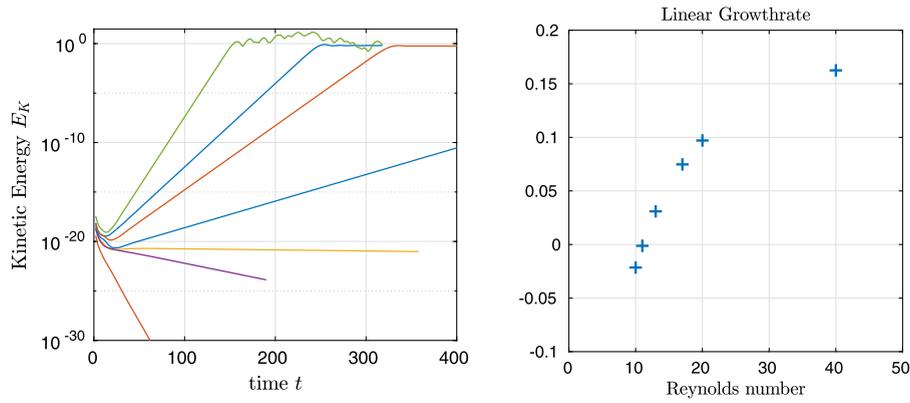
Fig. 11. Stability of the base flow $U_0(y)\,\boldsymbol{e}_x$ (see equation (83)). Left panel: kinetic energy time series for various Reynolds numbers. From bottom to top $\mathrm{Re} = 5, 10, 11, 13, 17, 20, 40$. Right panel: growthrate λ during the exponential growth phase as a function of Re.
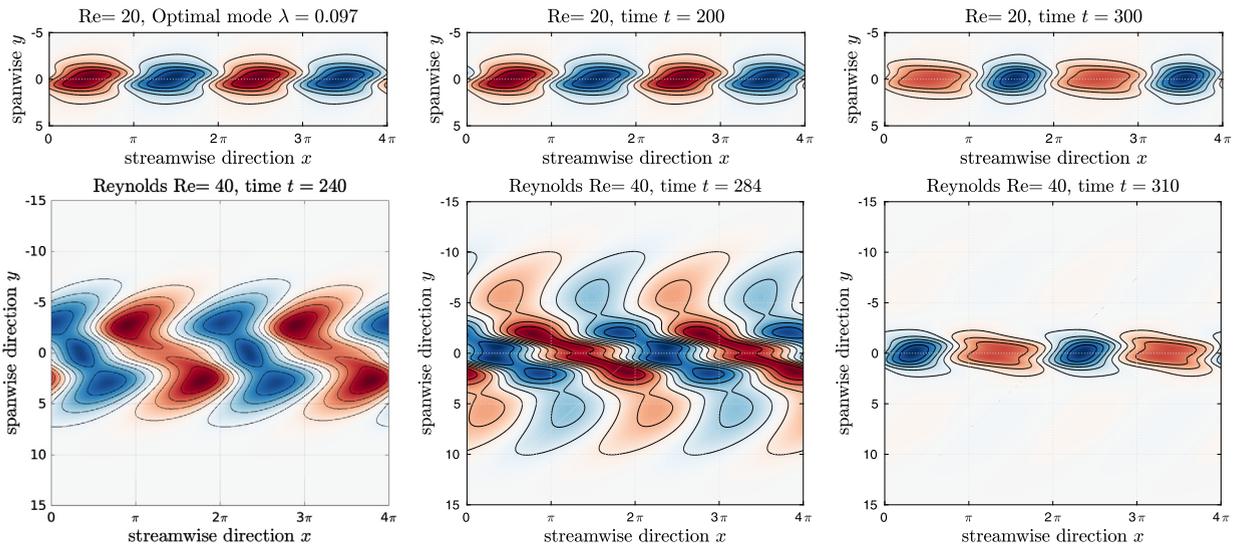
Fig. 12. Planform of the streamfunction of the perturbation $\psi(x, y)$. Top row, Re $= 20$: Optimal mode obtained by solving the eigenproblem (left); snapshot during the exponential growth phase (center); and snapshot after saturation (right). Bottom row, Re $= 40$: snapshots at three different instants after saturation, illustrating the intermittency of the flow.

the mode that dominates the temporal dynamics is depicted in Fig. 12 for Re $= 20$ and compared with the optimal mode obtained by means of the corresponding eigenproblem. As expected, the agreement is excellent. As long as the dynamics remains in the linear regime ($|\psi| \ll 1$). The total kinetic energy is dominated by the energy of the optimal mode alone which follows an exponential growth trend in time $E_K \propto \exp(2\lambda t)$. The optimal growth rate $\lambda(\mathrm{Re})$ can be fitted for each values of Re. Fig. 11 shows how the optimal growth rate $\lambda(\mathrm{Re})$ varies as a function of Re. We obtain the critical Reynolds number Re$_c \approx 11.066$ corresponding to the threshold of the instability by fitting the growthrate λ with a polynomial in Re around the onset.

