

A hybrid Fourier–Chebyshev method for partial differential equations

Rodrigo B. Platte* Anne Gelb†

November 19, 2008

Abstract

We propose a pseudospectral hybrid algorithm to approximate the solution of partial differential equations (PDEs) with non-periodic boundary conditions. Most of the approximations are computed using Fourier expansions that can be efficiently obtained by fast Fourier transforms. To avoid the Gibbs phenomenon, super-Gaussian window functions are used in physical space. Near the boundaries, we use local polynomial approximations to correct the solution. We analyze the accuracy and eigenvalue stability of the method for several PDEs. The method compares favorably to traditional spectral methods, and numerical results indicate that for hyperbolic problems a time step restriction of $O(1/N)$ is sufficient for stability.

Keywords. Hybrid methods, exponential convergence, Fourier spectral method, non-periodic boundary conditions, time-dependent PDEs

AMS subject classification. 65M70 - 65M12 - 65T40

1 Introduction

Spectral methods have been widely used in the numerical solution of partial differential equations, [2, 10, 14, 24]. These methods are known for their exponential convergence and are popular in the study of turbulent flows and meteorological simulations, among other applications. Fourier spectral methods, in particular, produce outstanding results for smooth problems with periodic boundary conditions. For non-periodic problems, Chebyshev spectral methods are commonly used to obtain exponential rates of convergence. However, while Fourier methods require two points per wavelength to resolve a function,

*Oxford University Computing Laboratory, Wolfson Bldg., Parks Road, Oxford, OX1 3QD, United Kingdom (rod@comlab.ox.ac.uk). After December 2009: Arizona State University, Department of Mathematics and Statistics, Tempe, AZ, 85287-1804.

†Arizona State University, Department of Mathematics and Statistics, Tempe, AZ, 85287-1804 (gelb@math.la.asu.edu).

Chebyshev methods require π , [17, 27]. Moreover, the clustering of nodes near the boundaries required by polynomial methods impose severe restrictions on the time step size when these methods are used for time dependent problems.

The purpose of this article is to obtain better approximation properties for non-periodic problems. We propose a hybrid method that relies on the Fourier partial sum expansion in most of the domain except near the boundaries, where polynomials are used. In order to avoid the Gibbs phenomenon, we multiply the function to be approximated, say u , by a smooth window function w , where w and its derivatives are close to zero at boundary points. The product uw can be accurately determined by a truncated trigonometric series, and then u can be obtained by dividing the Fourier expansion of uw by w . Since w is close to zero near the boundaries, polynomials are used to correct the approximations in these regions. The advantage of this hybrid strategy is that polynomial approximations are required only at a very small part of the domain. Hence we can take advantage of the superior approximation properties of Fourier spectral methods even though the underlying function may not be periodic. This process is illustrated in Figure 1 for $u(x) = \exp(x/\pi) + \exp(-50x^2)$. In this case the window function $w(x) = \exp(-32(x/\pi)^{10})$ was used.

Although this algorithm is simple to describe and implement, choosing an effective window function, size of the correction regions and other parameters, requires a detailed investigation. In this article we use theoretical insight and numerical experiments to provide guidelines for choosing these parameters. We then show that this scheme is often more efficient than standard polynomial spectral methods, in particular when approximating functions that require fine grids. Eigenvalue stability for time dependent PDEs indicate that a time step size proportional to $1/N$ is possible for hyperbolic problems. Several numerical experiments are used to demonstrate the performance of the method, including the simulation of a nonlinear incompressible fluid flow.

We point out that several alternatives to traditional spectral methods have been proposed recently. Among them are mapped polynomial methods, in which quadrature nodes are mapped to more evenly spaced grids, [16, 18], radial basis function schemes, [8], prolate spheroidal basis techniques, [3], and Fourier extension/embedding [4]. Fourier embedding is closely related to the present work, with the distinction that the hybrid scheme does not extend the function outside its interval of definition. In [7] Boyd describes seven strategies to defeat the Runge phenomenon – oscillations near endpoints of interpolating polynomials on equispaced grids. Unfortunately, each of these strategies has liabilities such as ill-conditioning and lack of theoretical justification. In addition, the convergence rates are usually subgeometric even when the underlying function is analytic.

Hybrid methods have been proposed in [6] and in [12]. In [6], Boyd proposes a strategy that uses trigonometric interpolation with filtering and polynomial interpolation near the boundaries to recover functions from discrete data. His method is implemented on equispaced nodes, including the correction near the endpoints. As a consequence, convergence is $O(\exp(-1/\sqrt{N}))$. Moreover, because differentiation matrices based on polynomial interpolation on equally

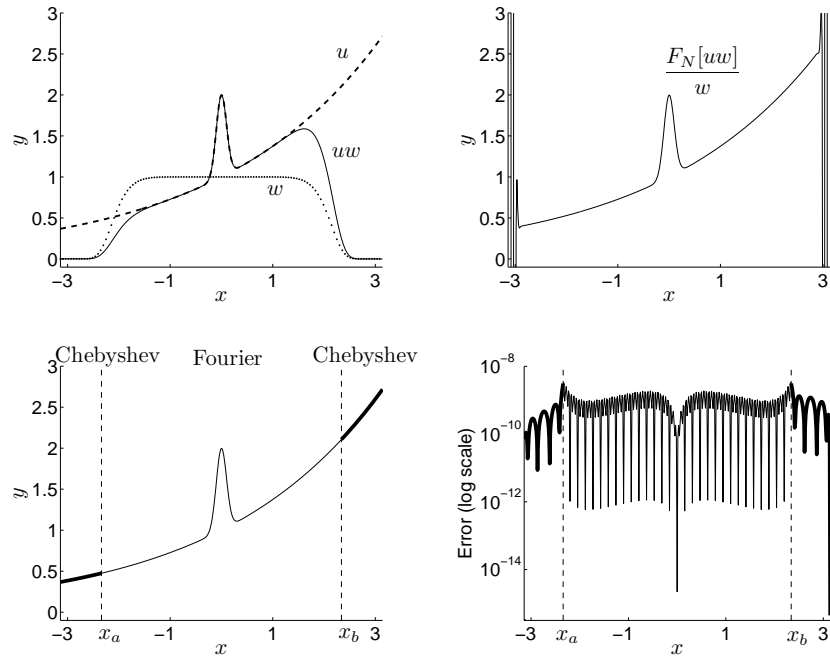


Figure 1: A hybrid approximation of $u(x) = \exp(x/\pi) + \exp(-50x^2)$. Top-left: plots of u , w , and uw . Top-right: Fourier approximation of uw divided by w . Bottom-left: approximation near the ends is replaced by local polynomial approximation. Bottom-right: plot of the error.

spaced nodes have eigenvalues with positive real part, it is unclear whether this approach can be implemented to solve time dependent problems.

