Chapter 7.
Fourier spectral methods

Think globally. Act locally.
— *BUMPER STICKER (1985)*

Finite difference methods, like finite element methods, are based on local representations of functions—usually by low-order polynomials. In contrast, spectral methods make use of global representations, usually by high-order polynomials or Fourier series. Under fortunate circumstances the result is a degree of accuracy that local methods cannot match. For large-scale computations, especially in several space dimensions, this higher accuracy may be most important for permitting a coarser mesh, hence a smaller number of data values to store and operate upon. It also leads to discrete models with little or no artificial dissipation, a particularly valuable feature in high Reynolds number fluid flow calculations, where the small amount of physical dissipation may be easily overwhelmed by any dissipation of the numerical kind. Spectral methods have achieved dramatic successes in this area.

Some of the ideas behind spectral methods have been introduced several times into numerical analysis. One early proponent was Cornelius Lanczos, in the 1950s, who showed the power of Fourier series and Chebyshev polynomials in a variety of problems where they had not been used before. The emphasis was on ordinary differential equations. Lanczos's work has been carried on, especially in Great Britain, by a number of colleagues such as C. W. Clenshaw.

More recently, spectral methods were introduced again by Kreiss and Oliger, Orszag, and others in the 1970s for the purpose of solving the partial differential equations of fluid mechanics. Increasingly they are becoming viewed within some fields as an equal competitor to the better established finite difference and finite element approaches. At present, however, they are less well understood.

Spectral methods fall into various categories, and one distinction often made is between "Galerkin," "tau," and "collocation" (or "pseudospectral") spectral methods. In a word, the first two work with the coefficients of a global expansion, and the latter with its values at points. The discussion in this book is entirely confined to collocation methods, which are probably used the most often, chiefly because they offer the simplest treatment of nonlinear terms.

Spectral methods are affected far more than finite difference methods by the presence of boundaries, which tend to introduce stability problems that are ill-understood and sometimes highly restrictive as regards time step. Indeed, difficulties with boundaries, direct and indirect, are probably the primary reason why spectral methods have not replaced their lower-accuracy competition in most applications. Chapter 8 considers spectral methods with boundaries, but the present chapter assumes that there are none. This means that the spatial domain is either infinite—a theoretical device, not applicable in practice—or periodic. In those cases where the physical problem naturally inhabits a periodic domain, spectral methods may be strikingly successful. Conspicuous examples are the global circulation models used by meteorologists. Limited-area meteorological codes, since they require boundaries, are often based on finite difference formulas, but as of this writing almost all of the global circulation codes in use— which model flow in the atmosphere of the entire spherical earth— are spectral.

7.1. An example

Spectral methods have been most dramatically successful in problems with periodic geometries. In this section we present two examples of this kind that involve elastic wave propagation. Both are taken from B. Fornberg, "The pseudospectral method: comparisons with finite differences for the elastic wave equation," *Geophysics 52* (1987), 483–501. Details and additional examples can be found in that paper.*

Elastic waves are waves in an elastic medium such as an iron bar, a building, or the earth, and they come in two varieties: "P" waves (pressure or primary), characterized by longitudinal vibrations, and "S" waves (shear or secondary), characterized by transverse vibrations. The partial differential equations of elasticity can be written in various forms, such as a system of two second-order equations involving displacements. For his numerical simulations, Fornberg chose a formulation as a system of five first-order equations.

Figures 7.1.1 and 7.1.2 show the results of calculations for two physical problems. In the first, a P wave propagates uninterruptedly through a periodic, uniform medium. In the second, an oblique P wave oriented at 45° hits a horizontal interface at which the wave speeds abruptly cut in half. The result is reflected and transmitted P and S waves. For this latter example, the actual computation was performed on a domain of twice the size shown — which is a hint of the trouble one may be willing to go to, with spectral methods, to avoid coping explicitly with boundaries.

The figures show that spectral methods may sometimes decisively outperform second-order and fourth-order finite difference methods. In particular, spectral methods are *nondispersive,* and in a wave calculation, that property can be of great importance. In these examples the accuracy achieved by the spectral calculation on a 64×64 grid is not matched by fourth-order finite differences on a 128×128 grid, or by second-order finite differences on a 256×256 grid. The corresponding differences in work and storage are enormous.

Fornberg picked his examples carefully; spectral methods do not always perform so convincingly. Nevertheless, sometimes they are extremely impressive. Although the reasons are not fully understood, their advantages often hold not just for problems involving smooth functions, but even in the presence of discontinuities.

---

*The figures in this section will appear in the published version of this book only with permission.

(a) Schematic initial and end states

(b) Computational results

Figure 7.1.1. Spectral and finite difference simulations of a P wave propagating through a uniform medium (from Fornberg, 1987).

(a) Schematic initial and end states

(b) Computational results

Figure 7.1.2. Spectral and finite difference simulations of a P wave incident obliquely upon an interface (from Fornberg, 1987).

7.2. Unbounded grids

We shall begin our study of spectral methods by looking at an infinite, unbounded domain. Of course, real computations are not carried out on infinite domains, but this simplified problem contains many of the essential features of more practical spectral methods.

Consider again ℓ_h^2, the set of square-integrable functions $v = \{v_j\}$ on the unbounded regular grid $h\mathbb{Z}$. As mentioned already in §3.3, the foundation of spectral methods is the **spectral differentiation operator** $D: \ell_h^2 \to \ell_h^2$, which can be described in several equivalent ways. One is by means of the Fourier transform:

SPECTRAL DIFFERENTIATION BY THE SEMIDISCRETE FOURIER TRANS.

(1) *Compute* $\hat{v}(\xi)$*;*

(2) *Multiply by* $i\xi$*;*

(3) *Inverse transform:*

$$Dv = \mathcal{F}_h^{-1}(i\xi\mathcal{F}_h(v)). \qquad (7.2.1)$$

Another is in terms of band-limited interpolation. As described in §2.3, one can think of the interpolant as a Fourier integral of band-limited complex exponentials or, equivalently, as an infinite series of sinc functions:

SPECTRAL DIFFERENTIATION BY SINC FUNCTION INTERPOLATION.

(1) *Interpolate* v *by a sum of sinc functions* $q(x) = \sum_{k=-\infty}^{\infty} v_k S_h(x - x_k)$*;*

(2) *Differentiate the interpolant at the grid points* x_j*:*

$$(Dv)_j = q'(x_j) \qquad (7.2.2)$$

Recall that the sinc function $S_h(x)$, defined by

$$S_h(x) = \frac{\sin(\pi x/h)}{\pi x/h}, \qquad (7.2.3)$$

is the unique function in L^2 that interpolates the discrete delta function e_j,

$$e_j = \begin{cases} 1 & j=0, \\ 0 & j\neq 0, \end{cases} \tag{7.2.4}$$

and that moreover is band-limited in the sense that its Fourier transform has compact support contained in $[-\pi/h,\pi/h]$.

For higher order spectral differentiation, we multiply $\mathcal{F}_h(v)$ by higher powers of $i\xi$, or equivalently, differentiate $q(x)$ more than once.

Why are these two descriptions equivalent? The fundamental reason is that $S_h(x)$ is not just any interpolant to the delta function e, but the *band-limited* interpolant. For a precise argument, note that both processes are obviously linear, and it is not hard to see that both are shift-invariant in the sense that $D(K^m v) = K^m Dv$ for any m. (The shift operator K was defined in (3.2.8).) Since an arbitrary function $v\in \ell_h^2$ can be written as a convolution sum $v_j = \sum_{k=-\infty}^{\infty} v_k e_{j-k}$, it follows that it is enough to prove that the two processes give the same result when applied to the particular function e. That equivalence results from the fact that the Fourier transform of e is the constant function h, whose inverse Fourier transform is in turn precisely $S_h(x)$.

