Journal of Computational Physics **178,** 118–160 (2002)
doi:10.1006/jcph.2002.7023, available online at http://www.idealibrary.com on IDEAL®

A Comparison of Numerical Algorithms for Fourier Extension of the First, Second, and Third Kinds

John P. Boyd

Department of Atmospheric, Oceanic and Space Science, University of Michigan, 2455 Hayward Avenue, Ann Arbor, Michigan 48109-2143

Received July 6, 2001; revised February 12, 2002

The range of Fourier methods can be significantly increased by extending a nonperiodic function $f(x)$ to a periodic function $\tilde{f}$ on a larger interval. When $f(x)$ is analytically known on the extended interval, the extension is straightforward. When $f(x)$ is unknown outside the physical interval, there is no standard recipe. Worse still, like a radarless aircraft groping through fog, the algorithm may wreck on the "mountain-in-fog" problem: a function $f(x)$ which is perfectly well behaved on the physical interval may very well have singularities in the extended domain. In this article, we compare several algorithms for successfully extending a function $f(x)$ into the "fog" even when the analytic extension is singular. The best third-kind extension requires singular value decomposition with iterative refinement but achieves accuracy close to machine precision.

Key Words: local Fourier basis; wavelets; pseudospectral algorithms; Fourier extension.

1. INTRODUCTION

1.1. The Need for Fourier Extension: Applying Fourier Methods to Nonperiodic Functions

If $f(x)$ can be expanded in a rapidly convergent Fourier series, then it is easy to find a particular solution to the inhomogeneous Poisson equation

$$u_{xx} - u = -f(x) \tag{1}$$

since

$$f(x) = \sum_{j=-N}^{N} f_j \exp(ijx) \to u_P(x) = \sum_{j=-N}^{N} \frac{f_j}{1+j^2} \exp(ijx), \tag{2}$$

0021-9991/02 $35.00

as follows by term-by-term matching. Boundary conditions such as

$$u(\pm\chi) = 0 \tag{3}$$

can easily be enforced by writing

$$u(x) = u_P + A\exp(x) + B\exp(-x) \tag{4}$$

and choosing the constants A and B so that the boundary conditions are satisfied. Because the Fourier coefficients of $f(x)$ can be computed by the fast Fourier transform (FFT) in $O(N \log_2 N)$ operations, the overall algorithm is extremely fast.

The flaw is that if $f(x)$ is nonperiodic, even so simple a function as $f(x) \equiv x$, then the Fourier series will converge very slowly, usually with $f_j \sim O(1/j)$ as $j \to \infty$ [6]. However, if $f(x)$ can be extended to a periodic function $\tilde{f}$ on the larger interval $x \in [-\Theta, \Theta]$, then the Fourier–Poisson algorithm will work as advertised. This and similar applications have motivated a great deal of interest in what we dub the "Fourier extension problem": how to systematically construct such periodic functions $\tilde{f}$ from a given nonperiodic $f(x)$ [5, 10–15, 20].

1.2. The Definition of "Fourier Extension"

DEFINITION 1.1 (Fourier Extension Problem). Given a (generally) nonperiodic function $f(x)$ on $x \in [-\chi, \chi]$, the Fourier extension problem is to define a function $\tilde{f}$ on a larger interval $x \in [-\Theta, \Theta]$ such that

1.

$$\tilde{f} \equiv f \quad \forall x \in [-\chi, \chi], \tag{5}$$

2. $\tilde{f}$ is periodic with period 2Θ, and
3. $\tilde{f}$ has a rapidly convergent Fourier series.

The extension is of the first kind when $f(x)$ is known and analytical everywhere on the extended interval $x \in [-\Theta, \Theta]$, of the second kind when $f(x)$ is known but has singularities on either or both of the "extension intervals" $x \in [-\Theta, -\chi]$ and $x \in [\chi, \Theta]$, and of the third kind when $f(x)$ is not known outside the "physical" interval $x \in [-\chi, \chi]$.

Fourier extension of the first kind is rather easy because one can multiply $f(x)$ by a function $\mathcal{T}$ to obtain $\tilde{f} \equiv \mathcal{T} f(x)$ provided that the "bell" or "smoothed top hat" $\mathcal{T}$ has the following properties:

1. $\mathcal{T} \equiv 1\ \forall x \in [-\chi, \chi]$.
2. $\mathcal{T}$ is infinitely differentiable everywhere on the extended interval $x \in [-\Theta, \Theta]$.
3. $\mathcal{T}$ is "infinitely flat" at $x = \pm\chi, \pm\Theta$ in the sense that the bell and all its derivatives are zero at these points.

These properties guarantee that $\tilde{f}$ is a good extension in the sense that its Fourier series converges exponentially fast, and yet even so $\tilde{f} \equiv f(x)$ everywhere on the physical interval, $x \in [-\chi, \chi]$. Bell functions with these properties are extremely useful in wavelets and local Fourier basis methods and so several good families are available, as explained in [5, 18].

Parenthetically, it should be noted that the extension can be more efficiently performed by slightly modifying the simple product construction [5]; more sophisticated first-kind methods are explained and compared in Section 2. The $\tilde{f} \equiv \mathcal{T} f$ variant suffices to show that extension of the first kind is a solved problem.

When $f(x)$ has known singularities on the extension zones, that is, at some point x_s in either of the intervals $x \in [-\Theta, -\chi]$ and $x \in [\chi, \Theta]$, the product formula also works provided only that the bell is narrowed so that it becomes infinitely flat at the location of the singularity closest to the physical domain instead of at $x = \pm\Theta$. (Other, sometimes better strategies are described and compared in Section 3.) Thus, although the rate of Fourier convergence is slowed when the bell is narrowed, extension of the second kind is a solved problem, too.

In many applications, however, $f(x)$ is known only on the physical domain $x \in [-\chi, \chi]$. For example, if $f(x)$ is an ocean stress due to wind and the coasts are at $x = \pm\chi$, then we have no data—and no water—in the extension zones $|x| > \chi$. Nevertheless, we can still mathematically define an extended function $\tilde{f}$, specified both on land and water, which matches the measured $f(x)$ in the ocean. However, outside the physical domain we are, metaphorically, flying in the unknown as much as a radarless 1920s mail plane groping its way through thick fog.

1.3. Uniqueness of Extension

Fortunately, there is an infinite variety of extensions which are infinitely differentiable (and therefore in the function class "C^∞"), but are not analytic, and therefore not obliged to blow up at the singularities of the analytic continuation of $f(x)$. Indeed, the space of C^∞ extensions is so wide that the extension problem is rather ill conditioned.