In [12], motivated by [11, 13, 23], we presented an algorithm that relies on weighted least-squares approximations for the solutions of PDEs. To circumvent Runge oscillations, super-Gaussian weights were used. Because the weight vanishes near the ends of the interval, a similar polynomial correction was employed. In contrast, the hybrid scheme presented here uses Fourier interpolation rather than polynomial least-square approximations, and as a result, we are able to make effective use of fast Fourier transforms and provide a better analysis of the accuracy and stability of the method. We also anticipate that a more evenly spaced point distribution throughout the domain may improve the resolution requirements for many PDE applications. In this article, we loosely use the term ‘‘Chebyshev method’’ to refer to schemes based on polynomial interpolation at Chebyshev extreme points.

This article is organized as follows. In the next section we motivate our choice of window function. In Section 3 we derive guidelines for implementing the method based on accuracy considerations. In Section 4 we derive differentiation matrices and use eigenvalue analysis to show that the method is suitable for time dependent PDEs. Examples of numerical solution of PDEs are presented in Section 5. Our conclusions and final remarks are provided in Section 6.

2 Fourier approximations of non-periodic functions

In this section we consider the truncated pseudospectral Fourier series,

$$F_N[u](x) = \sum_{n=-N}^N \prime \hat{u}_n \exp(inx), \quad -\pi \leq x \leq \pi, \quad (1)$$

where the coefficients \hat{u}_n are computed so that $u(x_j) = F_N[u](x_j)$ at $2N$ equally spaced nodes. The prime indicates that the terms $n = \pm N$ are multiplied by $\frac{1}{2}$, [22, 24]. This series converges exponentially fast, as $N \rightarrow \infty$, to smooth periodic functions (see e.g. [2, 14, 24]). When the target function u is not periodic, however, (1) will yield oscillations near the boundaries. Furthermore, its global construction means that convergence rates are extremely slow even away from the boundaries, [17]. This behavior is known as the Gibbs phenomenon.

A well known strategy to improve Fourier approximations of non-periodic functions away from the boundaries is the use of filters, [17, 23]. In this approach, (1) is replaced with

$$F_N^\sigma[u](x) = \sum_{n=-N}^N \prime \sigma(n/N) \hat{u}_n \exp(inx), \quad -\pi \leq x \leq \pi. \quad (2)$$

where σ is the filter function. A common choice for σ is the exponential filter

$$\sigma(k) = \exp(-\alpha(k/N)^p), \quad (3)$$




Figure 2: Super-Gaussian window functions: $\exp(-32(x/\pi)^{2\lambda})$.

where α is positive and p is even. An inherent characteristic of filtering highly oscillatory modes is the loss of resolution since the number of “effective” modes is decreased.

An alternative to filtering Fourier coefficients is to multiply u by a periodic function w so that wu and its derivatives become periodic to machine precision. In general, this implies that $w(\pm\pi) \approx w'(\pm\pi) \approx w''(\pm\pi) \approx \dots \approx 0$. The idea behind this approach is that if $F_N[wu]$ is a good approximation to wu , then

$$G_N^w[u](x) = \frac{1}{w(x)} F_N[wu](x) \quad (4)$$

is also a good approximation to $u(x)$ as long as $w(x)$ is not “close to zero”, e.g. $|w(x)| > 0.1$. In this article, we study the family of functions

$$w(x) = \exp(-\alpha(x/\pi)^{2\lambda}), \quad (5)$$

where λ is a positive integer. We call (5) the super-Gaussian window functions. In essence we are using the window (3) in physical rather than frequency space. In double precision, $\alpha = 32$ is a suitable choice. One may choose other values of α based on the desired accuracy, however.

The difficulty in finding a suitable window is that a nonzero function that has all derivatives vanishing at a point cannot be analytic on any neighborhood of the interval in consideration. Therefore, there must be a compromise between enforcing periodicity and the stiffness of the product uw . This aspect is investigated in more detail in Section 3. Figure 2 shows super-Gaussian window functions for $\lambda = 2, 5$, and 20 . Notice that $\lambda = 20$ gives large support for the Fourier approximation but also stiff gradients near the boundaries.

Figure 3 shows the error in the approximation of $u(x) = \exp(x/\pi) + \exp(-50x^2)$ with $N = 100$, using a filtered Fourier approximation and the Fourier approximation of uw . Notice that the windowed approach produces better results than




Figure 3: Error in the numerical approximation of $u(x) = \exp(x/\pi) + \exp(-50x^2)$ with $N = 100$. Left: Filtered Fourier approximation (exponential filter was used, $F_N^\sigma[u]$, with $\alpha = -32$). Right: Fourier approximation using window functions ($G_N^w[u]$).

the filtered approximation except at a neighborhood of $x = \pm\pi$. This figure also illustrates that a suitable choice of window, in this case determined by λ , is dependent on the number of nodes being used. Notice that $\lambda = 8$ works better than $\lambda = 2$ or $\lambda = 12$, because the former has small support and the latter introduces large gradients that are not resolved to machine precision with $N = 100$.

Clearly there are other possible choices of window functions. One could, for instance, choose $C^\infty[-\pi, \pi]$ functions that vanish exactly in value and derivatives at the endpoints. One example of this class of functions is $\exp(-c_1/(\pi - x^2)^{c_2})$, [23], where c_1 and c_2 are positive constants. Notice, however, that these functions are singular at $\pm\pi$. Boyd, in [5], also proposed window functions generated with the error function (erf) for the Fourier embedding in one dimension. Although we have not made extensive comparisons, we found that the super-Gaussian window, (5), gives favorable or similar results when compared to these options.

3 Accuracy and the choice of parameters

The accuracy and performance of the method depends mostly on how well the product uw is approximated. Assuming that the product is at least as stiff as the window w , we first study the convergence of the trigonometric polynomial that interpolates w .

3.1 Accuracy of the Fourier approximation of w

Figure 4 shows the number of Fourier modes required to approximate w to a specified accuracy ε for several values of λ . A linear regression on this data is




Figure 4: Number of nodes required to approximate $w(x) = e^{-32(x/\pi)^{2\lambda}}$ for several values of λ . The tolerance used to obtain this data was $\|F_N(w) - w\|_\infty \leq \varepsilon$. The lines are connected by the linear fit $N = C_\varepsilon \lambda$ (see table 1).

ε	linear fit $N = C_\varepsilon \lambda$ (Fig. 4)
10^{-4}	$N = 4.6\lambda$
10^{-6}	$N = 6.5\lambda$
10^{-8}	$N = 8.3\lambda$
10^{-10}	$N = 10.1\lambda$
10^{-12}	$N = 12.1\lambda$

Table 1: Choice of parameters required to approximate w to accuracy ε , i.e., $\|F_N[w] - w\|_\infty < \varepsilon$. The linear fit was found using the data presented in Fig 4.

shown in the figure and is presented in Table 1. Notice that the relationship is approximately linear in all cases. In particular, approximations accurate to almost machine precision can be obtained with $N > 12\lambda$.