Since spectral differentiation constitutes a linear operation on ℓ_h^2, it can also be viewed as multiplication by a biinfinite Toeplitz matrix:

$$D = \frac{1}{h}\begin{pmatrix} & & \ddots & & & & & & \\ & & & & & & & & \\ \cdots & -\frac{1}{3} & \frac{1}{2} & -1 & 0 & 1 & -\frac{1}{2} & \frac{1}{3} & \cdots \\ & & & & & & & & \\ & & & & & & \ddots & & \end{pmatrix}. \tag{7.2.5}$$

As discussed in §3.3, this matrix is the limit of banded Toeplitz matrices corresponding to finite difference differentiation operators of increasing orders of accuracy; see Table 3.3.1 on p. 131. (In this chapter we drop the subscript on the symbol D_∞ used in §3.3.) We shall be careless in this text about the distinction between the operator D and the matrix D that represents it. Another way to express the same thing is to write

$$Dv = a*v, \qquad a = \frac{1}{h^2}(\ \cdots\ \tfrac{1}{3}\ -\tfrac{1}{2}\quad 1\quad 0\ -1\quad \tfrac{1}{2}\ -\tfrac{1}{3}\ \cdots\ ) \tag{7.2.6}$$

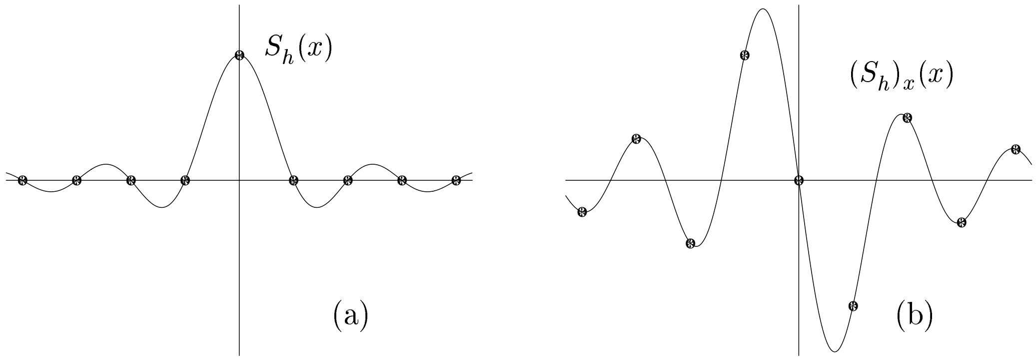

Figure 7.2.1. The sinc function $S_h(x)$ and its derivative $(S_h)_x(x)$.

as in (3.3.16).

The coefficients of (7.2.5)–(7.2.6) can be derived from either the Fourier transform or the sinc function interpretation. Let us begin with the latter. The sinc function has derivative

$$(S_h)_x(x) = \frac{\cos(\pi x/h)}{x} - \frac{\sin(\pi x/h)}{\pi x^2/h}, \tag{7.2.7}$$

with values

$$(S_h)_x(x_j) = \begin{cases} 0 & \text{if } j = 0, \\ \frac{(-1)^j}{jh} & \text{if } j \neq 0 \end{cases} \tag{7.2.8}$$

at the grid points. See Figure 7.2.1. This is precisely the "zeroth column" of D, since that column must by definition contain the values on the grid of the spectral derivative of the delta function.* The other columns, corresponding to delta functions centered at other points x_j, contain the same entries appropriately shifted.

Now let us rederive (7.2.5)–(7.2.6) by means of the Fourier transform. If $Dv = a * v$, then $\widehat{Dv}(\xi) = \hat{a}(\xi)\hat{v}(\xi)$, and for spectral differentiation we want $\hat{a}(\xi) = i\xi$. Therefore by the inverse semidiscrete Fourier transform (2.2.7),

$$a_j = \frac{1}{2\pi} \int_{-\pi/h}^{\pi/h} i\xi e^{i\xi jh} d\xi.$$

For $j = 0$ the integral is 0, giving $a_0 = 0$, while for $j \neq 0$, integration by parts

*Think about this. Make sure you understand why (7.2.8) represents a column rather than a row of D.

yields

$$a_j = \frac{1}{2\pi} i\xi \, \frac{e^{i\xi jh}}{ijh}\bigg|_{-\pi/h}^{\pi/h} - \frac{1}{2\pi}\int_{-\pi/h}^{\pi/h} \frac{ie^{i\xi jh}}{ijh} d\xi.$$

The integral here is again zero, since the integrand is a periodic exponential, and what remains is

$$\begin{aligned} a_j &= \frac{1}{2\pi jh}\left(\frac{\pi}{h}e^{i\pi j} - (-\frac{\pi}{h})e^{-i\pi j}\right) \\ &= \frac{1}{2jh^2}(e^{i\pi j} + e^{-i\pi j}) = \frac{(-1)^j}{jh^2}, \end{aligned} \tag{7.2.9}$$

as in (7.2.6).

The entries of D are suggestive of the Taylor expansion of $\log(1+x)$, and this is not a coincidence. In the notation of the spatial shift operator K of (3.2.8), D can be written

$$D = \frac{1}{h}(K - \tfrac{1}{2}K^2 + \tfrac{1}{3}K^3 - \cdots) - \frac{1}{h}(K^{-1} - \tfrac{1}{2}K^{-2} + \tfrac{1}{3}K^{-3} - \cdots),$$

which corresponds formally to

$$\begin{aligned} D &= \frac{1}{h}\log(1+K) - \frac{1}{h}\log(1+K^{-1}) \\ &= \frac{1}{h}\log\left(\frac{1+K}{1+K^{-1}}\right) = \frac{1}{h}\log K. \end{aligned} \tag{7.2.10}$$

Therefore formally, $e^{hD} = K$, and this makes sense: by integrating the derivative over a distance h, one gets a shift. See the proof of Theorem 1.2.

If $v_j = e^{i\xi jh}$ for some $\xi \in [-\pi/h, \pi/h]$, then $Dv = i\xi v$. Therefore $i\xi$ is an eigenvalue of the operator D.* On the other hand, if v has the same form with $\xi \notin [-\pi/h, \pi/h]$, then ξ will be indistinguishable on the grid from some alias wave number $\xi' \in [-\pi/h, \pi/h]$ with $\xi' = \xi + 2\pi\nu/h$ for some integer ν, and the result will be $Dv = i\xi' v$. In other words in Fourier space, the spatial differentiation operator becomes multiplication by a periodic function, thanks to the discrete grid, and in this sense is only an approximation to the exact differentiation operator for continuous functions. Figure 7.2.2 shows the situation graphically. For band-limited data, however, the spectral differentiation operator is exact, in contrast to finite difference differentiation operators, which are exact only in the limit $\xi \to 0$ (dashed line in the Figure).

* Actually, this is not quite true: by definition, an eigenvector must belong to ℓ^2, and $e^{i\xi jh}$ does not. Strictly speaking, $i\xi$ is in the spectrum of D but is not an eigenvalue. However, this technicality is unimportant for our purposes, and will be ignored in the present draft of this book.