THEOREM 1.1 (Uniqueness/Nonuniqueness of Extension). *If a function $f(x)$ is analytic everywhere on the extended interval $x \in [-\Theta, \Theta]$ except perhaps for isolated singularities, then the analytic continuation of $f(x)$ from the physical interval to the larger extended interval is unique, modulo choices of branch cuts, and so forth. If analyticity is relaxed to the milder condition of being C^∞, that is, infinitely differentiable with bounded derivatives of all orders everywhere on the extended interval, then any C^∞ extension is not unique. If $\tilde{f}$ is a C^∞ extension, then another extension in this same class is given by*

$$\tilde{g} \equiv \tilde{f} + \Psi(x), \tag{6}$$

where $\Psi(x)$ is any C^∞ function which is (i) infinitely flat at $x = \pm\chi$ and (ii) identically zero everywhere on the physical domain $x \in [-\chi, \chi]$.

Proof. With specification of all branch cuts, complex variable theory shows that the analytic continuation of a function $f(x)$ is unique, or, in other words, two analytic functions cannot agree everywhere on an *interval* of finite length and yet disagree elsewhere (p. 65 of [8]). The nonuniqueness of C^∞ functions is obvious since $\tilde{g} \equiv f(x)\ \forall x \in [-\chi, \chi]$; a C^∞ bell (such as $\mathcal{T}$ in the next section) of sufficiently narrow width to fit in the extension zones provides an explicit example of a function $\Psi(x)$ which is zero everywhere on the physical interval. ■

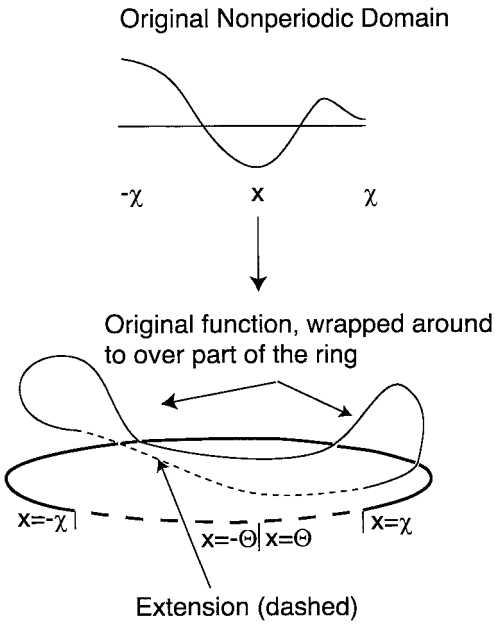

FIG. 1. Schematic of Fourier extension from the original, physical interval $x \in [-\chi, \chi]$ to a larger interval. Because the extended function $\tilde{f}$ is periodic on the extended interval $x \in [-\Theta, \Theta]$, it is conceptually useful to wrap the extended interval into a circle so that the points $x = \pm\Theta$ coincide. In both panels, the thin curve is the original function $f(x)$; the extension is dashed in the lower panel. The extension zones, $x \in [-\Theta, -\chi]$ and $x \in [\chi, \Theta]$, become a single continguous interval (heavy dashes) when the periodic domain is wrapped into a circle.

1.4. Derivative Matching Conditions: Why a Good Extension Must Be HIP

DEFINITION 1.2 (Hermite Interpolation Property (HIP)). An extension is said to have the "Hermite interpolation property" to order k, or, for brevity, to be "HIP" to order k, if $\tilde{f}$, $\tilde{f}^{(1)}, \ldots, \tilde{f}^{(k)}$ for the extended function match the corresponding derivatives of the function $f(x)$ which is being extended through order k at both endpoints of the physical interval.

The reason for the name is that the usual two-point Hermite polynomial interpolant of degree $2k + 1$ is HIP in the sense defined here, although this property applies to extensions which are not necessarily polynomials, or Hermite interpolants in any sense except this derivative-matching-at-the-endpoints property. The significance of being HIP is expressed by the following.

THEOREM 1.2 (Necessity of Derivative Matching (HIP)). *If the extended function $\tilde{f}$ has Fourier cosine coefficients a_j and sine coefficients b_j which decrease at least as fast as $O(1/j^{k+2})$ and if $\tilde{f}$ perfectly matches $f(x)$ everywhere on the physical interval, then the extension must possess the Hermite interpolation property (HIP) of order k.*

Proof. If $\tilde{f}$ is HIP to order k but not order $k + 1$ and exactly equals f everywhere on the physical interval, then the $(k + 1)$st derivative of $\tilde{f}$ must be discontinous at the endpoints of the physical domain. The usual integration-by-parts argument used to derive the Fourier asymptotic coefficient expansion (FACE) in Ch. 2 of [6] then shows that the Fourier coefficients of $\tilde{f}$ will decrease no faster than $O(1/j^{k+2})$. ■

Thus, if the extended function is to have a rapidly converging Fourier series, then $\tilde{f}$ must be HIP to high order.

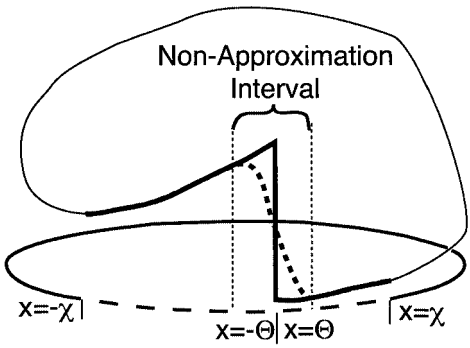

FIG. 2. The upper thin curve is $f(x)$ on the physical interval. The thick solid curve is the analytic extension, which has a jump discontinuity at $x = \pm\Theta$. The thick dashed curve is a "good" extension in which the discontinuity has been smoothed so that the Fourier expansion of the function defined by the thin-solid-plus-thick-dashed-curve converges rapidly. The vertical dotted lines bounded the region where the two extensions differ; this is the nonapproximation interval, where the smooth extension is a poor approximation to $f(x)$, the analytic extension.

The bell–imbricate methods described in the next two sections are HIP to *infinite* order. The reason is that $\tilde{f}$ is the product of f with a "bell" function $\mathcal{T}$ that is "infinitely flat" at the endpoints of the physical interval in the sense that all the derivatives of $\mathcal{T}$ are zero at these points.