This linear dependence can be explained using standard error estimates for Fourier interpolation based on the analytic properties of functions on the complex plane. If a periodic function f can be extended to an analytic function in the complex strip $|\text{Im } z| < \eta$, $z = x + iy$, with $\|f(z)\| < M$ uniformly in this strip, then we have (see e.g. [22, 24])

$$\|f - F_N[f]\|_\infty = O\left(Me^{-N(\eta-\epsilon)}\right) \text{ as } N \rightarrow \infty, \text{ for every } \epsilon > 0. \quad (6)$$

If we overlook the fact that super-Gaussians defined in (5) are not exactly periodic in all derivatives, we may use this result to estimate the rate of decay




Figure 5: Contour plot of the absolute value of $w(z) = \exp(-32(z/\pi)^{2\lambda})$ in the complex plane for $\lambda = 2$ (left) and $\lambda = 10$ (right). In the white region $|w| < 10$, gray region $10 < |w| < 10^{15}$, and black region $10^{15} < |w|$. The dashed lines mark the strip given by Lemma 3.1.

of $\|w - F_N[w]\|_\infty$. Notice that although super-Gaussians are entire functions, in absolute value, they go to infinity very quickly as z moves away from the interval $[-\pi, \pi]$. This is illustrated in Figure 5, where contours of $|w(z)|$ are presented for $\lambda = 2$ and $\lambda = 10$. In order to use (6), therefore, we must find a strip around the real line where $|w(z)|$ is bounded (by a not very large number M). To this end we state the following:

Lemma 3.1. *Suppose $\lambda > 0$ and $\alpha > \ln 10$. Then $|\exp(-\alpha(z/\pi)^{2\lambda})| < 10$ for all z inside the strip $|\operatorname{Im} z| < \pi^2/(4\lambda)$.*

Proof. Let (r, θ) be the polar coordinates of z , $r = |z|$ and $\theta = \arg(z)$, then

$$|\exp(-\alpha(z/\pi)^{2\lambda})| = \exp(-\alpha(r/\pi)^{2\lambda} \cos(2\lambda\theta)).$$

and $|w(z)| < 10$ whenever $-\alpha(r/\pi)^{2\lambda} \cos(2\lambda\theta) < \ln 10$. This inequality is satisfied if

$$|\theta| \leq \frac{\pi}{4\lambda} \quad \text{or} \quad r < \pi \left(\frac{\ln 10}{\alpha} \right)^{\frac{1}{2\lambda}}.$$

Since $|\operatorname{Im} z| = r |\sin \theta|$, either condition is fulfilled if

$$|\operatorname{Im} z| < \pi \left(\frac{\ln 10}{\alpha} \right)^{\frac{1}{2\lambda}} \sin \left(\frac{\pi}{4\lambda} \right) < \frac{\pi^2}{4\lambda}.$$

□

Using the result in Lemma 3.1, the estimate (6) gives

$$\|w - F_N[w]\|_\infty = O \left(10 \exp \left(- \left(\frac{N\pi^2}{4\lambda} - \epsilon \right) \right) \right).$$

Notice that choosing $N = 12\lambda$ yields $10 \exp\left(-\frac{N\pi^2}{4\lambda}\right) \approx 10^{-12}$, which is in good agreement with the experimental results in Figure 4.

3.2 The resolution needed to approximate uw

Now that we have established that we need about 24λ points to approximate w close to machine precision, we consider the Fourier interpolation of the product uw . For simplicity, we take $u(x) = \cos(kx)$ and write w in terms of its Fourier approximation. Then

$$uw \approx \frac{\exp(ikx) + \exp(-ikx)}{2} \sum_{n=-N_w}^{N_w} \hat{w}_n \exp(inx) = \sum_{n=-N_w-k}^{N_w+k} \hat{v}_n \exp(inx), \quad (7)$$

where $N_w \approx 12\lambda$ and \hat{v}_n are the Fourier coefficients of uw .

The second sum indicates that $2(N_w + k)/k$ points per wavelength of u are needed to resolve uw . If we choose λ such that $N_w = k$, we have that machine (double) precision accuracy is achieved with four points per wavelength. This implies that $N = 24\lambda$ is needed. In other words, although $N = 12\lambda$ suffices to approximate w to machine precision, (7) indicates that twice as many points for a given λ are needed to resolve the product uw to machine accuracy and avoid aliasing. Notice that

$$|\hat{v}_{|n|}| = \frac{1}{2} |\hat{w}_{|n|-k}|, \quad |n| \geq k,$$

and because the magnitude of the Fourier coefficients of w decay geometrically, convergence starts as soon as the minimum requirement of two points per wavelength is satisfied.

Figure 6 shows the error in the Fourier interpolation of uw , where $k = 60$ and $N \approx 24\lambda$ were used. Notice that exponential convergence starts at two points per wavelength of u and the approximation is fully converged at four points per wavelength of u . The standard Chebyshev interpolation requires on average π points per wavelength to start convergence, [27]. The dashed line in this figure presents the error for Chebyshev interpolation of u as a function of the number of nodes per wavelength of u .

3.3 Chebyshev correction

Once the product uw is approximated, the recovery of u is obtained by a point-wise division by w . Notice that the magnitude of the error in this process is inversely proportional to the values of $w(x)$, i.e.

$$\left| u(x) - \frac{F_N[uw](x)}{w(x)} \right| = \frac{1}{w(x)} |u(x)w(x) - F_N[uw](x)|.$$

Therefore, if corrections are made to the approximation in the regions where $w(x) < 0.1$, the error at the cutoff points would be about ten times larger than




Figure 6: Semilog plot of the maximum error in the approximation of uw as a function of $2N/k$, where $u(x) = x^2 \cos(60x)$, and $\lambda \approx N/24$ (\bullet –). For comparison, we also present the error in the Chebyshev interpolant of u (dashed line).

at the center of the interval. One may choose, instead, the cutoff points, x_a and x_b , to satisfy $w(x_a) = w(x_b) = 0.2$, allowing the error to be five times as large at the cutoff points as in the interior of the interval. In this article, we use 0.2 as the cutoff value. In the case of super-Gaussians, x_a and x_b are then given by

$$x_a = -\pi \left(\frac{\ln 5}{\alpha} \right)^{1/2\lambda} \quad \text{and} \quad x_b = \pi \left(\frac{\ln 5}{\alpha} \right)^{1/2\lambda}. \quad (8)$$

The approximation in the intervals $[-\pi, x_a]$ and $[x_b, \pi]$ must be corrected. In our algorithm we use local polynomial interpolation (or collocation for PDEs, [2, 17, 24]) on each subinterval. Polynomial approximation requires special set of nodes that are clustered more densely at the boundaries of an interval. We use Chebyshev nodes, given by

$$x_j = \frac{a-b}{2} \cos(j\pi/N_{ch}) + \frac{b+a}{2}, \quad j = 0, 1, \dots, N_{ch},$$

on the interval $[a, b]$.