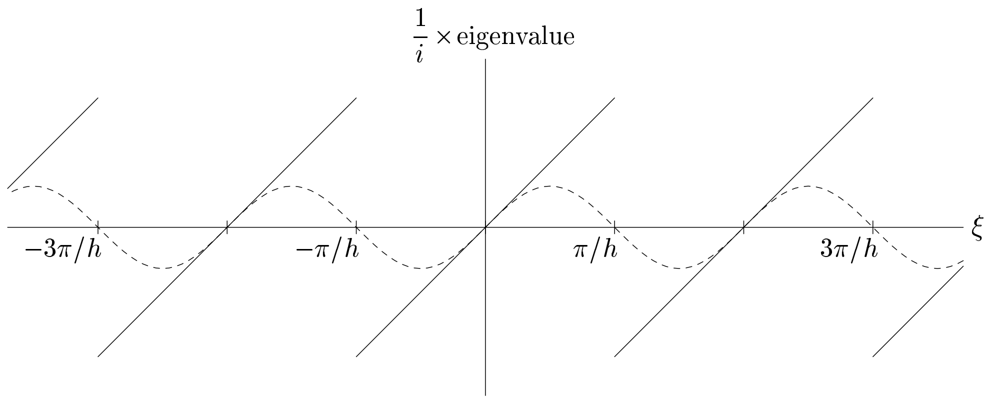

Figure 7.2.2. Eigenvalue of D (divided by i) corresponding to the eigenfunction $e^{i\xi x}$, as a function of ξ. The dashed line shows corresponding eigenvalues for the finite difference operator D_2.

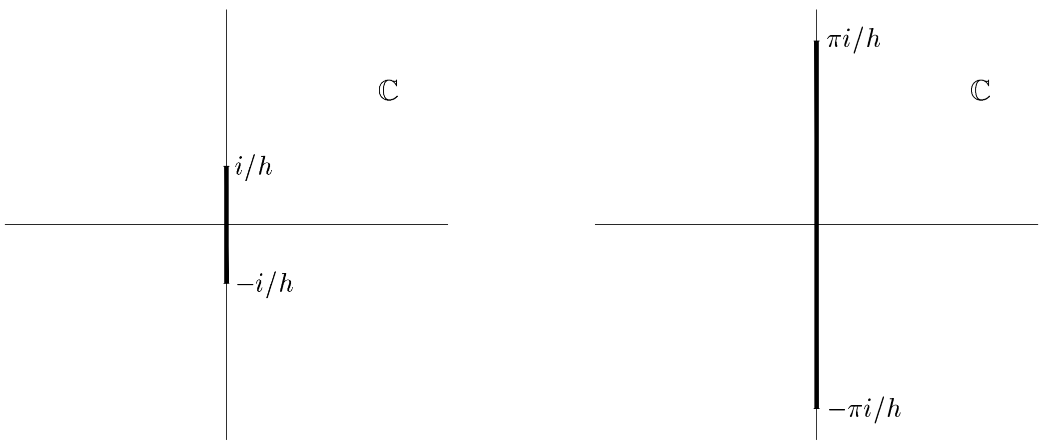

(a) D_2 (finite difference) *(b)* D (spectral)

Figure 7.2.3. Eigenvalues of finite difference and spectral first-order differentiation matrices, as subsets of the complex plane.

Figure 7.2.3 compares the spectrum of D to that of the second-order finite difference operator $D_2 = \delta_0$ of §3.3.

D is a bounded linear operator on ℓ_h^2, with norm

$$\|D\| = \max_{\xi\in[-\pi/h,\pi/h]} |i\xi| = \frac{\pi}{h} \tag{7.2.11}$$

(see §§2.5,3.5). Notice that the norm increases to infinity as the mesh is refined.

This is inevitable, as the differentiation operator for continuous functions is unbounded.

So far we have described the spectral approximation to the first derivative operator $\partial/\partial x$, but it is an easy matter to approximate higher derivatives too. For the second derivative, the coefficients turn out to be

$$D^2 = \frac{1}{h^2}\begin{pmatrix} & & \ddots & & & & & & \\ & & & & & & & & \\ \cdots & \frac{2}{9} & -\frac{2}{4} & \frac{2}{1} & -\frac{\pi^2}{3} & \frac{2}{1} & -\frac{2}{4} & \frac{2}{9} & \cdots \\ & & & & & & & & \\ & & & & & & \ddots & & \end{pmatrix}. \tag{7.2.12}$$

To derive the entries of this matrix, one can simply square D; this leads to infinite series to be summed. One can differentiate (7.2.7) a second time. Or one can compute the inverse Fourier transform of $\hat{a}(\xi) = -\xi^2$, as follows. Two integrations by parts are involved, and terms that are zero have been dropped. For $j \neq 0$,

$$\begin{aligned} a_j &= \frac{1}{2\pi}\int_{-\pi/h}^{\pi/h} -\xi^2 e^{i\xi jh} d\xi \\ &= \frac{1}{2\pi}\int_{-\pi/h}^{\pi/h} 2\xi \frac{e^{i\xi jh}}{ijh} d\xi \\ &= \frac{\xi}{\pi}\,\frac{e^{i\xi jh}}{(ijh)^2}\bigg|_{-\pi/h}^{\pi/h} = -\frac{e^{ij\pi} + e^{-ij\pi}}{j^2h^3} = \frac{2(-1)^{j+1}}{j^2h^3}. \end{aligned} \tag{7.2.13}$$

For $j = 0$ the integral is simply

$$a_0 = -\frac{1}{2\pi}\left(\frac{2\pi^3}{3h^3}\right) = -\frac{\pi^2}{3h^3}.$$

The effect of D^2 on a function $v_j = e^{i\xi jh}$ is to multiply it by the square of the factor associated with D, as illustrated in Figure 7.2.4. Again one has a periodic multiplier that is exactly correct for $\xi \in [-\pi/h, \pi/h]$. The dashed line shows the analogous curve for the standard centered three-point finite

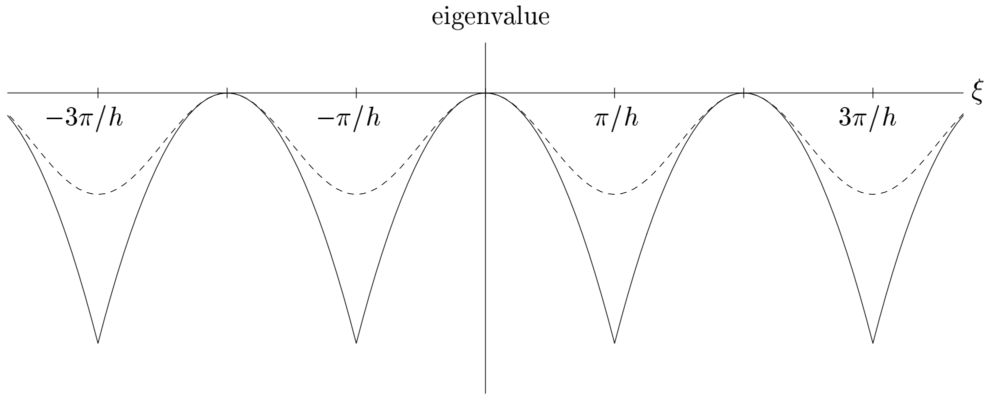

Figure 7.2.4. Eigenvalue of D^2 corresponding to the eigenfunction $e^{i\xi x}$, as a function of ξ. The dashed line shows corresponding eigenvalues for the finite difference operator $\delta_\times$.