Finally, the instability saturates as ψ becomes of order $O(1)$ (Fig. 11). As can be observed from the timeseries at modest supercriticality Re ≤ 20, the saturated state becomes steady after a short overshoot. The steady planform of the saturated state is shown in Fig. 12 for Re $= 20$ and displays a noticeable but vertically contained distortion compared to the linear optimal mode. By contrast, when the supercriticality is increase (e.g. for Re $= 40$), the saturated state becomes unsteady, showing intermittent bursts of activity, as reported in the bottom row of Fig. 12. The unstable layer of fluid around $y = 0$ then sheds vortices to the quiescent surrounding regions. Here, we notice the limitations of our method. As no proper physical trapping mechanism is present, viscous damping remains the only source of control, thus the outwardly propagating vortices can violate the CFL criterion as they propagate fast enough to reach regions where the collocation grid becomes coarser and coarser. If viscosity acts on timescales longer than the time needed for these vortices to reach large values of $|y|$, these vortices will become underresolved before they are damped out. The accuracy and stability of the methods is then jeopardized. If no additional physics is introduced to resolve this issue, the solution is to increase the number of basis function along y and to use a larger L to increase the size of the region where the dynamics is safely resolved. Note that a similar increase in resolution would be necessary if one were to chose to introduce boundary conditions to tackle

the problem of an unstable shear flow in an unbounded y-direction: these boundary conditions would have to be placed "far enough", so that they do not affect the dynamics. Increasing the distance between boundaries would equally lead to increasing the number of grid points or Chebyshev polynomials in the y-direction.

8. Conclusion

In this paper, we have illustrated on fundamental examples originated from topics as varied as quantum mechanics, geophysical fluid dynamics, or fluid mechanics, the implementation of sparse spectral methods for solving PDE systems in unbounded domains of arbitrary dimensions.

Central to our technique are hybrid expansions $\{TS_n\}$ and $\{ST_n\}$ (equation (28)), composed of Chebyshev functions $\{TB_n\}$ and $\{SB_n\}$ which preserve the sparsity of the discretization of differential operators on the infinite line even in presence of parity flipping differential operators in the general form of polynomial or rational fractions. Chebyshev functions correspond to remapped cos and sin functions to the infinite line, so that our hybrid expansions are remapped bases of interleaved sines and cosines, $\{\check{\beta}_n\}$ and $\{\mathring{\beta}_n\}$ defined in equations (17). For exponentially decaying functions, we have shown that our hybrid Chebyshev expansions inherit the exponential convergence properties of standard Chebyshev expansions.

Numerical analyses of systems of differential equations are greatly facilitated, or made possible at all, by the sparsity of the discretization of rational differential operators that results from this choice of hybrid bases. We have shown examples of investigation of the linear dynamics by solving sparse eigenproblems or by means of strongly accurate and stable implicit time-marching schemes, along with examples of analyses of nonlinear dynamics, tackled by means of popular IMEX schemes. The method is easily generalized to domains of higher dimensionality by means of Kronecker products.

To facilitate the implementation of this methodology and its adaptation to other problems, a toolbox written in both Python and Matlab is available on the personal webpage of B.M. for the interested reader: http://www.normalesup.org/~benmiquel/resources_en.html.

Acknowledgements

Appendix A. Convergence of the expansions

We discuss in this appendix the speed of convergence of the expansion of functions $\psi(y)$ on the $TS(y)$, and $ST(y)$ bases (eq. (30)). We first present tools to assess the speed of convergence of functions on the $[0, \pi]$ interval using the $\mathring{\beta}$ and $\check{\beta}$ bases (eq. (18)). We then discuss a few cases of functions on the infinite line, distinguished by their asymptotic behavior at $|y| \to \infty$.

A.1. Convergence of Fourier series for 2π-periodic functions

It is well-known (see e.g. [21]) that a 2π periodic function which is of $\mathscr{C}^{k-1}$ class (i.e. with a continuous $k-1$ derivative) on $\mathbb{R}$ and piecewise of $\mathscr{C}^k$ class on $\mathbb{R}$ possesses Fourier coefficients $\tilde{f}_n = o(1/n^k)$. This behavior is referred to as algebraic convergence. By extension, functions of $\mathscr{C}^\infty$ class on $\mathbb{R}$ possess Fourier coefficients that decay asymptotically faster than any power of $1/n$. Among the class of $\mathscr{C}^\infty$ functions on $\mathbb{R}$, the most coveted are analytic functions, the Fourier coefficients of which decay exponentially with $\tilde{f}_n \sim \exp(-an)$, with $a > 0$. Such a convergence is dubbed geometric convergence. Remapping infinite intervals onto the $[0, \pi]$ segment shall typically bring us to consider an intermediate class of functions: non-analytic $\mathscr{C}^\infty$ functions with essential singularities. For instance, $\exp(-\cot(\theta)^2)$ is a 2π-periodic function of class $\mathscr{C}^\infty$ on $\mathbb{R}$ with essential singularities at $\theta = n\pi$ (with $n \in \mathbb{Z}$). It has been shown in [22] that Fourier coefficients of $\mathscr{C}^\infty$ functions which possess a singularity of the form $\exp(-1/|\theta|^k)$ decay like $\tilde{f}_n \sim \exp(-an^{2k/[2k+1]})$, with $a > 0$.