In contrast, the FPIC-SU third-kind method of Section 4 does not explicitly impose HIP conditions. The resulting Fourier coefficient spectra (see Fig. 27, right panel) are rather flat except very near the truncation limit N. Even so (and fortunately), the approximations are still very accurate on the physical interval.

Because the derivatives that are matched are also the leading terms in the ordinary power series of both $f(x)$ and $\tilde{f}$ about the endpoints of the physical domain, it follows that the extension must accurately approximate $f(x)$ at least in a small neighborhood of both endpoints of the physical interval.

And yet, in dramatic contrast, all useful extension methods must also have a "nonapproximation property," defined as follows.

1.5. Nonapproximation Property (NAP)

One strategy for continuation is analytic continuation: $\tilde{f} \equiv f(x)$ everywhere on the entire extended domain, $x \in [-\Theta, \Theta]$. However, as illustrated in Fig. 1, a Fourier series approximation is *always periodic*; the best way to visualize this is to bend the interval into a circle so that the points $x = \Theta$ and $x = -\Theta$ are the *same* point, exactly as they are treated as the same point by a Fourier series with period 2Θ. This implies that if $f(\Theta) \neq f(-\Theta)$, then the extended function $\tilde{f}$ on the *circle* $x \in [-\Theta, \Theta]$ will have a jump discontinuity at the "join," $x = \pm\Theta$, as illustrated in Fig. 2. The Fourier coefficients a_j will then converge at the awful rate of $a_j \sim O(1/j)$. Thus, the analytic extension is a horrible strategy for nonperiodic $f(x)$.[1]

This argument has larger implications: any extension scheme that accurately approximates $f(x)$ over the whole extended interval will inherit the curse of the analytic extension:

[1] In the special case that f is periodic with periodic 2Θ, the Fourier series converges rapidly, but the extension problem is completely trivial.

terribly slow convergence. It follows that a useful Fourier extension method must necessarily have the following property.

DEFINITION 1.3 (Nonapproximation Property (NAP)). An extension $\tilde{f}$ is said to have the nonapproximation property (NAP) if it fails to uniformly approximate $f(x)$ on the extension interval in the limit that the number of terms in the Fourier series tends to infinity.

In other words, the extended function $\tilde{f}$ must break away from $f(x)$—nonapproximating $f(x)$—in a neighborhood of $x = \pm\Theta$ so that $\tilde{f}$ is everywhere infinitely differentiable. The vertical dotted lines bound the nonapproximation interval where the smooth extension (thick dashes) is a poor approximation to $f(x)$ (thick solid curve with discontinuity).

In addition, a function $f(x)$ which is well behaved on the physical interval $x \in [-\chi, \chi]$ may have singularities in the extension domain. A good extension $\tilde{f}(x)$ must also have the nonapproximation property in the vicinity of all such singularities, as well as the neighborhood of $x = \pm\Theta$.

For third-kind extensions, by definition nothing is known about f in the extension interval. It follows that nothing is known in this case about the singularities of $f(x)$ and therefore nothing about the neighborhoods where a good extension must be a poor approximation to the singular function $f(x)$. This is the "mountain-in-fog" problem, discussed at greater length in Section 7.

These simultaneous conditions of HIP and NAP imply that a good extension must be both an approximation to $f(x)$ (near the endpoints of the physical interval) and simultaneously *not* an approximation where the two pieces of the extension interval are joined at $x = \pm\Theta$ and also in the neighborhoods of any singularities of $f(x)$ on the extension zone. It might seem impossible to satisfy two such conflicting requirements simultaneously.

For Fourier extensions of the first and second kind, however, it turns out that one can easily create explicit, well-conditioned $\tilde{f}$ which are both HIP and NAP by multiplying $f(x)$ by a bell which has two properties. First, the bell is *infinitely flat* at the ends of the physical interval (in the sense that all derivatives of the bell $\mathcal{T}$ are zero at that point); this automatically enforces continuity of the derivatives (HIP) to all orders for the product $f\mathcal{T}$, as follows by applying the Leibnitz rule that the derivative of a product $f(x)\mathcal{T}(x)$ is $fd\mathcal{T}/dx + df/dx\mathcal{T}$ for any pair of functions $f(x)$ and $\mathcal{T}(x)$. Second, the bell decays exponentially fast away from the physical interval, thus removing the discontinuity at $x = \pm\Theta$. One can further prove, as explained below, that such extensions have Fourier coefficients a_n that converge faster than $O(n^{-k})$ for any finite k provided only that $f(x)$ is analytic for real x on the physical and extension intervals.

For an extension of the third kind, that is, when $f(x)$ is known only on the physical interval, it is much harder to enforce the HIP and NAP properties simultaneously. We have tried several methods and one is successful. However, even this optimum third-kind algorithm requires singular value decomposition of the pertinent matrix. Nevertheless, its performance is roughly the same as the best bell–imbricate first-kind method on a set of test functions, as explained in Section 6.

1.6. Perfectly Faithful Extension

In this article, we voluntarily restrict ourselves to extensions that have the following property.

DEFINITION 1.4 (Perfectly Faithful Extension (PFE)). An extended function is a "perfectly faithful extension" if it exactly agrees with the original function on the original domain $x \in [-1, 1]$.

The desirability of a PFE extension is obvious, but there is a subtle drawback. Since a perfectly faithful extension exactly agrees with $f(x)$ over an interval, but is different from it elsewhere (the NAP), it follows that a PFE cannot be an analytic function over the whole extended domain. At most, it can be C^∞, that is, infinitely differentiable. This in turn implies [6] that Fourier coefficients of the extended function cannot converge geometrically, that is, as $\exp(-qj)$ for some constant $q > 0$, but only as $\exp(-pj^r)$ for some exponent $r < 1$, a "subgeometric" rate of convergence.

If one employs a bell which is analytic everywhere, then the extended function $\tilde{f}$ can be analytic on the entire extended interval, too, and can converge geometrically fast. One choice for such a bell, from an almost infinite range of possibilities, is

$$B(x; \chi, L) \equiv \frac{1}{2}\{\mathrm{erf}(L[x + \chi]) - \mathrm{erf}(L[x - \chi])\}. \tag{7}$$

In the limit that $L \to \infty$, this becomes a step function which is nonzero between $x \in [-\chi, \chi]$, and the product of $f(x)$ with B is a perfectly faithful extension, but the discontinuity at $x = \pm\chi$ destroys the exponential convergence of the Fourier series of $\tilde{f} \equiv fB$. However, for finite L, the Fourier convergence is geometric but the extension is no longer perfectly faithful because B is only *approximately* equal to one on the physical interval.