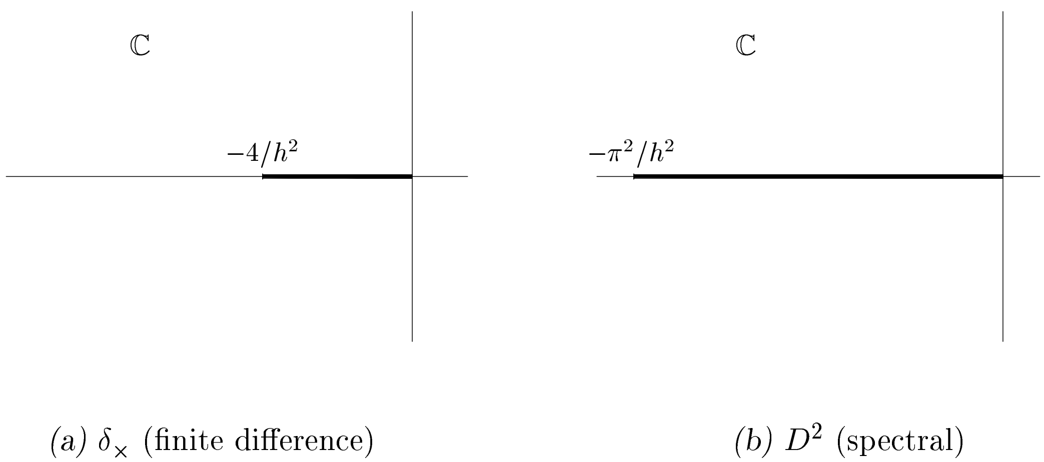

Figure 7.2.5. Eigenvalues of finite difference and spectral second-order differentiation matrices, as subsets of the complex plane.

difference operator $\delta_\times$ of §3.2. The spectrum of D^2 must be real, since D is symmetric; it is the interval $[-\pi^2/h^2, 0]$.

The developments of this section are summarized, and generalized to the mth-order case, in the following theorem:

SPECTRAL DIFFERENTIATION ON AN UNBOUNDED REGULAR GRID

Theorem 7.1. *The* m*th-order spectral differentiation operator* D^m *is a bounded linear operator on* ℓ_h^2 *with norm*

$$\|D^m\| = (\frac{\pi}{h})^m. \qquad (7.2.14)$$

If m *is odd,* D^m *has the imaginary spectrum* $[-i(\pi/h)^m, i(\pi/h)^m]$ *and can be represented by an infinite skew-symmetric Toeplitz matrix with entries*

$$a_0 = 0, \qquad a_j = h^{-m}(\qquad ? \qquad) \qquad \text{for} \quad j \neq 0. \qquad (7.2.15)$$

If m *is even,* D^m *has the real spectrum* $(-1)^{m/2} \times [0, (\pi/h)^m]$ *and can be represented by an infinite symmetric Toeplitz matrix with entries*

$$a_0 = \quad ? \quad, \qquad a_j = h^{-m}(\qquad ? \qquad) \qquad \text{for} \quad j \neq 0. \qquad (7.2.16)$$

The purpose of all of these spectral differentiation matrices is to solve partial differential equations. In a spectral collocation computation this is done in the most straightforward way possible: one discretizes the continuous problem as usual and integrates forward in time by a discrete formula, usually by finite differences.* Spatial derivatives are approximated by the matrix D. This same prescription holds regardless of whether the partial differential equation has variable coefficients or nonlinear terms. For example, to solve $u_t = a(x)u_x$ by spectral collocation, one approximates $a(x)u_x$ at each time step by $a(x_j)Dv$. For $u_t = (u^2)_x$, one uses $D(v^2)$, where v^2 denotes the pointwise square $(v^2)_i = (v_i)^2$. (Alternative discretizations may also be used for better stability properties; see....) This is in contrast to Galerkin or tau spectral methods, in which one adjusts the coefficients in the Fourier expansion of v to satisfy the partial differential equation globally.

* Spectral approximations with respect to time can also sometimes be used; see H. Tal-Ezer, "Spectral methods in time for hyperbolic equations," *SIAM J. Numer. Anal.* *23* (1986), pp. 11–26.

EXERCISES

▷ *7.2.1. Coefficients of D^3.* Determine the matrix coefficients of the third-order spectral differentiation matrix D^3. Compare your result with the coefficients of Table 3.3.1.

▷ *7.2.2.*

(a) Compute the integral $\int_{-\infty}^{\infty} S_h(x)dx$ of the sinc function (7.2.3). One could do this by complex contour integration, but instead, be cleverer than that and find the answer by considering the Fourier transform. The argument is quite easy; be sure to state it precisely.

(b) By considering the trapezoid rule for integration (Exercise 2.2.4), explain why the answer above had to come out as it did. (*Hint:* what is the integral of a constant function?)

7.3. Periodic grids

To be implemented in practice, a spectral method requires a bounded domain. In this section we consider the case of a periodic domain—or equivalently, an unbounded domain on which we permit only functions with a fixed periodicity. The underlying mathematics of discrete Fourier transforms was described in §2.4. The next chapter will deal with more general bounded problems.

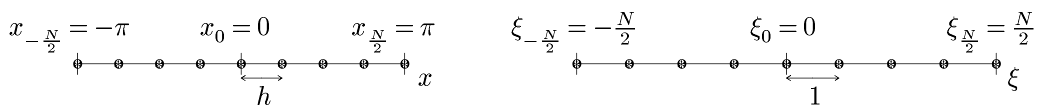

Figure 7.3.1. Space and wave number domains for the discrete Fourier transform.

To repeat some of the material of §2.4, our fundamental spatial domain will now be $[-\pi,\pi)$, as illustrated in Figure 7.3.1. Let N be a positive even integer, set

$$h = \frac{2\pi}{N} \qquad (N \text{ even}), \tag{7.3.1}$$

and define $x_j = jh$ for any j. The grid points in the fundamental domain are

$$x_{-N/2} = -\pi, \ldots, \quad x_0 = 0, \ldots, \quad x_{N/2-1} = \pi - h,$$

and the "invaluable identity" is this:

$$\frac{N}{2} = \frac{\pi}{h}. \tag{7.3.2}$$

The 2-norm $\|\cdot\|$ and space ℓ_N^2 were defined in §2.4, as was the **discrete Fourier transform,**

$$\hat{v}_\xi = (\mathcal{F}_N v)_\xi = h \sum_{j=-N/2}^{N/2-1} e^{-i\xi jh} v_j, \tag{7.3.3}$$

the **inverse discrete Fourier transform,**

$$v_j = (\mathcal{F}_N^{-1} \hat{v})_j = \frac{1}{2\pi} \sum_{\xi=-N/2}^{N/2-1} e^{i\xi jh} \hat{v}_\xi, \tag{7.3.4}$$

and the **discrete convolution,**

$$(v * w)_k = h \sum_{j=-N/2}^{N/2-1} v_{k-j} w_j. \tag{7.3.5}$$

Recall that since the spatial domain is periodic, the set of wave numbers ξ is discrete—namely the set of integers $\mathbb{Z}$—and this is why we have chosen to use ξ itself as a subscript in (7.3.3) and (7.3.4). We take $\xi = -N/2, \dots, N/2-1$ as our fundamental wave number domain. The properties of the discrete Fourier transform were summarized in Theorems 2.9 and 2.10; recall in particular the convolution formula

$$(\widehat{v*w})_\xi = \hat{v}_\xi \hat{w}_\xi. \tag{7.3.6}$$

As was described in §2.4, the discrete Fourier transform can be computed with great efficiency by the fast Fourier transform (FFT) algorithm, for which a program was given on p. 104. The discovery of the FFT in 1965 was an impetus to the development of spectral methods for partial differential equations in the 1970's. Curiously, however, practical implementations of spectral methods do not always make use of the FFT, but instead sometimes perform an explicit matrix multiplication. The reason is that in large-scale computations, which typically involve two or three space dimensions, the grid in each dimension may have as few as 32 or 64 points, or even fewer in so-called "spectral element" computations, and these numbers are low enough that the costs of an FFT and of a matrix multiplication may be roughly comparable.