A.2. Convergence of cosine, sine, or mixed trigonometric series on the $[0, \pi]$ segment

Let $f(\theta)$ be a function of $\mathscr{C}^\infty$ class on $[0, \pi]$, that satisfies $f(0) = f(\pi) = 0$. The behavior of the trigonometric series of f depends on symmetry considerations. To arm ourselves for manipulating such functions, we define the projectors on the symmetric and antisymmetric components of f with respect to the center of the interval $\pi/2$ (see Fig. 13, panels a and c):

$$\mathscr{S}_{\pi/2}[f(\theta)] = \frac{f(\theta) + f(\pi - \theta)}{2}, \tag{88a}$$

$$\mathscr{A}_{\pi/2}[f(\theta)] = \frac{f(\theta) - f(\pi - \theta)}{2} = f(\theta) - \mathscr{S}_{\pi/2}[f(\theta)]. \tag{88b}$$

Using these projectors, four distinct continuations of f to the interval $[0, 2\pi]$, with different symmetries, can be built. The four continuations (illustrated on Fig. 13, panels b and d) naturally coincide with f on the $[0, \pi]$ segment and are as defined as follows for $\theta \in [-\pi, 0]$:
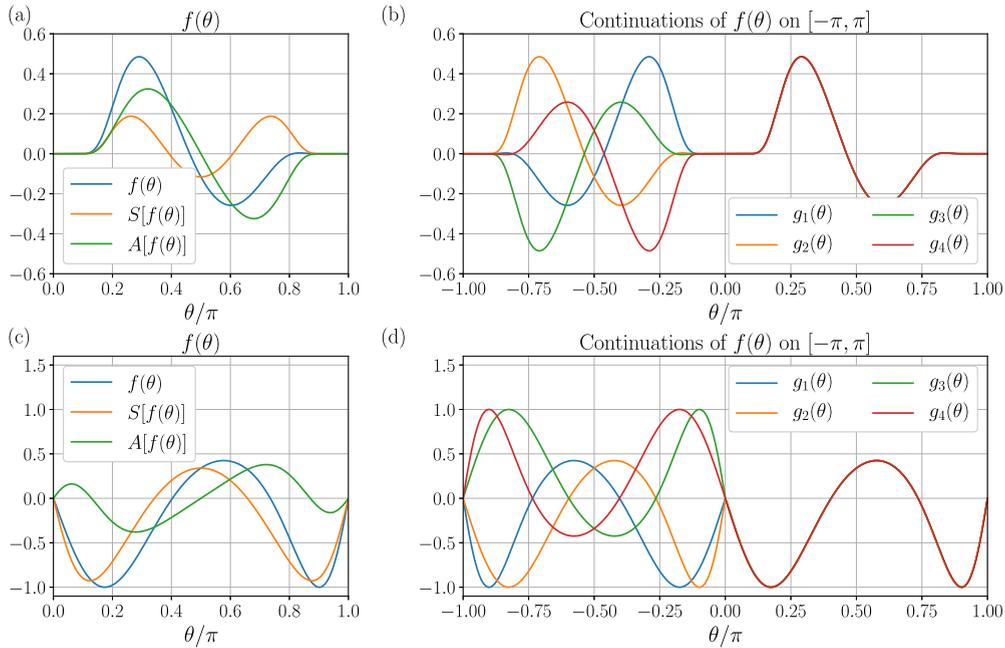
Fig. 13. Example of continuations for a function with essential singularities at the end of the interval (top row, panels (a) and (b)) and for an analytic function (bottom row, panels (c) and (d)). Left panels (a) and (c) show the decomposition $f(\theta) = \mathscr{S}[f(\theta)] + \mathscr{A}[f(\theta)]$ (see equation (88)). Right panels (b) and (d) show the four possible continuations (see equation (89)). Note that all the continuations are of $\mathscr{C}^\infty$ class in (b) whereas it is so for g_4 solely in panel (d): other continuations are piecewise of class $\mathscr{C}^1$, with a discontinuous slope at $\theta = 0, \pm\pi, \pm 2\pi, \dots$.

$$g_1(\theta) = +\mathscr{S}_{\pi/2}[f(-\theta)] + \mathscr{A}_{\pi/2}[f(-\theta)], \tag{89a}$$

$$g_2(\theta) = +\mathscr{S}_{\pi/2}[f(-\theta)] - \mathscr{A}_{\pi/2}[f(-\theta)], \tag{89b}$$

$$g_3(\theta) = -\mathscr{S}_{\pi/2}[f(-\theta)] + \mathscr{A}_{\pi/2}[f(-\theta)], \tag{89c}$$

$$g_4(\theta) = -\mathscr{S}_{\pi/2}[f(-\theta)] - \mathscr{A}_{\pi/2}[f(-\theta)]. \tag{89d}$$

Each of these continuations corresponds to one of our four choices of basis: $\{\cos(2p\theta)\}$ or $\{\sin([2p+1]\theta)\}$ for the symmetric component $\mathscr{S}_{\pi/2}[f(\theta)]$, and $\{\cos([2p+1]\theta)\}$ or $\{\sin(2p\theta)\}$ for the antisymmetric component $\mathscr{A}_{\pi/2}[f(\theta)]$. Inspecting the parity around $\theta = 0$, one easily identifies that g_1 corresponds to the pure cosine expansion, g_4 to the pure sine expansion, g_2 to the mixed $\check{\beta}$ expansion, and g_3 to the $\mathring{\beta}$ expansion:

$$g_1(\theta) = \sum_{n=0}^{N-1} \tilde{f}_n \cos(n\theta), \quad g_2(\theta) = \sum_{n=1}^{N} \check{f}_n \check{\beta}_n(\theta), \tag{90a}$$

$$g_3(\theta) = \sum_{n=1}^{N} \mathring{f}_n \mathring{\beta}_n(\theta), \quad g_4(\theta) = \sum_{n=1}^{N} \widehat{f}_n \sin(n\theta). \tag{90b}$$

The behavior of these four expansions, and notably their speed of convergence, is bound with the class of the continued function g_i they correspond to. We discuss several cases below.