It is not clear that this is necessarily bad; a truncated Fourier series can only *approximate* $\tilde{f}$ anyway, so a tiny difference between f and $\tilde{f}$ on the physical interval need not significantly increase the error of the approximation of f by the truncated Fourier series of $\tilde{f}$. Nevertheless, PFEs are simpler and we shall stick to them for the present. The usefulness of imperfectly faithful extensions will be left as an open research problem.

2. A COMPARISON OF METHODS FOR FOURIER EXTENSION OF THE FIRST KIND

2.1. Description of Four Algorithms

In this section, we describe and then compare four different schemes for first-kind extensions:

1. "Naive" or "nonoverlapped"
2. "Overlapped"
3. "Alternating"
4. Hermite interpolant/two-point Taylor.

The fourth method is in principle a scheme for third-kind extensions, too, at least sometimes, and has a very different justification than the others. Consequently, the Hermite interpolant/two-point Taylor method is discussed in Section 5 below. However, this scheme is included in the numerical comparisons between the four first-kind methods given later in this present section.

The first three extension schemes are based on a pair of key elements: (i) a bell function and (ii) an imbricate series.

As noted earlier, a bell is a function which is approximately one on the physical region and then smoothly decays to zero as x moves away from the physical interval. There are

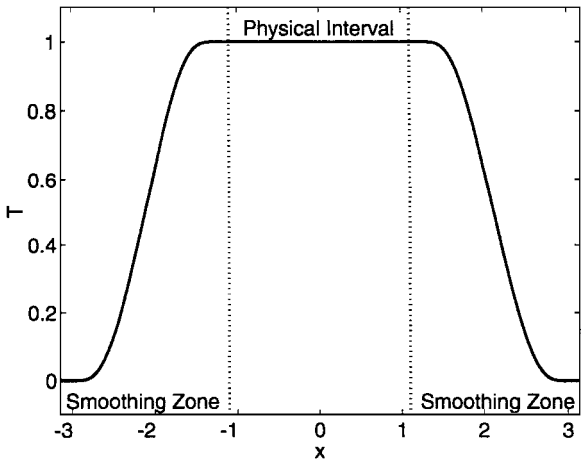

FIG. 3. Erf-smoothed top-hat function or bell $\mathcal{T}$. The function $\mathcal{T} \equiv 1$ for all $x \in [-\chi, \chi]$. χ can be chosen arbitrarily and is 1.1 in this graph, where $\chi = 1.1$, the physical region. The function smoothly varies from to 1 to 0 on the intervals $x \in [-\Theta, -\chi]$ and $x \in [\chi, \Theta]$, which are labeled "smoothing zones." Θ is also arbitrary and is here $\Theta = \pi$. The vertical dotted lines are the boundaries between the physical and smoothing intervals.

many possible choices, as noted earlier [3, 4, 18]. For expository purposes, it is sufficient to consider the particular example (Fig. 3)

$$\mathcal{T}(x; L, \chi, \Theta) \equiv \begin{cases} \mathcal{H}([x + \chi + \Xi]/\Xi; L), & x \in [-\Theta, -\chi] \\ 1, & x \in [-\chi, \chi] \\ \mathcal{H}(-[x - \chi - \Xi]/\Xi; L), & x \in [\chi, \Theta], \end{cases} \tag{8}$$

where $\Xi \equiv (\Theta - \chi)/2$. This bell tapers from one to zero on the intervals $x \in [-\Theta, -\chi]$ and $x \in [\chi, \Theta]$. The function $\mathcal{H}$ is a "ramp," that is, a smoothed approximation to a step function defined by

$$\mathcal{H}(x; L) \equiv (1/2)\{1 + \mathcal{E}(x; L)\}, \tag{9}$$

where $\mathcal{E}$ is an error-function-like function ("erfoid"),

$$\mathcal{E}(x; L) = \begin{cases} -1, & x < -1 \\ \operatorname{erf}\left(L \frac{x}{\sqrt{1 - x^2}}\right), & x \in [-1, 1] \\ 1, & x > 1. \end{cases} \tag{10}$$

The parameter L is a scaling factor that specifies how rapidly the erflike, ramp, and bell functions tend to their limits. (We return to the issue of choosing L later.) Thanks to the $1/\sqrt{1 - x^2}$, the argument of the error function varies from $-\infty$ to ∞ as x ranges from -1 to 1 so that $\mathcal{E}$ is infinitely flat (in the sense that all its derivatives vanish) at $x = \pm 1$. This allows the erflike function, and, by inheritance, the ramp and bell functions, to be infinitely differentiable for all real x (that is, to belong to the function class C^∞). The bell $\mathcal{T}$ is not analytic at the breakpoints $x = \pm\chi, \pm\Theta$ which bound the intervals where the ramp varies from 1 to 0.

The function $f(x)$ is multiplied by the bell $\mathcal{T}$ to create a "pattern" function $G(x)$:

$$G(x) \equiv \mathcal{T} f(x). \tag{11}$$

For the first two methods, the extended $\tilde{f}$ is defined as an "imbricate series,"

$$\tilde{f} \equiv \sum_{m=-\infty}^{\infty} G(x - mP), \tag{12}$$

where P is the spatial period. In words, the imbricate series is the result of duplicating this pattern function, placing copies with an even spacing P over the whole interval, and superimposing the copies. By construction, an imbricate series is always periodic with period P. Because $\tilde{f}$ is periodic, it is possible to expand it in a Fourier series with coefficients which decrease *exponentially* fast with degree.