Now we are ready to investigate $D_N : \ell_N^2 \to \ell_N^2$, the spectral differentiation operator for N-periodic functions. As usual, D_N can be described in various ways. One is by means of the discrete Fourier transform:

SPECTRAL DIFFERENTIATION BY THE DISCRETE FOURIER TRANSFORM.

(1) *Compute* $\hat{v}_\xi$;

(2) *Multiply by* $i\xi$, *except that* $\hat{v}_{-N/2}$ *is multiplied by* 0,

(3) *Inverse transform:*

$$D_N v = \mathcal{F}_N^{-1}(\,0 \text{ for } \xi = -N/2; \quad i\xi\hat{v}_\xi \text{ otherwise}). \tag{7.3.8}$$

The special treatment of the value $\hat{v}_{-N/2}$ is required to maintain symmetry, and appears only in spectral differentiation of *odd* order.

Another is in terms of interpolation by a finite series of complex exponentials or, equivalently, periodic sinc functions:

PERIODIC SPECTRAL DIFFERENTIATION BY SINC INTERPOLATION.

(1) *Interpolate* v *by a sum of periodic sinc functions*

$$q(x) = \sum_{k=-N/2}^{N/2-1} v_k S_N(x - x_k);$$

(2) *Differentiate the interpolant at the grid points* x_j:

$$(D_N v)_j = q'(x_j). \tag{7.3.9}$$

In the second description we have made use of the **periodic sinc function** on the N-point grid,

$$S_N(x) = \frac{\sin \frac{\pi x}{h}}{\frac{2\pi}{h} \tan \frac{x}{2}}, \tag{7.3.10}$$

which is the unique 2π-periodic function in L^2 that interpolates the discrete delta function e_j on the grid,

$$e_j = \begin{cases} 1 & j = 0, \\ 0 & j = -N/2, \dots, -1, 1, \dots, N/2-1, \end{cases} \tag{7.3.11}$$

and which is band-limited in the sense that its Fourier transform has compact support contained in $[-\pi/h, \pi/h]$ (and furthermore satisfies $\hat{v}_{-N/2} = \hat{v}_{N/2}$). Compare (7.2.3).

For higher order spectral differentiation on the periodic grid, we multiply $\hat{v}_\xi$ by higher powers of $i\xi$, or equivalently, differentiate $q(x)$ more than once. If the order of differentiation is odd, $\hat{v}_{-N/2}$ is multiplied by the special value 0 to maintain symmetry.

As in the last section, the spectral differentiation process can be viewed as multiplication by a skew-symmetric Toeplitz matrix D_N (compare (7.2.5)):

$$D_N = \begin{pmatrix} 0 & \frac{1}{2}\cot\frac{h}{2} & -\frac{1}{2}\cot\frac{2h}{2} & & & & -\frac{1}{2}\cot\frac{h}{2} \\ & \ddots & & & & & \\ & \frac{1}{2}\cot\frac{2h}{2} & -\frac{1}{2}\cot\frac{h}{2} & 0 & \frac{1}{2}\cot\frac{h}{2} & -\frac{1}{2}\cot\frac{2h}{2} & \\ & & & & & \ddots & \\ \frac{1}{2}\cot\frac{h}{2} & & & & \frac{1}{2}\cot\frac{2h}{2} & -\frac{1}{2}\cot\frac{h}{2} & 0 \end{pmatrix}. \qquad (7.3.12)$$

In contrast to the matrix D of (7.2.5), D_N is finite: it applies to vectors $(v_{-N/2},\dots,v_{N/2-1})$, and has dimension $N\times N$. D_N is not only a Toeplitz matrix, but is in fact **circulant**. This means that its entries $(D_N)_{ij}$ "wrap around," depending not merely on $i-j$ but on $(i-j)(\mathrm{mod}\,N)$.*

As in the last section, the entries of (7.3.12) can be derived either by the inverse discrete Fourier transform or by differentiating a sinc function. The latter approach is illustrated in Figure 7.3.2, which shows S_N and S_N' for $N=16$. Since N is even, symmetry implies that $S_N'(x)=0$ for $x=\pm\pi$ as well as for $x=0$. Differentiation yields

$$S_N'(x) = \frac{\cos\frac{\pi x}{h}}{2\tan\frac{x}{2}} - \frac{\sin\frac{\pi x}{h}}{\frac{4\pi}{h}\sin^2\frac{x}{2}}, \qquad (7.3.13)$$

and at the grid points the values are

$$S_N'(x_j) = \begin{cases} 0 & \text{if } j=0, \\ \frac{1}{2}(-1)^j\cot\frac{jh}{2} & \text{if } j\neq 0. \end{cases} \qquad (7.3.14)$$

Notice that for $|jh|\ll 1$, these values are approximately the same as in (7.2.8). Thus the (i,j) entry of D_N, as indicated in (7.3.12), is

$$(D_N)_{ij} = \begin{cases} 0 & \text{if } i=j, \\ \frac{1}{2}(-1)^{i+j}\cot(\frac{x_i-x_j}{2}) & \text{if } i\neq j. \end{cases} \qquad (7.3.15)$$

* *Any* circulant matrix describes a convolution on a periodic grid, and is equivalent to a pointwise multiplication in Fourier space. In the case of D_N, that multiplication happens to be by the function $i\xi$.

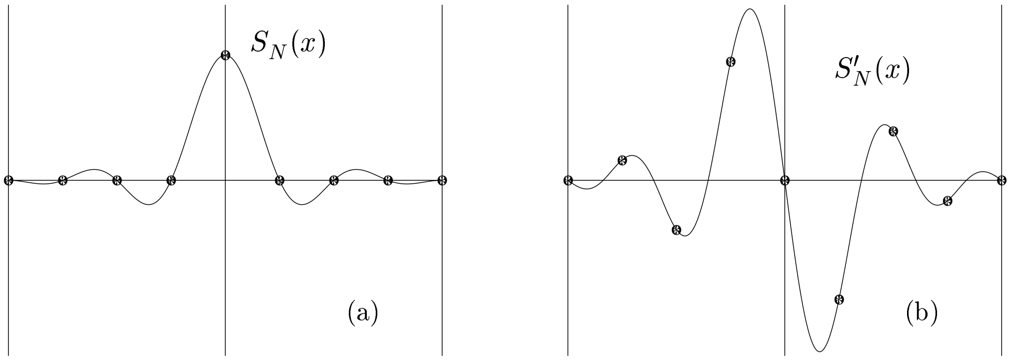

Figure 7.3.2. The periodic sinc function $S_N(x)$ and its derivative $S'_N(x)$.

To derive (7.3.15) by the Fourier transform, one can make use of "summation by parts," the discrete analog of integration by parts. See Exercise 7.3.1.