Functions with essential singularities at 0 and π. Functions with essential singularities at the end of the $[0, \pi]$ segment (common in the context of the remapped infinite line) are advantageous and give complete flexibility concerning the choice of the basis. Indeed, consider f such that:

$$f(\varepsilon) \sim \exp\left(-\frac{1}{|\varepsilon|^k}\right), \quad \text{and} \quad f(\pi - \varepsilon) \sim \exp\left(-\frac{1}{|\varepsilon|^q}\right). \tag{91}$$

Each of the for continuations $g_1, \dots, g_4$ will be of class $\mathscr{C}^\infty$ on $\mathbb{R}$, with essential singularities. Therefore, the coefficients of each expansion (pure or mixed trigonometric) will converge at a subgeometric rate. In this case, the user has a total freedom of choice: mixed expansions will inherit the desirable properties of pure expansions.

Non-singular functions. Non-singular functions are ironically the least flexible as far as the choice of the basis is concerned. Let the Taylor expansions of f at $\theta = 0$ and $\theta = \pi$ be:

$$f(\varepsilon) = \sum_{n=0}^{\infty} a_n \varepsilon^n, \quad \text{and} \quad f(\pi - \varepsilon) = \sum_{n=0}^{\infty} b_n \varepsilon^n \,. \tag{92}$$

Our decomposition based on the symmetry around $\pi/2$ can be taken one step further, to distinguish between even and odd functions around 0. The function f can be decomposed into:

$$f(\theta) = \mathscr{SE}(\theta) + \mathscr{SO}(\theta) + \mathscr{AE}(\theta) + \mathscr{AO}(\theta) \tag{93}$$

with

$$\mathscr{SE}(\theta) = \sum_{p=0}^{\infty} \left(\frac{a_{2p} + b_{2p}}{2} \right) \theta^{2p}, \quad \mathscr{SO}(\theta) = \sum_{p=0}^{\infty} \left(\frac{a_{2p+1} + b_{2p+1}}{2} \right) \theta^{2p+1} \,, \tag{94a}$$

$$\mathscr{AE}(\theta) = \sum_{p=0}^{\infty} \left(\frac{a_{2p} - b_{2p}}{2} \right) \theta^{2p}, \quad \mathscr{AO}(\theta) = \sum_{p=0}^{\infty} \left(\frac{a_{2p+1} - b_{2p+1}}{2} \right) \theta^{2p+1} \,. \tag{94b}$$

As one attempts to continue f over $[-\pi, 0]$ by building g of the form:

$$g(-\theta) = \sigma_{\mathscr{S}} \left(\mathscr{SE}(\theta) + \mathscr{SO}(\theta) \right) + \sigma_{\mathscr{A}} \left(\mathscr{AE}(\theta) + \mathscr{AO}(\theta) \right) \,, \tag{95}$$

it becomes obvious that none of the four choices for $\sigma_{\mathscr{S}}, \sigma_{\mathscr{A}} = \pm 1$ will yield an analytic function unless $\mathscr{S}[f(\theta)]$ and $\mathscr{A}[f(\theta)]$ are purely odd or purely even at $\theta = 0$: some derivatives will be discontinuous at 0. Such functions $f(\theta)$ will invariably have algebraically decaying Fourier coefficients, regardless of the expansion chosen.

Functions such that $\mathscr{S}[f(\theta)]$ and $\mathscr{A}[f(\theta)]$ have a definite parity around zero have potential for being continued into a 2π-periodic analytic function on $\mathbb{R}$, but will offer little flexibility concerning the choice of basis that yields geometric convergence. As an illustration, consider f such that $f(\theta) = \mathscr{SE}(\theta) + \mathscr{AO}(\theta)$. Building $g_2(\theta) = +\mathscr{S}_{\pi/2}[f(-\theta)] - \mathscr{A}_{\pi/2}[f(-\theta)]$ is the sole way to continue f into an analytic function. As a consequence, a $\{\breve{\beta}\}$ expansion alone yields geometrically decaying coefficients.