In order to maintain the resolution criteria of the previous section – that convergence should start at two points per wavelength and fully converge at about four points per wavelength – we replace the number of Fourier nodes in each interval, $\approx (1 - x_b/\pi)N$, by approximately

$$N_{ch} \approx \frac{\pi}{2} \left(1 - \frac{x_b}{\pi} \right) N \quad (9)$$

Chebyshev nodes. The reason being that polynomial interpolation requires a minimum of π points per wavelength to start convergence, [27].

Asymptotically, as $N \rightarrow \infty$ and $\lambda \rightarrow \infty$, we have that

$$x_a \rightarrow -\pi \left(1 + \frac{\ln \frac{\ln 5}{\alpha}}{2\lambda} \right) \quad \text{and} \quad x_b \rightarrow \pi \left(1 + \frac{\ln \frac{\ln 5}{\alpha}}{2\lambda} \right)$$

and choosing $\lambda \sim N/(2C_\varepsilon)$ gives

$$N_{ch} \rightarrow \frac{\pi N}{4\lambda} (\ln \alpha - \ln \ln 5) = \frac{C_\varepsilon \pi}{2} (\ln \alpha - \ln \ln 5). \quad (10)$$

Using the data in Table 1, we obtain that $N_{ch} \rightarrow 8$ for a desired accuracy $\varepsilon = 10^{-4}$, $N_{ch} \rightarrow 16$ ($\varepsilon = 10^{-6}$), $N_{ch} \rightarrow 26$ ($\varepsilon = 10^{-8}$), $N_{ch} \rightarrow 35$ ($\varepsilon = 10^{-10}$), and $N_{ch} \rightarrow 48$ ($\varepsilon = 10^{-12}$).

Finally, we summarize the choice of parameters and algorithm as follows.

Given a target function u , number of nodes $2N$ for the Fourier approximation, and desired accuracy ε :

1. Set $\alpha = 32$, although other values are possible.
2. Choose C_ε according to Table 1; and $\lambda = N/(2C_\varepsilon)$.
3. Use the window $w(x) = \exp(-\alpha(x/\pi)^{2\lambda})$ to compute the windowed Fourier approximation of u , $G_N^w[u](x) = F_N[wu](x)/w(x)$.
4. Correct the approximation in the intervals $[-\pi, x_a]$ and $[x_b, \pi]$ using local Chebyshev interpolation. The values of x_a and x_b are given by (8) and the number of Chebyshev nodes, N_{ch} , in each subinterval is given by (9).

We point out that these are general guidelines and one may choose different parameters if additional information about the function being approximated is available. If a function, for instance, is known to be very smooth near the endpoints, one may choose N_{ch} smaller than the value provided by (9).

3.4 Approximation of analytic functions

Contours of equiconvergence in the complex plane give an a priori estimate of the error of approximations based on analytic properties (e.g. singularity locations) of the target function, [19, 2, 24, 25]. Given the choices of parameters provided above for the hybrid scheme, how do the equiconvergence levels compare to those of polynomial interpolation on Chebyshev nodes? The theorem that follows can be used to partially answer this question.

In order to state the theorem, we consider the following linear representation of the function f on $[a, b]$

$$\mathcal{L} [\{f(x_j)\}] (x) = \sum_{j=0}^m f(x_j) L_j(x), \quad (11)$$

where $L_j \in C[a, b]$. In the case of interpolation, L_j are the cardinal functions.

Theorem 3.2. *Let f be analytic in a neighborhood of $[a, b]$ and let Γ be a contour contained in the region of analyticity of f that encloses the interval $[a, b]$ once in the counterclockwise direction. The error at x in the approximation (11) is*

$$f(x) - \mathcal{L}[\{f(x_j)\}](x) = \frac{1}{2\pi i} \int_{\Gamma} f(z) r_m(z, x) dz, \quad (12)$$

where

$$r_m(z, x) = \frac{1}{z - x} - \sum_{j=0}^m \frac{L_j(x)}{z - x_j}.$$

Proof. The Cauchy's integral formula gives

$$f(x) = \frac{1}{2\pi i} \int_{\Gamma} \frac{f(z)}{z - x} dz$$

and using the complex residue calculus in (11) we have

$$\mathcal{L}[\{f(x_j)\}](x) = \frac{1}{2\pi i} \int_{\Gamma} f(z) \sum_{j=0}^m \frac{L_j(x)}{z - x_j} dz.$$

□

Using (12) we obtain the following bound on the L_{∞} error

$$\|f - \mathcal{L}[\{f(x_j)\}]\|_{\infty} \leq \frac{1}{2\pi} \int_{\Gamma} |f(z)| \|r_m(z, \cdot)\|_{\infty} dz, \quad (13)$$

where $\|\cdot\|_{\infty}$ is the sup norm with respect to $x \in [a, b]$.

Figure 7 displays the level sets of $\phi_m(z) = \log(\|r_m(z, \cdot)\|_{\infty})$ for the Chebyshev and hybrid methods with $m = 128$. Because of symmetry, and for better comparison, contours are shown in the half-plane in both cases. Notice that these level curves enclose the interval $[-1, 1]$ and taking the line integral in (13) over one of these contours yields

$$\|f - \mathcal{L}[\{f(x_j)\}]\|_{\infty} \leq \frac{|\Gamma| \max_{z \in \Gamma} |f|}{2\pi} \exp(\phi_m(\Gamma)),$$

where $\phi_m(\Gamma)$ is independent of f and $|\Gamma|$ is the arc-length of Γ .

By looking at the contours of ϕ_m in Figure 7, we therefore observe that the hybrid method provides a more accurate solution than standard polynomial interpolation when the function being approximated has singularities located near the imaginary axis (e.g. $f(x) = 1/(1 + 36x^2)$). Polynomial interpolation, on the other hand, is more accurate when singularities are closer to $z = \pm 1$ than $z = 0$. Error plots for the derivative of a Runge function are provided in the next section and show good agreement with this prediction.

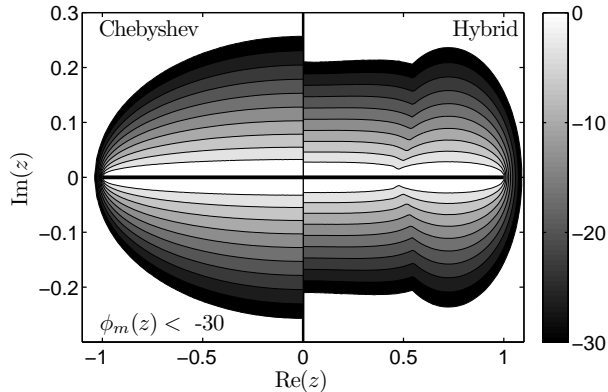


Figure 7: Equi-convergence contours for the Chebyshev (left) and hybrid (right) methods for $m = 128$. The thick line represents the interval $[-1, 1]$.