In the naive or nonoverlapped method, the width of the bell is chosen so that there is no overlap between the copies of the pattern function,

$$G_{nonoverlapped} \equiv \mathcal{T}(x; L, \chi, \Theta) f(x), \tag{13}$$

where the spatial period is $P = 2\Theta$. Because the bell is identically zero for $|x| > \Theta$, the infinite series is unnecessary and we can simply write

$$\tilde{f}_{nonoverlapped} \equiv \mathcal{T}(x; L, \chi, \Theta) f(x), \quad x \in [-\Theta, \Theta]. \tag{14}$$

However, this strategy is not very efficient. The periodic interval, $x \in [-\Theta, \Theta]$, is best visualized as a *circle*, as in Fig. 1; the two smoothing zones are actually continguous with one another at the point $x = \Theta$, which is identical on the circle with $x = -\Theta$. It is more efficient to double the width of the smoothing zones so as to generate a more rapidly convergent Fourier series. This overlapped strategy is equivalent to choosing the pattern function to be

$$G_{overlapped} \equiv \mathcal{T}(x; L, \chi, 2\Theta - \chi) f(x), \tag{15}$$

$$\tilde{f}_{overlapped} \equiv \begin{cases} \mathcal{T}(x; L, \chi, 2\Theta - \chi) f(x) + \mathcal{T}(x + 2\Theta; L, \chi, 2\Theta - \chi) f(x + 2\Theta), & x \in [-\Theta, -\chi] \\ f(x), & x \in [-\chi, \chi] \\ \mathcal{T}(x; L, \chi, 2\Theta - \chi) f(x) + \mathcal{T}(x - 2\Theta; L, \chi, 2\Theta - \chi) f(x - 2\Theta), & x \in [\chi, \Theta]. \end{cases} \tag{16}$$

A variation on this has been employed by Averbuch *et al*. [1], who preferred a so-called "alternating imbricate" series,

$$\tilde{f} \equiv \sum_{m=-\infty}^{\infty} (\pm 1)^m G_{alternating}(x - mP/2), \tag{17}$$

where the pattern function is exactly the same as for the overlapped (but nonalternating) extension:

$$G_{alternating} = G_{overlapped} = \mathcal{T}(x; L, \chi, 2\Theta - \chi) f(x). \tag{18}$$

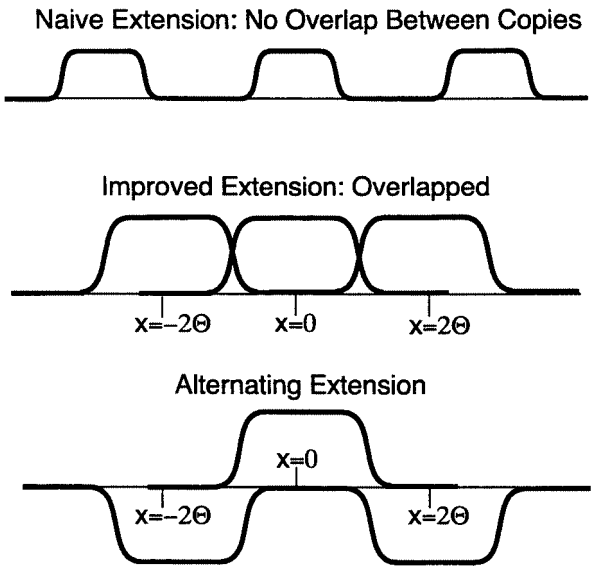

FIG. 4. The pattern function and two of its copies with the correct spacing, width, and sign for each of the three first-kind-extension schemes that are based on a bell function/imbricate series combination.

The apparent bad news is that the spatial period is *doubled* so that $P = 4\Theta$, twice the spacing between copies of the pattern function $G_{alternating}$. The good news is that because the cancellation between positive and negative copies of the pattern forces $\tilde{f}$ to have zeros at the ends of the extension interval, it is no longer necessary to use a general Fourier series to approximate it; only a sum of the *odd* sines and *odd* cosines is required. As we shall see in a later section, the result is a Fourier series that converges at almost exactly the same rate as for the nonalternating expansion. Figure 4 compares the three bell-shaped methods.

2.2. Width of the Extension Region

In the rest of the article, we assume that the width of the extension region, $\Theta - \chi$, is *fixed*. This is unrealistic in the sense that one always has the freedom to choose the width of the extension zone, and this does have an impact on the convergence of the Fourier series of the extended function.

However, a little reflection should convince us that the optimum width is highly dependent on the particular $f(x)$ which is being extended. If $f(x) \equiv 1$ or another very smooth function, then it is best to use a very wide extension zone so that the bell function can be as wide and as smooth as possible, thus maximizing the convergence rate of the extended function. However, if $f(x)$ is wildly oscillatory or has singularities in the complex plane that are very close to points on the physical interval, then it will be necessary to use a large number of Fourier coefficients in any event. A small extension zone is then best; for a "wiggly" $f(x)$, the smoothness of the bell function does not control the rate of convergence of $\mathcal{T}f$ unless the extension zone is very narrow.

In the analysis that follows, we always take it as a given that the boundaries of the physical zone and of the extended interval, χ and Θ, are fixed. A comprehensive formula for choosing the size of the extension zone for different classes of $f(x)$ is an open problem beyond the scope of this article.

2.3. Imbricate Series Theory

As explained in [3, 4], the Poisson summation theorem implies that the Fourier coefficients of an imbricate series are given by the Fourier transform of the pattern function where the transform is defined by

$$g(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} G(k) \exp(-ikx)\, dk. \tag{19}$$

For the naive, nonoverlapped bell/imbricate method and also for the overlapped scheme, Poisson summation gives

$$f(x) = \sum_{m=-\infty}^{\infty} G([x - m2\Theta]) = \frac{\pi}{\Theta} \sum_{n=-\infty}^{\infty} g\left(\frac{\pi}{\Theta} n\right) \exp(i\pi nx/\Theta). \tag{20}$$

The crucial point is that the transition region in the pattern function in the nonoverlapped method is half as wide as that in the overlapped method. Because the width of a Fourier transform is inversely proportional to the width of the function being transformed, this implies that the Fourier transform $g(n)$ decreases only half as fast in the nonoverlapped method, which is therefore always inferior to the overlapped scheme.

The alternating method gives

$$\begin{aligned} f(x) &= \sum_{m=-\infty}^{\infty} (-1)^m G_{alternating}([x - m2\Theta]) \\ &= \frac{\pi}{\Theta} \sum_{n=-\infty}^{\infty} g_{alternating}\left(\frac{\pi}{\Theta}[n + 1/2]\right) \exp(i[2n + 1]\pi x/(2\Theta)). \end{aligned} \tag{21}$$

Because $G_{alternating} \equiv G_{overlapped}$, we conclude that the alternating and nonalternating schemes have virtually *identical* rates of convergence.

This is confirmed by Fig. 5, where the curves for Fourier coefficients of the overlapped and alternating methods fall on top of one another. The coefficients for the naive, nonoverlapped method fall off with degree j at half the rate of the other methods, as predicted.

The right panel shows that the extended function, $\tilde{f}$, is very different for each of the three methods on the extension zone, even for the two that have almost identical Fourier coefficient magnitudes. This is yet another reminder that the extension is highly nonunique.