A.3. Convergence of SB, TB, or mixed TS or ST bases on $(-\infty, \infty)$

By virtue of the discussion above, the speed of convergence of pure $\{TB\}$ and $\{SB\}$ or hybrid $\{TS\}$ and $\{ST\}$ Chebyshev expansions for a function $\psi(y)$ is controlled by the asymptotic behavior of ψ as $|y| \to \infty$. Functions $\psi(y)$ dominated at $|y| \to \infty$ by an exponentially decaying function $\exp(-b|y|^a)$ with $a, b > 0$, are mapped onto functions $f(\theta)$ with essential singularities at both ends of the $[0, \pi]$ segment. Indeed, upon remapping $y = L \cot \theta$ and choosing $L = 1$ for simplicity and without loss of generality, we find that $f(\theta) = \psi(y(\theta))$ is dominated around $\theta \downarrow 0$ by $\exp(-b/|\theta|^a)$ (by virtue of $\cot \theta = 1/\theta + O(\theta)$). In the light of our reasoning on the $[0, \pi]$ segment, the four mixed and hybrid expansions will converge subgeometrically.

Functions that decay algebraically $\psi(y) \sim 1/y^k$ for $|y| \to \infty$ will map to analytic functions on $[0, \pi]$. As stressed above, such a situation is unfortunate as it might not allow to choose freely the basis of expansion (unless one is to pay the price of algebraic convergence).

Appendix B. Orthogonality of the hybrid bases $\{\breve{\beta}_n\}$ and $\{\mathring{\beta}_n\}$.

Recalling from equation (17a) that:

$$\breve{\beta}_{2p} = \sin(2p\theta), \quad \text{and} \quad \breve{\beta}_{2p+1} = \cos(2p\theta) \,, \tag{96}$$

the orthogonality of the basis is immediate. We prove that $\left\langle \breve{\beta}_m, \breve{\beta}_n \right\rangle \propto \delta_m^n$, where δ_m^n is the Kronecker delta symbol, by remarking the following: we have $\left\langle \breve{\beta}_1, \breve{\beta}_1 \right\rangle = 1$, and for p or q not both equal to zero:

$$\left\langle \breve{\beta}_{2p}, \breve{\beta}_{2q} \right\rangle = \frac{1}{2} \delta_q^p, \quad \left\langle \breve{\beta}_{2p+1}, \breve{\beta}_{2q+1} \right\rangle = \frac{1}{2} \delta_q^p, \quad \left\langle \breve{\beta}_{2p}, \breve{\beta}_{2q+1} \right\rangle = 0, \quad \left\langle \breve{\beta}_{2p+1}, \breve{\beta}_{2q} \right\rangle = 0 \,. \tag{97}$$

Similarly, using equation (17b):

$$\mathring{\beta}_{2p-1} = \cos([2p-1]\theta), \quad \text{and} \quad \mathring{\beta}_{2p} = \sin([2p-1]\theta) \,, \tag{98}$$

one immediately gets the orthogonality $\langle \mathring{\beta}_m, \mathring{\beta}_n \rangle = \frac{1}{2} \delta_m^n$:

$$\langle \mathring{\beta}_{2p}, \mathring{\beta}_{2q} \rangle = \frac{1}{2} \delta_q^p, \quad \langle \mathring{\beta}_{2p+1}, \mathring{\beta}_{2q+1} \rangle = \frac{1}{2} \delta_q^p, \quad \langle \mathring{\beta}_{2p}, \mathring{\beta}_{2q+1} \rangle = 0, \quad \langle \mathring{\beta}_{2p+1}, \mathring{\beta}_{2q} \rangle = 0 \,. \tag{99}$$

Appendix C. Matrix elements

C.1. Discretization along an unbounded direction: Chebyshev functions

This appendix documents a method to obtain discretized matrices for operators composed of superpositions of powers of cosines, sines, and derivatives. Such operators are typically obtained after the procedure described in 4.2 and exemplified in 4.3. We will consider the four matrices $\mathcal{S}^{sc}$, $\mathcal{S}^{cs}$, $\mathcal{C}^{cc}$ and $\mathcal{C}^{ss}$ as our elementary building blocks. These matrices represent the action of $\sin\theta$ on cosines and sines, and the action of $\cos\theta$ on cosines and sines, respectively. Their matrix elements are given by:

$$\mathcal{S}^{sc}_{mn} = \frac{\langle \sin(m\theta),\, \sin\theta\cos([n-1]\theta)\rangle}{\|\sin(m\theta)\|^2} = \begin{cases} \delta^1_m & \text{for } n=1 \\ \frac{1}{2}\left(\delta^n_m - \delta^{n-2}_m\right) & \text{for } n \geq 2 \end{cases} \tag{100a}$$

$$\mathcal{S}^{cs}_{mn} = \frac{\langle \cos([m-1]\theta),\, \sin\theta\sin(n\theta)\rangle}{\|\cos([m-1]\theta)\|^2} = \frac{1}{2}\left(-\delta^{n+1}_{m-1} + \delta^{n-1}_{m-1}\right) \tag{100b}$$

$$\mathcal{C}^{cc}_{mn} = \frac{\langle \cos([m-1]\theta),\, \cos\theta\cos([n-1]\theta)\rangle}{\|\cos([m-1]\theta)\|^2} = \begin{cases} \delta^2_m & \text{for } n=1 \\ \frac{1}{2}\left(\delta^n_{m-1} + \delta^{n-2}_{m-1}\right) & \text{for } n \geq 2 \end{cases} \tag{100c}$$

$$\mathcal{C}^{ss}_{mn} = \frac{\langle \sin(m\theta),\, \cos\theta\sin(n\theta)\rangle}{\|\sin(m\theta)\|^2} = \frac{1}{2}\left(\delta^{n+1}_m + \delta^{n-1}_m\right) \tag{100d}$$