4 Eigenvalue stability and differentiation matrices

Before applying the hybrid technique to the solution of PDEs, we study the eigenvalues and pseudospectra of differentiation matrices. Notice that one may expect instability issues for time dependent problems because the approximation is not smooth across the interfaces $x = x_a$ and $x = x_b$.

4.1 Computing derivatives

The derivative of the windowed Fourier approximation is computed using the product rule, $(wu)' = w'u + wu'$,

$$\frac{d}{dx}u(x) \approx \frac{u(x)\frac{d}{dx}w(x) - \frac{d}{dx}F_N[wu](x)}{w(x)}. \quad (14)$$

The correction of the derivative in $[-\pi, x_a]$ and $[x_b, \pi]$ is performed by standard Chebyshev approximations in both subintervals.

Because this procedure is a linear operation, a differentiation matrix that maps function values into derivative values at hybrid points can be obtained. Figure 8 shows the hybrid nodes used in this process. Now suppose that the entries of the vector \mathbf{u} are the values of the function u at these nodes. We seek a matrix D such that the entries of the vector $\mathbf{u}' = D\mathbf{u}$ are approximately the values of the derivatives of u at the same nodal set. The differentiation process consists of four steps:




Figure 8: Hybrid grid.

1. Project the Chebyshev data on $[-\pi, x_a]$ and $[x_b, \pi]$ onto equispaced nodes;
2. Determine the derivative (14) using the windowed Fourier approximation (4) on $[-\pi, \pi]$;
3. Compute the Chebyshev approximation of the derivatives in each Chebyshev interval, $[-\pi, x_a]$ and $[x_b, \pi]$;
4. Replace the windowed Fourier approximation in $[-\pi, x_a]$ and $[x_b, \pi]$ with the Chebyshev approximation.

We point out that in most practical situations, these steps can be performed in $O(N \log N)$. More specifically, step 1 can be computed using the Lagrange barycentric interpolation formula and steps 2 and 3 using fast Fourier transforms. When solving boundary value or eigenvalue problems, however, differentiation matrices may be required. These matrices are also useful in the study of stability properties of the method when applied to time dependent PDEs.

Since each step described above corresponds to a linear map, they can all be performed at once using a single matrix. To this end let us define the following matrices: P the projection of data on Chebyshev nodes onto equally spaced nodes; D_{ch} the Chebyshev differentiation matrix, [24]; and D_F the windowed Fourier differentiation matrix that performs the differentiation in (14) written in blocks as

$$D_F = \begin{bmatrix} D_F^{\ell\ell} & D_F^{\ell i} & D_F^{\ell r} \\ D_F^{i\ell} & D_F^{ii} & D_F^{ir} \\ D_F^{r\ell} & D_F^{ri} & D_F^{rr} \end{bmatrix}.$$

Here we denoted left, interior, and right with ℓ , i , and r respectively. Then the hybrid differentiation matrix that operates on the hybrid grid is given by

$$\mathbf{u}' = \begin{bmatrix} \mathbf{u}_\ell \\ \mathbf{u}_F \\ \mathbf{u}_r \end{bmatrix}' = \begin{bmatrix} D_{ch} & 0 & 0 \\ D_F^{i\ell} P & D_F^{ii} & D_F^{ir} P \\ 0 & 0 & D_{ch} \end{bmatrix} \begin{bmatrix} \mathbf{u}_\ell \\ \mathbf{u}_F \\ \mathbf{u}_r \end{bmatrix}. \quad (15)$$

When the differentiation is performed in connection to a PDE, appropriate boundary conditions must be imposed across the interfaces $x = x_a$ and $x = x_b$. This is explored in more detail in the next section.

In Figure 9 we compare the error in the approximation of the derivative of $f(x) = 1/(1 + 36x^2) + \exp(x)$ using the hybrid and Chebyshev methods. Notice that there is a significant gain in accuracy over Chebyshev interpolation for this

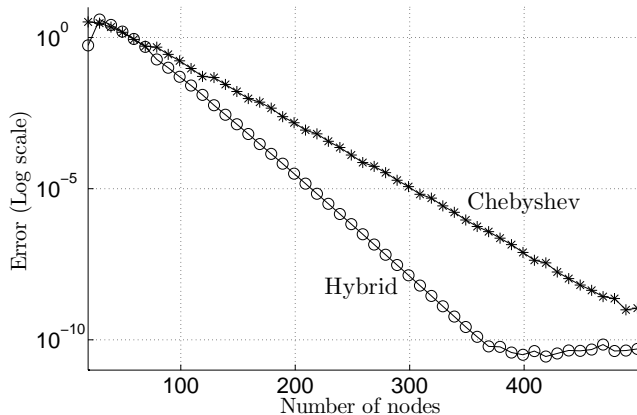


Figure 9: L_∞ error in $[-\pi, \pi]$ for the numerical differentiation of $f(x) = 1/(1 + 36x^2) + \exp(x)$ using the hybrid method with $\varepsilon = 10^{-12}$ (\circ) and the Chebyshev method ($*$).

Runge function. The analytic extension of this function on the complex plane is limited by singularities at $z = \pm i/6$, and as noted in Section 3.4, the more uniform node distribution of the hybrid method provides better resolution in this case. With a grid of 350 points the hybrid scheme is more accurate than the Chebyshev method to about five orders of magnitude.

4.2 The transport problem

We now consider the method of lines and study the eigenvalue stability for time dependent problems. Given a PDE and boundary conditions, we semidiscretize the derivatives in space and write the approximation in terms of a system of ordinary differential equations

$$u_t = Lu,$$

where L is the discretized spatial (linear) operator.

A rule of thumb is that the method of lines is stable if the eigenvalues of L , scaled by the time step Δt , lie in the stability region of the time discretization operator. When the matrix L is far from normal, however, stability conditions might be more restrictive, [15]. In this case, we shall rely on stability theorems based on pseudospectra instead, [26]. The ε -pseudospectrum of a matrix A is the set of $z \in \mathbb{C}$ such that

$$\|zI - A\|^{-1} > \varepsilon^{-1}.$$

In this context, the condition for stability is that the ε -pseudospectrum must lie within a distance $O(\varepsilon)$ of the stability region of the time stepping scheme as $\varepsilon \rightarrow 0$ (see [26] for details).




Figure 10: Pseudospectra (contours) and eigenvalues (dots) for the spatial operator in (16) using the hybrid discretization. Contour lines correspond to $\varepsilon = 10^{-13}, 10^{-11}, 10^{-9}, 10^{-7}, 10^{-5}, 10^{-3}, 10^{-1}, 10^1$. The outer contour is given by $\varepsilon = 10^1$ and the inner most contour by 10^{-13} . The right plot is a detailed view of the left graph near the origin.