2.4. A Proof That Bell–Imbricate Extensions Have Infinite-Order Convergence

THEOREM 2.1 (Infinite Order Fourier Convergence for Bell–Imbricate Extensions). *Suppose $f(x)$ is analytic everywhere on the interval $x \in [-\Theta, \Theta]$ (for the nonoverlapped extension) and everywhere on the interval $x \in [-2\Theta, 2\Theta]$ for the overlapped and alternating extensions. Let $\mathcal{B}$ denote a bell, not necessarily that defined earlier, which is infinitely differentiable (C^∞) everywhere on $x \in [-\infty, \infty]$ and is identically zero when $|x| \geq \Theta$ (nonoverlapping case) or when $|x| \geq 2\Theta$ (nonoverlapped and alternating cases). Define the extensions by the imbricate series*

$$\tilde{f}_{nonoverlapped} \equiv \sum_{m=-\infty}^{\infty} f(x - 2\Theta m)\mathcal{B}(x - 2\Theta m), \tag{22}$$

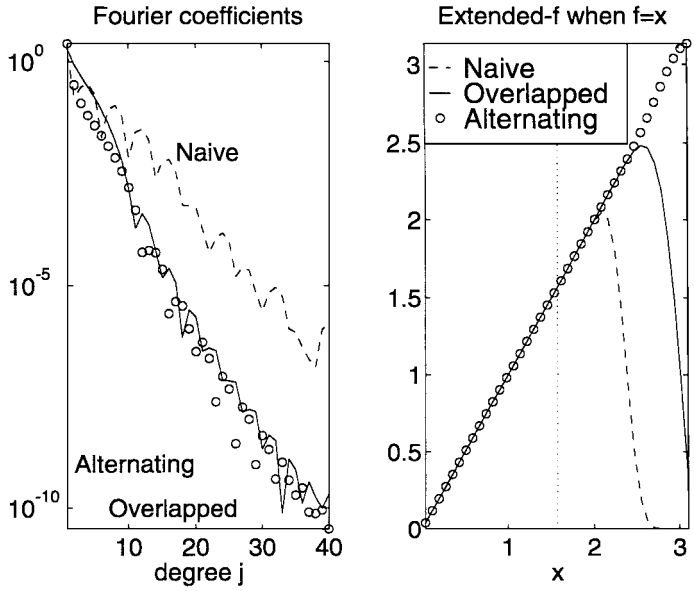

FIG. 5. Comparison of three methods of first-kind extension for $f(x) = x$. (Left) Fourier coefficients where the nonoverlapped method is labeled naive; the overlapped and alternating coefficients are not identical but are inextricably mixed on the graph. (Right) $\tilde{f}$ for each of the three methods. The vertical dotted lines mark $\chi = \pm\pi/2$; the spatial period of the extended function is 2π. Because $\tilde{f}(x) = -\tilde{f}(-x)$ for all three extensions of this $f(x)$, the extended functions are graphed only for $x \geq 0$. The bell width parameter is $L = 4$.

and with a wider bell

$$\tilde{f}_{overlapped} \equiv \sum_{m=-\infty}^{\infty} f(x - 2\Theta m)\mathcal{B}(x - 2\Theta m), \tag{23}$$

and by the alternating-imbricate series

$$\tilde{f}_{alternating} \equiv \sum_{m=-\infty}^{\infty} (-1)^m f(x - 2\Theta m)\mathcal{B}(x - 2\Theta m). \tag{24}$$

Then the Fourier series for all three extensions have Fourier series with "infinite order" convergence [6] *in the sense that coefficients c_n, or, equivalently, the cosine coefficients a_n and sine coefficients b_n, satisfy the bounds*

$$|a_n|, |b_n| \leq \text{constant } n^{-k} \tag{25}$$

for arbitrarily large-order k.

Proof. Because the bell $\mathcal{B}$ is infinitely differentiable for all x and is identically zero for $|x| > \Theta$ or $|x| > 2\Theta$, depending on the case, the product of $\mathcal{B}$ with $f(x)$ is infinitely differentiable for all real x. This allows us to integrate by parts the Fourier coefficient integrals as many times as we please. The coefficients of the complex exponential form of a Fourier series of period 2Θ are

$$c_n \equiv \frac{1}{2\Theta} \int_{-\Theta}^{\Theta} \tilde{f}(x) \exp(-in[\pi/\Theta]x)\, dx, \tag{26}$$

where the cosine and sine coefficients are $a_n = 2\Re(c_n)$ and $b_n = -2\Im c_n$. If we integrate by parts, repeatedly integrating the exponential and differentiating $f(x)$, we obtain without approximation after k steps

$$c_n = \frac{1}{2\Theta}\sum_{j=0}^{k-1}(-1)^{j+n}\left(\frac{i}{n}\right)^{j+1}\left(\frac{\Theta}{\pi}\right)^{j+1}\left\{\tilde{f}^{(j)}(\Theta) - \tilde{f}^{(j)}(-\Theta)\right\}$$
$$+ \frac{1}{2\Theta}\left(-\frac{i}{n}\right)^{k}\left(\frac{\Theta}{\pi}\right)^{k}\int_{-\Theta}^{\Theta}\tilde{f}^{(k)}(x)\exp(-in[\pi/\Theta]x)\,dx, \tag{27}$$

where $\tilde{f}^{(k)}(x)$ denotes the kth derivative of $\tilde{f}(x)$. Because the bell–imbricate extensions are defined as imbricate or alternating-imbricate series, $\tilde{f}$ is always periodic. The period is 2Θ for the overlapped and nonoverlapped extensions. This implies that all the boundary terms in the summation in Eq. (26) are zero since $\tilde{f}$ and its derivatives are the same at both $x = \Theta$ and $x = -\Theta$ because of the periodicity.

The integral can be bounded by a constant independent of degree n, namely

$$\left|\int_{-\Theta}^{\Theta}\tilde{f}^{(k)}(x)\exp(-in[\pi/\Theta]x)\,dx\right| \le 2\Theta\max_{x\in[-\Theta,\Theta]}\left|\tilde{f}^{(k)}(x)\right|. \tag{28}$$

It follows that Fourier coefficients must be bounded by a constant times n^{-k} after k integrations by parts. Since k is arbitrary, the theorem is proved.