We also define differentiation matrices:

$$\mathcal{D}^{sc} = \frac{\langle \sin(m\theta),\, \partial_\theta\cos([n-1]\theta)\rangle}{\|\sin(m\theta)\|^2} = -(n-1)\,\delta^{n-1}_m\,, \tag{100e}$$

$$\mathcal{D}^{cs} = \frac{\langle \cos([m-1]\theta),\, \partial_\theta\sin(n\theta)\rangle}{\|\sin([m-1]\theta)\|^2} = n\,\delta^n_{m-1}\,. \tag{100f}$$

From these building blocks, more complex operators are readily constructed using matrix products, with the only reservation that the lower right corner of the matrix, which corresponds to modes neighboring the truncation, deserves specific attention. Indeed, should no care be taken, the lower right corner of the matrix product would be inexact, due to the proximity of truncation. A rule of thumb is that, when two matrices $\mathbf{L}_1$ and $\mathbf{L}_2$ represent the action of operators $\mathscr{L}_1$ and $\mathscr{L}_2$ and have a bandwidth $(1+2p)$ (with p subdiagonals and p superdiagonals), the action of the operator $\mathscr{L}_3 = \mathscr{L}_1\mathscr{L}_2$ is *not* represented by the product $\mathbf{L}_1\mathbf{L}_2$, for the last p columns of this product are spurious. This problem can be understood as some kind of aliasing related to linear operators and is easily circumvented by using a larger truncation $N' > N$ to discretize $\mathscr{L}_1$ and $\mathscr{L}_2$, so that to obtain large square matrices $\mathbf{L}'_1$ and $\mathbf{L}'_2$ of size N'. Then, one computes the product $\mathbf{L}'_3 = \mathbf{L}'_1\mathbf{L}'_2$ of size N', and finally one extracts a smaller square matrix $\mathbf{L}_3$ of size N that rigorously discretize $\mathscr{L}_3$. Below we denote with an overbar $\overline{\mathbf{L}}^N$ the extraction from a given square matrix $\mathbf{L}$ of a square submatrix of size N. Using our prescription, we can safely use a matrix product to obtain the matrices $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, $\mathbf{D}$ and $\mathbf{E}$, defined in equation (43), as products of elementary matrices $\mathcal{E}$ of suitably enlarged size $N' > N$. We distinguish below the case of parity-conserving and parity-flipping operators.

The case of parity-conserving operators. Our prescription is first illustrated with the parity conserving operator $\mathscr{D} = \sin^2\theta$, in view of obtaining its discretization $\mathbf{D}$ on the basis $\{\breve{\beta}_n\}$, as defined in equation (43d). Recalling that this basis contains interleaved even harmonics of sines and cosines [see equations (17a) and (18a)]:

$$\left\{\breve{\beta}_n\right\}_{1\leq n\leq N} = \{1,\, \sin 2\theta,\, \cos 2\theta,\, \sin 4\theta,\, \cos 4\theta,\, \dots\}\,, \tag{101}$$

we easily deduce that even elements $\breve{\beta}_{2p} = \sin(2p\theta)$ are antisymmetric functions whereas odd elements $\breve{\beta}_{2p-1} = \cos([2p-2]\theta)$ are symmetric functions. The symmetry of the operator $\mathscr{D}$ ensures us that this operator cannot couple symmetric and antisymmetric functions. Hence we only compute the coupling among (symmetric) cosines $\mathbf{D}^c$ and among (antisymmetric) sines $\mathbf{D}^s$ as products of elementary matrices $\mathcal{E}$ of size $N' \geq N+1$:

$$\mathbf{D}^c = \frac{\left\langle \cos([m-1]\theta),\, \sin^2\theta\cos([n-1]\theta)\right\rangle}{\|\cos([m-1]\theta)\|^2} = \overline{\mathcal{S}^{cs}\mathcal{S}^{sc}}^N\,, \tag{102a}$$

$$\mathbf{D}^s = \frac{\left\langle \sin(m\theta),\, \sin^2\theta\sin(n\theta)\right\rangle}{\|\sin(m\theta)\|^2} = \overline{\mathcal{S}^{sc}\mathcal{S}^{cs}}^N\,. \tag{102b}$$

Finally, we obtain the expression of $\mathbf{D}$, the discretization of $\mathscr{D} = \sin^2\theta$ on the basis $\{\breve{\beta}_n\}$, by interweaving adequately the matrices $\mathbf{D}^c$ and $\mathbf{D}^s$:

$$\mathbf{D}_{2p-1,2q-1} = \mathbf{D}^c_{2p-1,2q-1} \,, \tag{103a}$$

$$\mathbf{D}_{2p,2q} = \mathbf{D}^s_{2p,2q} \,. \tag{103b}$$

We indicate below the construction by matrix products for matrices of operators $\mathscr{A} = \sin^6 \partial^2_{\theta\theta}$, $\mathscr{B} = \cos\theta \sin^5\theta \partial_\theta$, and $\mathscr{C} = \cos^2\theta$. These parity conserving operators are all treated in a similar fashion as $\mathscr{D}$: as justified above, we only compute coupling among cosines and sines:

$$\mathbf{A}^c = \frac{\langle \cos([m-1]\theta), \, \mathscr{A}\cos([n-1]\theta)\rangle}{\|\cos([m-1]\theta)\|^2} = \overline{\mathcal{S}^{cs}\mathcal{S}^{sc}\mathcal{S}^{cs}\mathcal{S}^{sc}\mathcal{S}^{cs}\mathcal{S}^{sc}\mathcal{D}^{cs}\mathcal{D}^{sc}}^N \,, \tag{104a}$$

$$\mathbf{A}^s = \frac{\langle \sin(m\theta), \, \mathscr{A}\sin(n\theta)\rangle}{\|\sin(m\theta)\|^2} = \overline{\mathcal{S}^{sc}\mathcal{S}^{cs}\mathcal{S}^{sc}\mathcal{S}^{cs}\mathcal{S}^{sc}\mathcal{S}^{cs}\mathcal{D}^{sc}\mathcal{D}^{cs}}^N \,, \tag{104b}$$

$$\mathbf{B}^c = \frac{\langle \cos([m-1]\theta), \, \mathscr{B}\cos([n-1]\theta)\rangle}{\|\cos([m-1]\theta)\|^2} = \overline{\mathcal{C}^{cc}\mathcal{S}^{cs}\mathcal{S}^{sc}\mathcal{S}^{cs}\mathcal{S}^{sc}\mathcal{S}^{cs}\mathcal{D}^{sc}}^N \,, \tag{104c}$$

$$\mathbf{B}^s = \frac{\langle \sin(m\theta), \, \mathscr{B}\sin(n\theta)\rangle}{\|\sin(m\theta)\|^2} = \overline{\mathcal{C}^{ss}\mathcal{S}^{sc}\mathcal{S}^{cs}\mathcal{S}^{sc}\mathcal{S}^{cs}\mathcal{S}^{sc}\mathcal{D}^{cs}}^N \,. \tag{104d}$$

The final interweaving to obtain **A** and **B** is identical as in the case of **D**, given in equations (103). The matrix **C** is directly obtained from the identity $\cos^2\theta = 1 - \sin^2\theta$. Thus, denoting **1** the identity matrix:

$$\mathbf{C} = \mathbf{1} - \mathbf{D}\,. \tag{105}$$

The case of parity-flipping operators. Finally, we illustrate the treatment of a parity breaking term by discretizing the operator $\mathscr{E} = \cos\theta \sin\theta$. The parity-flipping nature of $\mathscr{E}$ guarantees that, in contrast with parity-conserving operators, this operator does not couple elements of the basis with the same parity. Instead, the coupling occurs between sines and cosines, leading us to compute $\mathbf{E}^{sc}$ and $\mathbf{E}^{cs}$:

$$\mathbf{E}^{sc} = \frac{\langle \sin(m\theta), \, \mathscr{E}\cos([n-1]\theta)\rangle}{\|\sin(m\theta)\|^2} = \overline{\mathcal{C}^{ss}\mathcal{S}^{sc}}^N \,, \tag{106a}$$

$$\mathbf{E}^{sc} = \frac{\langle \cos([m-1]\theta), \, \mathscr{E}\sin(n\theta)\rangle}{\|\cos([m-1]\theta)\|^2} = \overline{\mathcal{C}^{cc}\mathcal{S}^{cs}}^N \,. \tag{106b}$$

Finally, the matrix **E** defined in equation (43e) is obtained by the following interweaving, that differs from the parity-conserving case:

$$\mathbf{E}_{2p,2q-1} = \mathbf{E}^{sc}_{2p,2q-1} \,, \tag{107a}$$

$$\mathbf{E}_{2p-1,2q} = \mathbf{E}^{sc}_{2p-1,2q} \,. \tag{107b}$$

C.2. Discretization along a bounded direction: Chebyshev polynomials and the quasi-inverse technique

Central to the quasi-inverse technique is the three-term relation for the antiderivative of Chebyshev polynomials, first emphasized in the context of numerical resolution of ODEs by Clenshaw [23]:

$$\int T_0(x)\mathrm{d}x = T_1(x) + k\,, \tag{108a}$$

$$\int T_1(x)\mathrm{d}x = \frac{1}{4}T_2(x) + k\,, \tag{108b}$$

$$\int T_n(x)\mathrm{d}x = \frac{1}{2}\left(\frac{T_{n+1}(x)}{n+1} + \frac{T_{n-1}(x)}{n-1}\right) + k, \text{ for } n \geq 2\,. \tag{108c}$$

Upon integration, an arbitrary constant k naturally appears. This constant will be set by enforcing boundary conditions: the consequence is that, for now, the top row of the matrix representing integration can be left equal to zero.

$$\mathbf{I}_{mn} = \frac{\langle T_{m-1}, \, \int T_{n-1}\rangle}{\|T_{m-1}\|^2} = \begin{cases} \delta^2_m \text{ for } n=1\,, \\ \dfrac{1}{4}\delta^3_m \text{ for } n=2\,, \\ \dfrac{1}{2m}\delta^{n+1}_m - \dfrac{1}{2m}\delta^{n-1}_m\,. \end{cases} \tag{109}$$