In this section we consider the transport equation given by

$$u_t + u_x = 0, \quad x \in (-\pi, \pi), \quad t > 0, \quad (16)$$

and for simplicity, we assume homogeneous boundary condition $u(-\pi, t) = 0$. We point out that this spatial differential operator has no physical eigenvalues and eigenvalues of corresponding matrices will provide only necessary conditions for stability, [20].

In order to generate a differentiation matrix for (16), we need to provide proper boundary conditions for the Chebyshev correction regions. Given that the information travels from left to right in this particular example, we used the prescribed boundary condition at $x = -\pi$. At the interface $x = x_b$, we enforce continuity so the function value obtained from the Chebyshev correction matches the Fourier approximation values. Note that the approximation is not smooth across the inflow interface, $x = x_b$, only continuous. The approximation is not even continuous at $x = x_a$, but this interface is “corrected” at each time step. Nevertheless, the method is accurate and stable.

Figure 10 shows the spectrum and ε -pseudospectrum for the hybrid differentiation matrix for 128 grid points. The ε -pseudospectra of the continuous operator, for each ε , is the half-plane lying to the left of some line $\text{Re } z = \beta_\varepsilon$, [26]. This is in good agreement with Figure 10 (right). It can also be observed that the outlier eigenvalues are stable under perturbations and that for small ε , the contours lie entirely on the left half of the complex plane indicating that the method is stable in the presence of round-off errors.

In Figure 11 we present the spectrum and ε -pseudospectrum for a Chebyshev differentiation matrix for (16) using 128 Chebyshev nodes. Comparing this figure with Fig. 10, we note that in this case the hybrid method provides a




Figure 11: Same as Fig. 10 but for a Chebyshev differentiation operator.

stable approximation for a time step that can be more than twice the size that is required by the standard Chebyshev collocation method.

Figure 12 shows the spectral radius for hybrid and Chebyshev differentiation matrices for several matrix orders. Notice that for a large number of nodes, the spectral radius of the hybrid differentiation matrices grows linearly, in contrast to Chebyshev methods, for which the radius grows quadratically with the degree of the approximations. This linear relationship between eigenvalue growth and number of nodes was predicted in the asymptotic result (10). Because the number of Chebyshev nodes required to correct the approximations near the boundaries is asymptotically constant for large N , the spacing between nodes decreases as $O(N^{-1})$ for the hybrid method. This indicates that the maximum allowed time step for stability is indeed $\Delta t = O(N^{-1})$.

4.3 The acoustic problem

We now consider the one dimensional acoustic equations

$$u_t = p_x, \quad p_t = u_x, \quad x \in (-\pi, \pi), \quad t > 0, \quad (17)$$

where u can be seen as the fluid velocity and p as the pressure. Usually boundary conditions are provided for u at $x = \pm\pi$, and for simplicity, we consider homogeneous Dirichlet boundary conditions. No boundary condition is directly imposed on p . Similarly, when implementing the hybrid scheme, we only explicitly imposed continuity of u across the interfaces $x = x_a$ and $x = x_b$.

In contrast to the transport problem, the continuous spatial operator in the acoustic problem has purely imaginary eigenvalues given by $\pm ki$, $k = 0, 0.5, 1, 1.5, 2, \dots$, which correspond to natural frequencies of the system. Moreover, discretization matrices for this operator should be close to normal so that the eigenvalues provide relevant information of the overall behavior of the numerical solution. Since the eigenvalues are imaginary, the numerical solution of

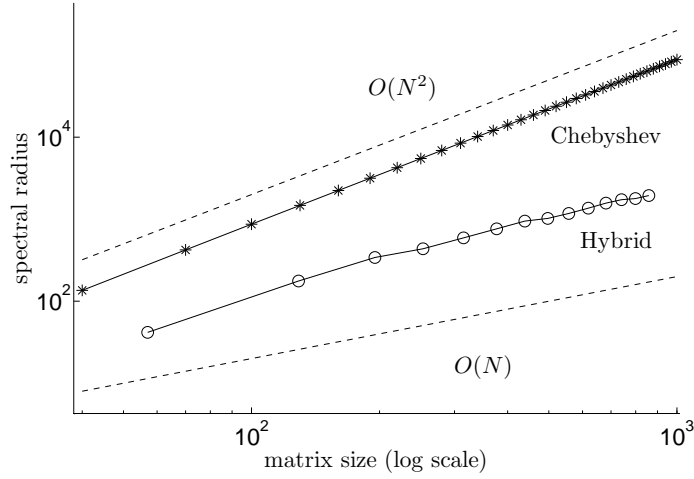


Figure 12: Spectral radius of hybrid and Chebyshev differentiation matrices as functions of the number of collocation points.

(17) can be stably computed with a mass-conserving time stepping scheme such as leap frog, in which case truncation errors would be purely dispersive.

Although we are not able to prove that the spatial operator matrix for the hybrid method has purely imaginary eigenvalues, our numerical results yield zero to machine precision for the real part of the eigenvalues. Figure 13 shows the error in the computed eigenvalues when compared to the eigenvalues of the operator in (17). In this figure we consider only eigenvalues of positive imaginary part and compare the accuracy of the eigenvalues of the hybrid matrix with those obtained with a Chebyshev matrix. For this approximation we used $N = 100$ in (1), giving a total of 200 equally spaced nodes, and a hybrid grid with 230 nodes. Therefore, if $k = 100$, the eigenfunction is approximated with four points per wavelength. Notice that the error increases for $k > 100$ because of the insufficient resolution in the higher modes. In particular, for $k = 200$ or higher, two points per wavelength or less are being used, resulting in large errors. The figure shows good agreement with the findings in Section 3. Although the Chebyshev method is more accurate for some wave numbers, as it converges supergeometrically once the minimum resolution is achieved, the hybrid method gives significantly more accurate results for larger wave numbers. Moreover, the large magnitude of the outlier eigenvalues obtained with the Chebyshev method imposes severe restriction on the size of the time step allowed for time integration – a difficulty that is circumvented by the hybrid method.




Figure 13: Error in the computed eigenvalue μ_k with the hybrid method (x) and Chebyshev collocation (o). Due to symmetry, only eigenvalues with positive imaginary values are considered ($\mu_k^* = ki$, $k = 0.5, 1, 1.5 \dots$).

5 Numerical solution of time dependent PDEs

In this section we consider three time dependent problems: the acoustic equation in one dimension, the wave equation in two dimensions, and the simulation of a thermally driven flow in a cavity of large aspect ratio. Our implementation uses the method of lines with three explicit time stepping schemes: Runge-Kutta, leap frog, and Adams-Bashforth.

5.1 The acoustic equation

Figure 14 presents errors in the computed solutions of the acoustic equation (17) with homogeneous Dirichlet boundary conditions in u and initial conditions

$$u(0, x) = \exp(-50x^2) \quad \text{and} \quad p(0, x) = 0.$$

To advance in time, a fourth order Runge-Kutta scheme was used.