For the alternating series where the period is 4Θ, the same argument applies except that Θ in (26) and (27) is replaced by 2Θ. ■

Note that (i) $f(x)$ need not be periodic and (ii) the bell $\mathcal{B}$ need not equal the particular bell $\mathcal{T}$ defined above, although the theorem applies to this bell as a special case. When the bell is narrowed to create a bell–imbricate extension for a second-kind problem, as described in the next section, the theorem and proof still apply with only trivial modifications.

2.5. Summary of First-Kind Extensions: When $f(x)$ Is Known Outside the Physical Interval

When $f(x)$ is known not only on the physical interval but also in the desired extension region, we have an "extension of the first kind." If $f(x)$ is free of singularities on and near the extended interval, then all four methods described here will yield rapidly convergent Fourier series for the extended function $\tilde{f}$.

Figure 6 compares all four schemes for a particular nonperiodic $f(x)$. As noted earlier, the nonoverlapping extension (thin dashed line) is greatly inferior to the overlapping and alternating bell–imbricate methods and therefore should never be used. The overlapping and alternating extensions are almost indistinguishable; both give Fourier series with an *exponential* rate of convergence, as shown by the almost-linear graphs of their coefficients on this log-linear plot.

The three Hermite interpolant curves show several things. First, the rate of convergence improves with increasing order; the ninth-degree polynomial (thick zigzag solid curve) is a Fourier series with sixth-order convergence. Second, even the ninth-degree polynomial gives coefficients which are much larger for high degree than those of the overlapped and

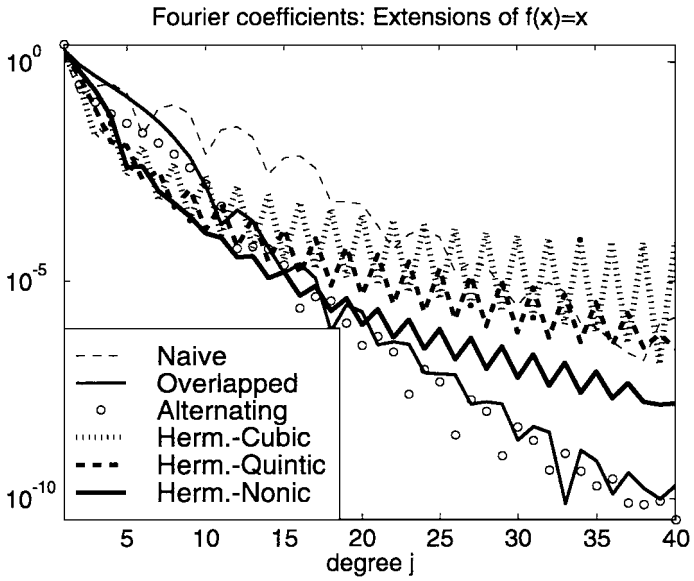

FIG. 6. Comparison of the Fourier coefficients of $\tilde{f}$ for four different methods for first-kind Fourier extensions for $f(x) = x$. The three thick, zigzag curves are computed by the Hermite/two-point Taylor method, but for orders $K = 3, 4, 6$, respectively. $L = 4$ for the three schemes that use $\mathcal{T}$. Note that only the first 40 nonzero coefficients are shown; Fourier terms that are identically zero because of symmetry [i.e., $f(x) = -f(x)\ \forall x$] are omitted.

alternating methods. This is exactly as predicted by theory; Theorem 2.1 proves that the bell–imbricate methods give Fourier coefficients that decay faster than $O(j^{-k})$ for any k, an infinite-order convergence [6]. In contrast, Theorem 5.1 of Section 5 shows that the Hermite approximant of degree K gives Fourier coefficients decaying no faster than $O(j^{-(K+3)/2})$, a finite-order convergence rate.

Thus, although the quintic Hermite polynomial worked well with fourth-order convergence for Garbey and Tromeur–Dervout, the clear winners are two bell–imbricate methods: overlapped and alternating. Our personal preference is for the overlapped method, but the alternating imbricate scheme used by [1, 16, 24, 25] is not at all inferior.

All third-kind-extension schemes can also be applied to first-kind problems, too. We compare the best first-kind and third-kind methods in Section 6; the third-kind method proved costly but accurate.

3. EXTENSION OF THE SECOND KIND: KNOWN SINGULARITIES IN THE EXTENSION INTERVAL

When the function $f(x)$ is singular at a point x_s on the extended interval (the definition of second-kind extension), three strategies for obtaining a nonsingular $\tilde{f}$ are the following:

1. Choosing a different $f(x)$, with its special case of singularity movement (SM).
2. Narrowed bell (NB).
3. Twice-residual singularity subtraction (TRSS).

3.1. Prudent Choice and Singularity Movement

Choosing a different function, nonsingular on both the physical and extended intervals, is an embarrassingly obvious tactic. If the first choice is singular but the only need is for a function that increases monotonically at a certain general rate, or has a narrow peak, or whatever, a function often can be invented that has the required qualitative behavior but is singularity-free on the whole extended interval.

A variation on this theme, which is illustrated in the comparisons below, is to modify the internal structure of the function so as to move the singularity off the extension interval. For example,

$$\hat{f} \equiv \frac{1}{\cos(x)+2} - \frac{x_s}{\mathcal{H}([-x+x_s]/(x_s-\chi), L)x - x_s} + \frac{x_s}{\mathcal{H}([x+x_s]/(x_s-\chi), L)x + x_s} \tag{29}$$

is a symmetric function which would be singular at both $x = \pm x_s$ if the ramps (smoothed step functions) $\mathcal{H}$ were replaced by the constant function. Outside the physical interval, however, the denominator factors of x are weakened rapidly as $|x|$ increases so that the denominator never changes sign and $\hat{f}(x)$ is nonsingular for all real x.

3.2. Narrowed-Bell (NB) Scheme

If one insists on employing an $f(x)$ whose *analytic* extension is singular in the extension zone, a simple and never-failing strategy is to multiply $f(x)$ by a bell function, as in the naive, nonoverlapping first-kind method, but *narrowing* the bell so that $\mathcal{T}$ is *infinitely flat* at the location of the singularity, $x = x_s$:

$$\tilde{f}_{NB} \equiv f(x)\mathcal{T}(x; L, \chi, |x_s|). \tag{30}$$

(Recall that by definition $\mathcal{T}(x; L, \chi, |x_s|)$ decays from one to zero on $x \in [\chi, |x_s|]$.) Thus, this extended function is identically zero at the location of the nearest singularity and for all larger $|x|$.