The eigenvectors of D_N are the vectors $e^{i\xi x}$ with $\xi \in \mathbb{Z}$, and the eigenvalues are the quantities $i\xi$ with $-N/2+1 \le \xi \le N/2-1$.* These are precisely the factors $i\xi$ in the definition of D_N by the Fourier transform formula (7.3.8). The number 0 is actually a double eigenvalue of D_N, corresponding to two distinct eigenvectors, the constant function and the sawtooth.

What about the spectral differentiation operator of second order? Again there are various ways to describe it. This time, because $-\xi^2$ is an even function, no special treatment of $\xi = -N/2$ is required to preserve symmetry in the Fourier transform description:

(1) Compute $\hat{v}_\xi$,
(2) Multiply by $-\xi^2$,
(3) Inverse transform:

$$D_N^{(2)} v = \mathcal{F}_N^{-1}(-\xi^2 \hat{v}_\xi). \tag{7.3.16}$$

The sinc interpolation description follows the usual pattern:

(1) Interpolate v by a sum of periodic sinc functions

$$q(x) = \sum_{k=-N/2}^{N/2-1} v_k S_N(x - x_k);$$

(2) Differentiate the interpolant twice at the grid points x_j,

$$(D_N^{(2)} v)_j = q''(x_j). \tag{7.3.17}$$

The matrix looks like this (compare (7.2.12)):

*Now that D_N is finite, they are truly eigenvalues; there are no technicalities to worry about as in the footnote on p. 241.

$$D_N^{(2)} = \begin{pmatrix} -\frac{\pi^2}{3h^2}-\frac{1}{6} & \frac{1}{2}\csc^2\frac{h}{2} & -\frac{1}{2}\csc^2\frac{2h}{2} & & & & \frac{1}{2}\csc^2\frac{h}{2} \\ & \ddots & & & & & \\ & -\frac{1}{2}\csc^2\frac{2h}{2} & \frac{1}{2}\csc^2\frac{h}{2} & -\frac{\pi^2}{3h^2}-\frac{1}{6} & \frac{1}{2}\csc^2\frac{h}{2} & -\frac{1}{2}\csc^2\frac{2h}{2} & \\ & & & & & \ddots & \\ \frac{1}{2}\csc^2\frac{h}{2} & & & -\frac{1}{2}\csc^2\frac{2h}{2} & \frac{1}{2}\csc^2\frac{h}{2} & & -\frac{\pi^2}{3h^2}-\frac{1}{6} \end{pmatrix} \tag{7.3.18}$$

Note that because $\xi = -N/2$ has been treated differently in the two cases, $D_N^{(2)}$ is not the square of D_N, which is why we have put the superscript in parentheses. In general, the mth-order spectral differentiation operator can be written as a power of $D_N^{(2)}$ if m is even, and as a power of D_N (or as D_N times a power of $D_N^{(2)}$) if m is odd. See Exercise 7.3.2.

Figures 7.3.3–7.3.6 are the analogs of Figures 7.2.2–7.2.5 for a periodic grid.

We summarize and generalize the developments of this section in the following theorem:

SPECTRAL DIFFERENTIATION ON A PERIODIC GRID

Theorem 7.2. *Let N be even. If m is odd, the mth-order spectral differentiation matrix $D_N^{(m)}$ is a skew-symmetric matrix with entries*

$$a_0 = 0, \qquad a_j = (\qquad ? \qquad) \qquad \text{for} \quad j \neq 0, \tag{7.3.19}$$

eigenvalues $[-i(\frac{\pi}{h}-1)^m, i(\frac{\pi}{h}-1)^m]$, and norm

$$\|D_N^{(m)}\| = \left(\frac{\pi}{h}-1\right)^m. \tag{7.3.20}$$

If m is even, $D_N^{(m)}$ is a symmetric matrix with entries

$$a_0 = \quad ? \quad, \qquad a_j = (\qquad ? \qquad) \qquad \text{for} \quad j \neq 0, \tag{7.3.21}$$

eigenvalues $(-1)^{m/2} \times [0, (\frac{\pi}{h})^m]$, and norm

$$\|D_N^{(m)}\| = \left(\frac{\pi}{h}\right)^m. \tag{7.3.22}$$

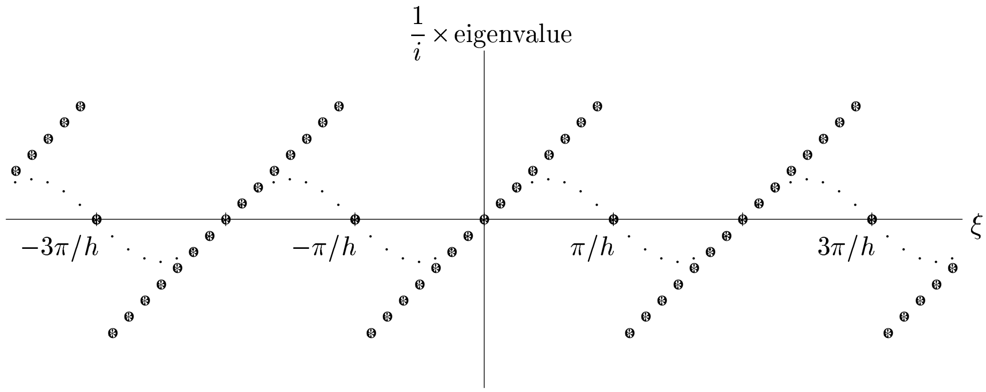

Figure 7.3.3. Eigenvalue of D_N (divided by i) corresponding to the eigenfunction $e^{i\xi x}$, as a function of ξ, for $N = 16$. The smaller dots show corresponding eigenvalues for the finite difference operator δ_0.

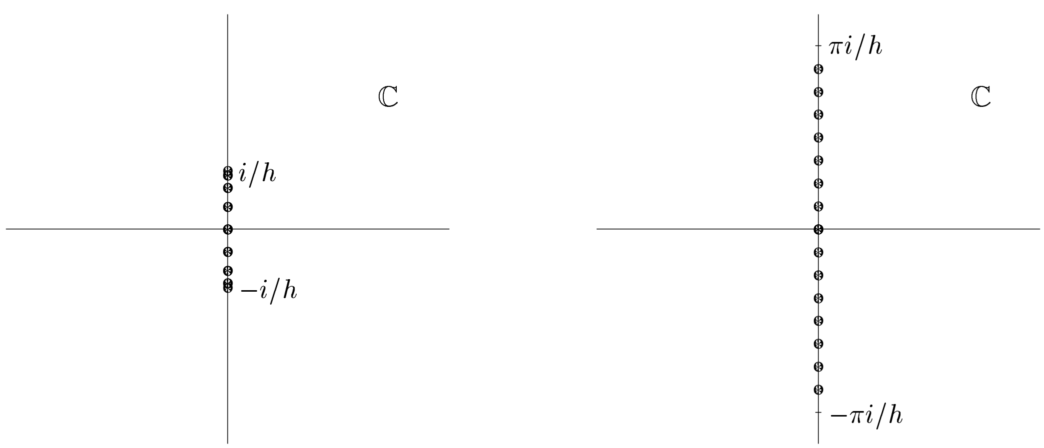

(a) δ_0 (finite difference) *(b)* DN (spectral)

Figure 7.3.4. Eigenvalues of finite difference and spectral first-order differentiation matrices on a periodic grid, as subsets of the complex plane, for $N = 16$.