Another handy three-term relationship corresponds to the multiplication by x:

$$xT_0(x) = T_1(x)\,, \tag{110a}$$

$$xT_n(x) = \frac{1}{2}\big(T_{n+1}(x) + T_{n-1}(x)\big),\ \text{for } n \geq 1\,, \tag{110b}$$

which has the discretized matrix:

$$\mathbf{x}_{mn} = \frac{\langle T_{m-1}, xT_{n-1}\rangle}{\|T_{m-1}\|^2} = \begin{cases} \delta_m^2 & \text{for } n = 1\,, \\ \frac{1}{2}\delta_m^{n+1} + \frac{1}{2}\delta_m^{n-1}\,. \end{cases} \tag{111}$$

Hence the matrices $\mathbf{Q}$ defined in equations (76a)–(76c) are again obtained by truncated matrix products of large matrices:

$$\mathbf{Q}^{(0)} = \overline{\mathbf{I2}}^N\,, \tag{112a}$$

$$\mathbf{Q}^{(1)} = \overline{\mathbf{I2x}}^N\,, \tag{112b}$$

$$\mathbf{Q}^{(2)} = \overline{\mathbf{I2x2}}^N\,. \tag{112c}$$

References

[1] John P. Boyd, Chebyshev and Fourier Spectral Methods, second edition, Dover, New York, 2001.
[2] Bengt Fornberg, A Practical Guide to Pseudospectral Methods, Cambridge University Press, Cambridge, 1996.
[3] D. Gottlieb, S.A. Orszag, Numerical Analysis of Spectral Methods: Theory and Applications, CBMS-NSF Reg. Conf. Ser. Appl. Math., Society for Industrial and Applied Mathematics, 1977.
[4] John P. Boyd, Spectral methods using rational basis functions on an infinite interval, J. Comput. Phys. 69 (1) (1987) 112–142.
[5] L.F. Richardson, On the approximate arithmetical solution by finite differences of physical problems involving differential equations, with an application to the stresses in a masonry dam, Proc. R. Soc., Math. Phys. Eng. Sci. 83 (563) (1910) 335–336.
[6] Randall J. LeVeque, Finite Volume Methods for Hyperbolic Problems, Cambridge University Press, 2002.
[7] Anthony T. Patera, A spectral element method for fluid dynamics: laminar flow in a channel expansion, J. Comput. Phys. 54 (3) (1984) 468–488.
[8] Uri M. Ascher, Steven J. Ruuth, Raymond J. Spiteri, Implicit–explicit Runge–Kutta methods for time-dependent partial differential equations, Appl. Numer. Math. 25 (2–3) (November 1997) 151–167.
[9] Ian Grooms, Keith Julien, Linearly implicit methods for nonlinear PDEs with linear dispersion and dissipation, J. Comput. Phys. 230 (9) (2011) 3630–3650.
[10] Daniele Cavaglieri, Thomas Bewley, Low-storage implicit/explicit Runge–Kutta schemes for the simulation of stiff high-dimensional ODE systems, J. Comput. Phys. 286 (2015) 172–193.
[11] J.W. Cooley, J. Tukey, An algorithm for the machine calculation of complex Fourier series, Math. Comput. 19 (1965) 297–301.
[12] Carl Friedrich Gauss, Theoria Interpolationis Methodo Nova Tractata, Werke, vol. 3, K. Gesellschaft der Wissenschaften zu Göttingen, 1870, p. 524.
[13] Geoffrey M. Vasil, Nicholas H. Brummell, Keith Julien, A new method for fast transforms in parity-mixed PDEs: Part I. Numerical techniques and analysis, J. Comput. Phys. 227 (17) (2008) 7999–8016.
[14] A.B. Cain, J.H. Ferziger, W.C. Reynolds, Discrete orthogonal function expansions for non-uniform grids using the fast Fourier transform, J. Comput. Phys. 56 (2) (1984) 272–286.
[15] Steven A. Orszag, Fast eigenfunction transforms, in: Science and Computers, Academic Press, New York, 1986, pp. 23–30.
[16] John P. Boyd, Multipole expansions and pseudospectral cardinal functions: a new generalization of the fast Fourier transform, J. Comput. Phys. 103 (1) (1992) 184–186.
[17] A.E. Gill, Atmosphere-ocean Dynamics, Int. Geophys. Ser., vol. 30, Academic Press, 1982.
[18] Wilhelm Heinrichs, Improved condition number for spectral methods, Math. Comput. 53 (187) (1989) 103–119.
[19] Keith Julien, Mike Watson, Efficient multi-dimensional solution of PDEs using Chebyshev spectral methods, J. Comput. Phys. 228 (5) (2009) 1480–1503.
[20] P.G. Drazin, W.H. Reid, Hydrodynamic Stability, 2nd edition, Cambridge University Press, Cambridge, 2004, p. 007.
[21] Walter Appel, Mathematics for Physics and Physicists, Princeton University Press, 2007.
[22] John P. Boyd, The optimization of convergence for Chebyshev polynomial methods in an unbounded domain, J. Comput. Phys. 45 (1) (1982) 43–79.
[23] C.W. Clenshaw, The numerical solution of linear differential equations in Chebyshev series, Math. Proc. Camb. Philos. Soc. 53 (1) (1957) 134–149.