The error displayed in Fig. 14 (left) shows the spectral convergence of the hybrid method. The maximum L_∞ error over $0 \leq t \leq 4\pi$ was computed. For comparison, the error for the Chebyshev spectral method is also shown. The solutions were computed with $\Delta t = 10^{-4}$ and several values of N . Both methods converge exponentially, with the hybrid scheme giving more accurate results.

Figure 14 (right) presents the temporal error for solutions computed with the hybrid scheme. As predicted in the previous section, the method is stable for long time simulations. In this case, calculations were carried out with $\Delta t = 0.1/N$ and $N = 50, 60$, and 70 . In this experiment, the time integration error dominates and is $O(1/N^4)$.

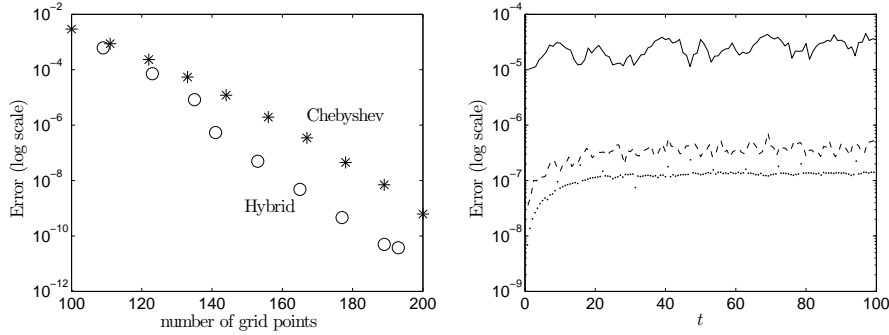


Figure 14: Error in the computed solution of the acoustic equation. Left: maximum error, $\max_{t \in [0, 4\pi], x \in [-\pi, \pi]} |u_{comp}(x, t) - u_{exact}(x, t)|$, as a function of N and $\Delta t = 10^{-4}$ for the Chebyshev and hybrid methods. Right: Temporal error of the hybrid method, $\max_{x \in [-\pi, \pi]} |u_{comp}(x, t) - u_{exact}(x, t)|$, for $\Delta t = 0.1/N$ and $N = 50$ (solid), 60 (dashed), and 70 (dotted).

5.2 The wave equation in 2D

Rectangular domains can be handled by tensor product approximations. Figure 15 (left) shows the 101×101 hybrid grid we use in our next example. In multiple dimensions, computations are carried as sequences of one dimensional approximations.

Our first two dimensional test is the linear wave equation, $u_{tt} = \Delta u$, defined on $[-\pi, \pi] \times [-\pi, \pi]$ with homogeneous Dirichlet boundary conditions. Solutions were obtained for the initial condition $u(0, x, y) = \exp(-25((x - 1.5)^2 + y^2))$. Figure 15 (right) depicts the solution at $t = 4$, when results obtained with the hybrid and Chebyshev methods are compared.

In Figure 16 we show the error in the computed solutions. Both methods were implemented using leap frog time integration with the same number of grid points and time step, $\Delta t = 10^{-4}$. These results demonstrate that the hybrid scheme is at least one order of magnitude more accurate than the Chebyshev method.

5.3 Natural convection in a cavity

Our final numerical example is the simulation of a buoyancy-driven flow in a deep cavity heated from below, [9, 21, 28]. We assume that the fluid is incompressible and use the following governing equations in non-dimensional form to simulate the flow, see [1, 9] for details,

$$\begin{aligned} \omega_t + \psi_y \omega_x - \psi_x \omega_y &= Pr \Delta \omega + Ra Pr T_x, \\ T_t + \psi_y T_x - \psi_x T_y &= \Delta T, \\ \Delta \psi &= -\omega, \end{aligned}$$




Figure 15: Left: the hybrid grid. Right: contour plot of the solution of the wave equation at $t = 4$.




Figure 16: Error in the solution of the wave equation at $t = 4$ using the hybrid (left) and Chebyshev (right) methods.

where $(x, y) \in (-1, 1) \times (-4, 4)$ and $t > 0$. Here ω is the fluid vorticity, ψ the stream function and T the temperature. The independent non-dimensional constants are the Rayleigh number Ra , which reflects the buoyant contribution, and the Prandtl number Pr , which is the ratio of viscous to thermal diffusion. In our simulation they are set to $Ra = 5 \times 10^4$ and $Pr = 0.71$ (air).

The fluid is assumed to be at rest at $t = 0$, with $T = \psi = \omega = 0$. The boundary condition for the stream function is $\psi_\Gamma = 0$, which implies that there is no mass transfer through the boundary Γ . The value of the vorticity at the walls is expressed as $\omega_\Gamma = -\Delta\psi|_\Gamma$ and the temperature at Γ is defined by

$$T(t, x, y) = \begin{cases} \tanh^4(100t) \exp(-30(x - 0.1)^2), & y = 0, \quad -1 \leq x \leq 1, \quad t > 0, \\ 0, & (x, y) \in \Gamma, \quad y \neq 0, \quad t > 0. \end{cases}$$

The boundary condition for T at the bottom wall was chosen so that it is compatible with the initial condition and quickly asymptotes to $\exp(-30(x - 0.1)^2)$. This non-symmetric heat source drives the flow into very interesting dynamics, with the fluid being heated at the bottom and cooled at the top and lateral walls.

Although Chebyshev methods are particularly suitable for the simulation of incompressible flows in enclosures, the large number of grid points needed to approximate the solution in the y direction in this tall cavity results in severe time step restrictions. Therefore, in our simulation we use the standard pseudospectral Chebyshev method in the x direction and the new hybrid scheme to approximate derivatives in the y direction. The grid is shown in Figure 17 (right). Here again we rely on tensor product approximations.

Figure 17 presents the numerical solution at $t = 2.7$ on a 40×140 grid. A third-order Adams-Bashforth scheme was used to advance the problem in time. These results showed good agreement with simulations performed using Chebyshev in both directions but required a third of the computational time. The performance of the hybrid method in other flow simulations is currently being investigated. The method seems to also be suitable for problems with absorbing boundary conditions.

6 Final Remarks

This paper introduced a hybrid scheme for non-periodic problems that combines Fourier and Chebyshev approximations. In contrast to standard filtering algorithms, the hybrid method uses a physical space window that is derived based on accuracy considerations (in the interior of the domain) and the number of nodes required for the Chebyshev correction near the boundaries. Such criteria ensure that the method is of practical value for approximating non-periodic functions and for solving time dependent problems. Eigenvalue stability analysis indicates that the method is stable for time dependent problems when boundary conditions are consistently applied at the interface of the windowed Fourier and Chebyshev regions. The spectral radius of hybrid differentiation matrices




Figure 17: Left: contour levels of the square root of the temperature at $t = 2.7$, black corresponds to $T = 1$ and white $T = 0$. Middle: streamlines. Right: velocity field on a hybrid grid.

has been shown to grow linearly with the number of grid points, in contrast with the purely Chebyshev case, where the radius grows quadratically. As a consequence, the new method is less restrictive on the time step size for the simulation of PDEs.