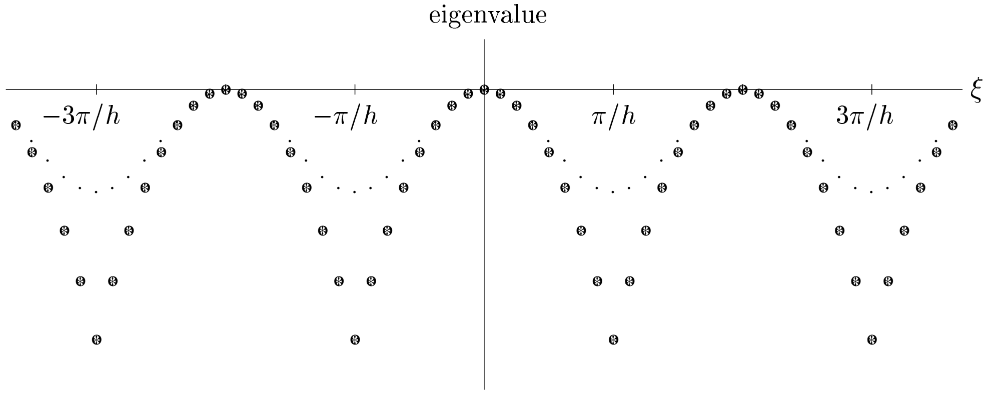

Figure 7.3.5. Eigenvalue of $D_N^{(2)}$ corresponding to the eigenfunction $e^{i\xi x}$, as a function of ξ, for $N=16$. The smaller dots show corresponding eigenvalues for the finite difference operator $\delta_\times$.

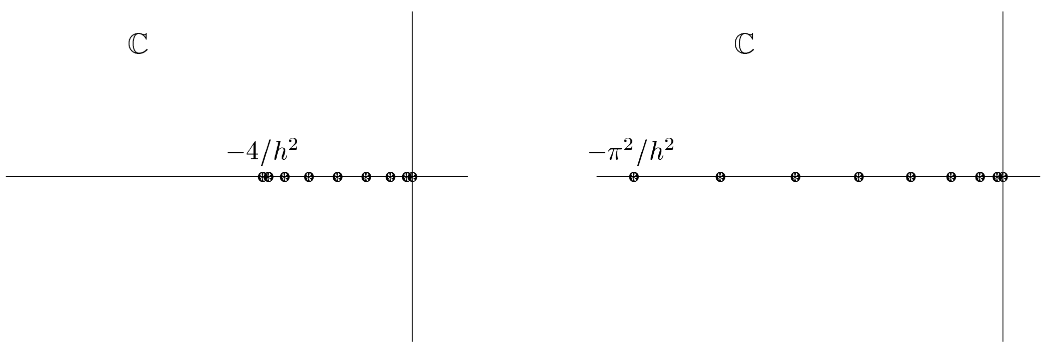

(a) $\delta_\times$ (finite difference) *(b)* D_N^2 (spectral)

Figure 7.3.6. Eigenvalues of second-order finite difference and spectral differentiation matrices on a periodic grid, as subsets of the complex plane, for $N=16$.

EXERCISES

▷ *7.3.1. Fourier transform derivation of D_N.* [Not yet written.]

▷ *7.3.2.* $D_N^{(2)} \neq D_N^2$. For most values of N, the matrix D_N^2 would serve as quite a good discrete second-order differentiation operator, but as mentioned in the text, it is not identical to $D_N^{(2)}$.

(a) Determine D_N, D_N^2, and $D_N^{(2)}$ for $N=2$, and confirm that the latter two are not equal.
(b) Explain the result of *(a)* by considering sinc interpolation as in Figure 7.3.2.
(c) Explain it again by considering Fourier transforms as in Figure 7.3.3 and 7.3.5.
(d) Give an exact formula for the eigenvalues of $D_N^{(J)}$ for arbitrary $J \geq 0$.

▶ *7.3.3. Spectral differentiation.*

Making use of a program for computing the FFT (in Matlab such a program is built-in; in Fortran one can use the program of p. 102), write a program DERIV that computes the mth-order spectral derivative of an N-point data sequence v representing a function defined on $[-\pi,\pi]$ or $[0,2\pi]$:

N: length of sequence (power of 2) (input)
m: order of derivative (integer ≥ 0) (input)
v: sequence to be differentiated (real sequence of length N) (input)
w: mth spectral derivative of v (real sequence of length N) (output)

Although a general FFT code deals with complex sequences, make v and w real variables in your program, since most applications concern real variables. Allow m to be any integer $m \geq 0$, and make sure to treat the distinction properly between even and odd values of m.

Test DERIV with $N=32$ and $N=64$ on the functions

$$u_1(x) = \exp(\sin 3x) \qquad \text{and} \qquad u_2(x) = |\sin x|^3$$

for the values $m=1,2$, and hand in two 2×2 tables—one for each function—of the resulting errors in the discrete ∞-norm. Put a star * where appropriate, and explain your results. Plot the computed derivatives with $m=1$.

▶ *7.3.4. Spectral integration.* Modify DERIV to accept negative as well as positive values of m. For $m<0$, DERIV should return a scalar representing the definite integral of v over one period, together with a function w representing the $|m|$th indefinite integral. Explore the behavior of DERIV with various v and $m<0$.

▶ *7.3.5. Hamming window.* Write a program FILTER that takes a sequence v and smooths it by transforming to $\hat{v}$, multiplying the transform by the "Hamming window",

$$\hat{w}_k = (0.54 + 0.46\cos\frac{2\pi k}{N})\hat{v}_k,$$

and inverse transforming. Apply FILTER to the function $|\sin x|$ and hand in a plot of the result.

-

▷ *8.3.1.* Fourier transform derivation of D_N. [Not yet written.]

▷ *8.3.2.* $D_N^{(2)} \neq D_N^2$. For most values of N, the matrix D_N^2 would serve as quite a good discrete second-order differentiation operator, but as mentioned in the text, it is not identical to $D_N^{(2)}$.

*(a)*Determine D_N, D_N^2, and $D_N^{(2)}$ for $N=2$, and confirm that the latter two are not equal.

*(b)*Explain the result of *(a)*by considering sinc interpolation as in Figure 8.3.3.

*(c)*Explain it again by considering Fourier transforms as in Figures 8.3.4 and 8.3.6.

▶ *8.3.3.* Spectral differentiation. Type the program FFT of Figure 8.3.2 into your computer and experiment with it until you are confident you understand how to compute both a discrete Fourier transform and an inverse discrete Fourier transform. For example, try as input a sine wave and a sinc function, and make sure the output you get is what you expect. Then make sure you can get the input back again by inversion.

Making use of FFT, write a program DERIV(N,m,v,w) which returns the mth-order spectral derivative of the N-point data sequence v representing a function defined on $[-\pi,\pi]$ or $[0,2\pi]$:

N: length of sequence (power of 2) (input)
m: order of derivative (integer ≥ 0) (input)
v: sequence to be differentiated (real sequence of length N) (input)
w: mth spectral derivative of v (real sequence of length N) (output)

Although FFT deals with complex sequences, make v and w real variables in your program, since most applications concern real variables. Allow m to be any integer $m \geq 0$, and make sure to treat the distinction properly between even and odd values of m.

Test DERIV with $N=32$ and $N=64$ on the functions

$$u_1(x) = \exp(\sin^3 x) \qquad \text{and} \qquad u_2(x) = |\sin x|^3$$

for the values $m=1,2$, and hand in two 2×2 tables—one for each function—of the resulting errors in the discrete ∞-norm. Explain your results. If possible, plot the computed derivatives with $m=1$.

▶ *8.3.4.* Spectral integration. Modify DERIV to accept negative as well as positive values of m. For $m<0$, DERIV should return a scalar representing the definite integral of v over one period, together with a function w representing the $|m|$th indefinite integral. Explore the behavior of DERIV with various v and $m<0$.

▶ *8.3.5.* Hamming window. Write a program FILTER(N,v,w) which takes a sequence v and smooths it by transforming to $\hat{v}$, multiplying the transform by the "Hamming window",

$$\hat{w}_k = (.54 + .46\cos\frac{2\pi k}{N})\hat{v}_k,$$

and inverse transforming. Apply FILTER to the function $|\sin x|$ and hand in the result. A plot would be nice.

7.4. Stability

[Just a few results so far. The finished section will be more substantial.]

Spectral methods are commonly applied to time-dependent problems according to the "method of lines" prescription of §3.3: first the problem is discretized with respect to space, and then the resulting system of ordinary differential equations is solved by a finite difference method in time. As usual, we can investigate the eigenvalue stability of this process by examining under what conditions the eigenvalues of the spectral differentiation operator are contained in the stability region of the time discretization formula. The separate question of stability in the sense of the Lax Equivalence Theorem is rather different, and involves some subtleties that were not present with finite difference methods; these issues are deferred to the next section.

In §§7.2,7.3 we have introduced two families of spectral differentiation matrices: D and its powers D^m for an infinite grid, and D_N and its higher-order analogs $D_N^{(m)}$ (not exactly powers) for a periodic grid. The spectra of all of these matrices lie in the closed left half of the complex plane, and that is the same region that comes up in the definition of A-stability. We conclude that if any equation

$$u_t = \frac{\partial^m u}{\partial x^m} \tag{7.4.1}$$

is modeled on a regular grid by spectral differentiation in space and an A-stable formula in time, the result is eigenvalue stable, regardless of the time step.

By Theorem 1.13, an A-stable linear multistep formula must be implicit. For a model problem as simple as (7.4.1), the system of equations involved in the implicit formula can be solved quickly by the FFT, but in more realistic problems this is often not true. Since spectral differentiation matrices are dense (unlike finite difference differentiation matrices), the implementation of implicit formulas can be a formidable problem. Therefore it is desirable to look for explicit alternatives.

On a regular grid, satisfactory explicit alternatives exist. For example, suppose we solve $u_t = u_x$ by spectral differentiation in space and the midpoint formula (1.2.6) in time. The stability region for the midpoint formula is the complex interval $[-i/k, i/k]$. From Theorem 7.1 or Figure 7.2.3, we conclude that the time-stability restriction is $\pi/h \le 1/k$, i.e.

$$k \le \frac{h}{\pi}. \tag{7.4.2}$$

This is stricter by a factor π than the time-stability restriction $k \le h$ for the leap frog formula, which is based on second-order finite difference differentiation. The explanation goes back to the fact that the sawtooth curve in Figure 7.2.2 is π times taller than the dashed one.

On a periodic grid, Theorem 7.2 or Figure 7.3.4 loosens (7.4.2) slightly to

$$k \le \frac{h}{\pi - h} = \frac{2}{N-2}. \tag{7.4.3}$$

Other explicit time-discretization formulas can also be used with $u_t = u_x$, so long as their stability regions include a neighborhood of the imaginary axis near the origin. Figure

1.7.4 reveals that this is true, for example, for the Adams-Bashforth formulas of orders 3–6. The answer to Exercise 1.7.2(b) can readily be converted to the exact stability bound for the 3rd-order Adams-Bashforth discretization.

For $u_t = u_{xx}$, the eigenvalues of D or D_N become real and negative, so we need a stability region that contains a segment of the negative real axis. Thus the midpoint rule will be unstable. On the other hand the Euler formula, whose stability region was drawn in Figure 1.7.3 and again in Figure 1.7.4, leads to the stability restriction $\pi^2/h^2 \le 2/k$, i.e.

$$k \le \frac{2h^2}{\pi^2}. \tag{7.4.4}$$

On an infinite grid this is $\pi^2/4$ stricter than the stability restriction for the finite difference forward Euler formula considered in Example 4.4.3.* (The cusped curve in Figure 7.2.4 is $\pi^2/4$ times deeper than the dashed one.) By Theorem 7.2, exactly the same restriction is also valid for a periodic grid:

$$k \le \frac{2h^2}{\pi^2} = \frac{8}{N^2}. \tag{7.4.5}$$

As another example, the answer to Exercise 1.7.2(a) can be converted to the exact stability bound for the 3rd-order Adams-Bashforth discretization of $u_t = u_{xx}$.

In general, spectral methods on a periodic grid tend to have stability restrictions that are stricter by a constant factor than their finite difference counterparts. (This is opposite to what the CFL condition might suggest: the numerical domain of dependence of a spectral method is unbounded, so there is no CFL stability limit.) The constant factor is usually not much of a problem in practice, for spectral methods permit larger values of h than finite difference methods in the first place, because of their high order of spatial accuracy. In other words, relatively small time steps k are needed anyway to avoid large time-discretization errors.

EXERCISES

▶ *7.4.1. A simple spectral calculation.*

Write a program to solve $u_t = u_x$ on $[-\pi,\pi]$ with periodic boundary conditions by the pseudospectral method. The program should use the midpoint time integration formula and spatial differentiation obtained from the program DERIV of Exercise 7.3.3.

(a) Run the program up to $t = 2\pi$ with $k = h/4$, $N = 8$ and $N = 16$, and initial data $v^0 = f(x)$, $v^1 = f(x+k)$, with both

$$f_1(x) = \cos^2 x \qquad \text{and} \qquad f_2(x) = \begin{cases} \cos^2 x & \text{for } |x| \le \pi/2 \,(\mathrm{mod}\, 2\pi), \\ 0 & \text{otherwise.} \end{cases}$$

List the four ℓ_N^∞ errors you obtain at $t = 2\pi$. Plot the computed solutions $v(x,2\pi)$ if possible.

(b) Rerun the program with $k = h/2$ and list the same four errors as before.

(c) Explain the results of *(a)* and *(b)*. What order of accuracy is observed? How might it be improved?

*Why is the ratio not π^2? Because $\delta_\times = \delta_0(h/2)^2$, not $\delta_0(h)^2$.

▶ *7.4.2. Spectral calculation with filtering.*

(a) Modify the program above so that instead of computing v^n at each step with DERIV alone, it uses DERIV followed by the program FILTER of Exercise 7.3.5. Take $k=h/4$ again and print and plot the same results as in the last problem. Are the errors smaller than they were without FILTER?

(b) Run the program again in DERIV/FILTER mode with $k=h/2$. Print and plot the same results as in *(c)*.

(c) What is the exact theoretical stability restriction on k for the DERIV/FILTER method? You may consider the limit $N=\infty$ for simplicity.

Note. Filtering is an important idea in spectral methods, but this simple linear problem is not a good example of a problem where filtering is needed.

▶ *7.4.3. Inviscid Burgers' equation.*

Write a program to solve $u_t=(u^2)_x$ on $[0,2\pi]$ with periodic boundary conditions by the pseudospectral method. The program should use forward Euler time integration formula and spatial differentiation implemented with the program DERIV.

(a) Run the program up to $t=2\pi$ with $k=h/8$, $N=32$, and initial data $v^0=0.3+0.08\sin x$. Plot the computed results (superimposed on a single plot) at times $t=0, \pi/4, \pi/2, \dots, 2\pi$.

(b) Explain the results of part *(a)* as well as you can.

(c) *(Optional.)* Can you find a way to improve the calculation?