Fast Fourier transforms can be used to compute the Fourier approximation and polynomial corrections. As the number of grid points is increased, wider windows can be used corresponding to smaller corrective regions. Thus the degree of the polynomial correction near endpoints is small even when the function being approximated has large gradients. Because of the more uniform node distribution, the hybrid method provides better resolution at the interior of the domain than pseudospectral polynomial methods, which require quadrature points. A more evenly spaced distribution may be more appropriate for some PDE applications.

The numerical examples in Section 5 show that the new method performs favorably to Chebyshev schemes. Solutions of the acoustic equation in one dimension showed good results when $\Delta t = O(1/N)$ is used to explicitly advance in time. For two dimensional regions, approximations can be obtained as tensor products. We have also observed significant improvement over traditional methods when solving the linear wave equation. To demonstrate that the method can be used to solve nonlinear PDEs with a diffusion term, we used the hybrid scheme to simulate a buoyancy driven fluid flow in a tall rectangular enclosure.

Acknowledgments

This work was partially supported by NSF grants DMS-0608844 and DMS-0510813. Anne Gelb was also supported by NSF FRG DMS-0652833, NSF RUI-0608844 and NSF SCREMS DMS-0421846. We thank John P. Boyd for useful discussions on window functions and for sharing his derivation and bounds of the Fourier coefficients of super-Gaussians.

References

- [1] O. AYDIN, A. ÜNAL, AND T. AYHAN, *Natural convection in rectangular enclosures heated from one side and cooled from the ceiling*, Int. J. Heat Mass Transfer, 42 (1999), pp. 2345–2355.
- [2] J. P. BOYD, *Chebyshev and Fourier Spectral Methods*, Dover Publications Inc., Mineola, NY, second ed., 2001.
- [3] ———, *Prolate spheroidal wavefunctions as an alternative to Chebyshev and Legendre polynomials for spectral element and pseudospectral algorithms*, J. Comput. Phys., 199 (2004), pp. 688–716.
- [4] ———, *Fourier embedded domain methods: extending a function defined on an irregular region to a rectangle so that the extension is spatially periodic and C^∞* , Appl. Math. Comput., 161 (2005), pp. 591–597.

- [5] ———, *Asymptotic Fourier coefficients for a C^∞ bell (smoothed-“top-hat”) & the Fourier extension problem*, J. Sci. Comput., 29 (2006), pp. 1–24.
- [6] ———, *Exponentially accurate Runge–Free approximation of non-periodic functions from samples on an evenly-spaced grid*, Appl. Math. Lett., 20 (2007), pp. 971–975.
- [7] J. P. BOYD AND J. R. ONG, *Exponentially-convergent strategies for defeating the Runge phenomenon for the approximation of non-periodic functions, part I: Single-interval schemes*, Commun. Comput. Phys., 5 (2009), pp. 484–497.
- [8] M. D. BUHMANN, *Radial Basis Functions: Theory and Implementations*, Cambridge University Press, Cambridge, 2003.
- [9] M. FARHANGNIA, S. BIRINGEN, AND L. J. PELTIER, *Numerical simulation of two-dimensional buoyancy-driven turbulence in a tall rectangular cavity*, Int. J. Numer. Methods Fluids, 23 (1996), pp. 1311–1326.
- [10] B. FORNBERG, *A Practical Guide to Pseudospectral Methods*, Cambridge University Press, Cambridge, 1996.
- [11] A. GELB, *Reconstruction of piecewise smooth functions from non-uniform grid point data*, J. Sci. Comput., 30 (2007), pp. 409–440.
- [12] A. GELB, R. B. PLATTE, AND W. S. ROSENTHAL, *The discrete orthogonal polynomial least squares method for approximation and solving partial differential equations*, Commun. Comput. Phys., 3 (2008), pp. 734–758.
- [13] A. GELB AND J. TANNER, *Robust reprojection methods for the resolution of the Gibbs phenomenon*, Appl. Comput. Harmon. Anal., 20 (2006), pp. 3–25.
- [14] D. GOTTLIEB AND S. A. ORSZAG, *Numerical Analysis of Spectral Methods: Theory and Applications*, SIAM, Philadelphia, 1977.
- [15] B. GUSTAFSSON, H.-O. KREISS, AND J. OLIGER, *Time Dependent Problems and Difference Methods*, John Wiley & Sons Inc., New York, 1995.
- [16] N. HALE AND L. N. TREFETHEN, *New quadrature methods from conformal maps*, SIAM J. Numer. Anal., 46 (2008), pp. 930–948.
- [17] J. S. HESTHAVEN, S. GOTTLIEB, AND D. GOTTLIEB, *Spectral Methods for Time-Dependent Problems*, Cambridge University Press, Cambridge, UK, 2007.
- [18] D. KOSLOFF AND H. TAL-EZER, *A modified Chebyshev pseudospectral method with an $O(N^{-1})$ time step restriction*, J. Comput. Phys., 104 (1993), pp. 457–469.

- [19] R. B. PLATTE AND T. A. DRISCOLL, *Polynomials and potential theory for Gaussian radial basis function interpolation*, SIAM J. Numer. Anal., 43 (2005), pp. 750–766.
- [20] S. C. REDDY AND L. N. TREFETHEN, *Lax-stability of fully discrete spectral methods via stability regions and pseudo-eigenvalues*, Comput. Methods Appl. Mech. Engrg., 80 (1990), pp. 147–164.
- [21] I. E. SARRIS, I. LEKAKIS, AND N. S. VLACHOS, *Natural convection in rectangular tanks heated locally from below*, Int. J. Heat Mass Transfer, 47 (2004), pp. 3549–3563.
- [22] E. TADMOR, *The exponential accuracy of Fourier and Chebyshev differencing methods*, SIAM J. Numer. Anal., 23 (1986), pp. 1–10.
- [23] E. TADMOR AND J. TANNER, *Adaptive filters for piecewise smooth spectral data*, IMA J. Numer. Anal., 25 (2005), pp. 635–647.
- [24] L. N. TREFETHEN, *Spectral Methods in MATLAB*, SIAM, Philadelphia, PA, 2000.
- [25] ———, *Is Gauss quadrature better than Clenshaw-Curtis?*, SIAM Rev., 50 (2008), pp. 67–87.
- [26] L. N. TREFETHEN AND M. EMBREE, *Spectra and Pseudospectra*, Princeton University Press, Princeton, NJ, 2005.
- [27] J. A. C. WEIDEMAN AND L. N. TREFETHEN, *The eigenvalues of second-order spectral differentiation matrices*, SIAM J. Numer. Anal., 25 (1988), pp. 1279–1298.
- [28] C. XIA AND J. Y. MURTHY, *Buoyancy-driven flow transitions in deep cavities heated from below*, J. Heat Transfer, 124 (2002), pp. 650–659.