As noted earlier, the narrower the bell, the slower the rate of convergence of the Fourier series. In any event, the narrowed bell still yields an *exponential* rate of Fourier convergence for the extended function, $\tilde{f}$.

A strength of this strategy is that one does not need to know the precise *type* of singularity (pole, branch point, etc.) nor even the precise *location*. If the point of the nearest singularity is known only approximately, one can set $|x_s|$ equal to the lower bound of the range where the singularity is thought to lie.

3.3. Twice-Residual Singularity Subtraction (TRSS)

Another possibility is to modify the function—but only outside the physical interval—so that the modified function is nonsingular. If the singularity is a simple pole, one may modify $f(x)$ by adding a term that (i) is identically zero on the physical inteval but (ii) cancels the pole on the extension zone.

To employ bells and ramps which are as wide as possible, it is most efficient to subtract a term which is equal to *twice* the residual. This will still cancel the singularity if it is

multiplied by a ramp $\mathcal{H}$ which is sufficiently wide so that it rises only to a value of $1/2$ at the singular point. Assuming $\mathcal{H}$ varies from the endpoint of the physical interval, this gives a ramp which is twice as wide as those employed in the narrowed-bell method: the ramp is not required to rise all the way to one, but only halfway, at $x = x_s$.

An example of such a modified function is

$$\bar{f} \equiv \frac{1}{\cos(x)+2} + \frac{x_s}{|x|-x_s} - \frac{x_s}{|x|-x_s}\left\{1 - 2\mathcal{H}\left(\frac{|x|-x_s}{|x_s|-\chi}, L\right)\right\}. \tag{31}$$

$\bar{f} \equiv f \; \forall x \in [-\chi, \chi]$, but is nonsingular for all real x.

3.4. Comparisons

Figure 7 compares three methods of second-kind extension. The left panel is a log-linear plot; the nearly linear slopes of the Fourier coefficients of $\tilde{f}$ for all three methods show that exponential accuracy is indeed recovered. However, because of its wider ramp, the TRSS strategy is roughly twice as efficient (or, equivalently, gives the same accuracy with only half as many Fourier terms) as the other two methods.

The singularity movement (SM) tactic [Eq. (29)] is no more efficient than the narrowed-bell procedure even though it uses a much wider ramp. We are at a loss to explain this. There is no theory for the asymptotic Fourier coefficients of a function which has been modified by a C^∞ ramp in this way except for some simple theorems that prove the rate of convergence must be exponential.

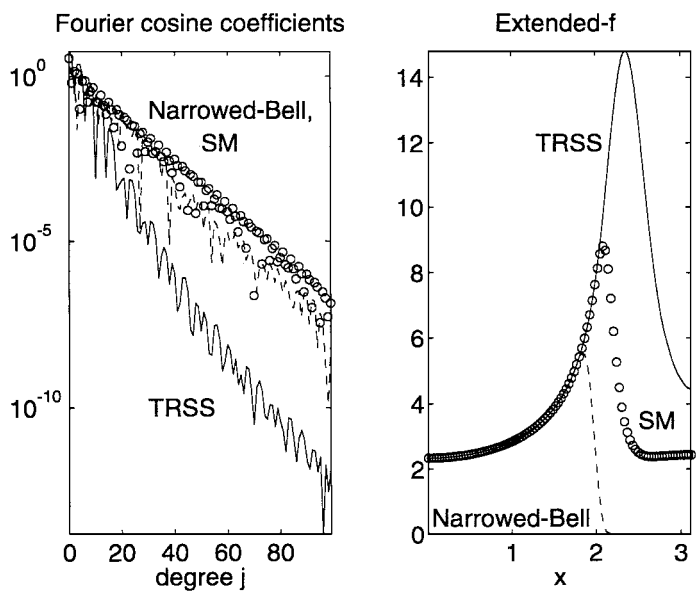

FIG. 7. Comparison of three methods of second-kind extension for $f(x) = 1/(\cos(x)+2) - x_s/(x-x_s) + x_s/(x+x_s)$, where the singularities are at $x = \pm x_s$, $x_s = (3/4)\pi$. (Left) Fourier coefficients. (Right) $\tilde{f}$ for each of the three methods. Because $\tilde{f}(x) = -\tilde{f}(-x)$ for all three extensions of this $f(x)$, the extended functions are graphed only for $x \geq 0$. The bell width parameter is $L = 4$.

4. EXTENSION OF THE THIRD KIND: AVOIDING SINGULARITIES HIDDEN IN THE FOG

When $f(x)$ is not known on the extended interval but only on the physical interval, the extension problem becomes much harder. First, the extension region becomes a "fog zone" where poles and branch points may lurk, as noted earlier. Second, the extension becomes ill conditioned both numerically and intrinsically, as explained below.

We discuss only four methods. One is the Hermite interpolation/two-point Taylor scheme of Garbey and Tromeur-Dervout, which is described in Section 5. The other three use Fourier physical interval collocation (FPIC): each algorithm derives a trigometric polynomial to approximate the extended function $\tilde{f}$ in such a way that the polynomial interpolates $f(x)$ at a set of collocation points which are on the physical interval only. Two of these methods were unsuccessful, so their description is relegated to a pair of brief appendixes (Appendix B and Appendix C). The best method found so far is described next.

4.1. Fourier Physical Interval Collocation—Spectral Coefficients as the Unknowns (FPIC-SU)

The best variety of Fourier physical interval collocation is solving a matrix equation in which the column vector $\vec{a}$, the unknown, stores the coefficients of the usual Fourier series for the interpolant. As in the previous section, assume for simplicity that $f(x)$ and $\tilde{f}$ are both symmetric with respect to $x = 0$. The Fourier approximation is then a truncated cosine series, and all interpolation points will be at nonnegative locations (and confined to the physical interval). The symmetry assumption is no real restriction because an arbitrary function $f(x)$ can be split into its symmetric and antisymmetric parts, $S(x) \equiv (f(x) + f(-x))/2$ and $A(x) \equiv (f(x) - f(-x))/2$, respectively, and the interpolants to S and A computed separately. This reduces the cost by about a factor of 4 from computing a single interpolant to a general unsymmetric function (Chaps. 8 and 10 of [6]).

The Fourier coefficients of the cosine interpolant are the solution of the matrix problem

$$\vec{\vec{M}}\vec{a} = \vec{f}, \tag{32}$$

where

$$M_{ij} = \cos\left([j-1]\frac{\pi}{\Theta}x_i\right), \quad i = 1, 2, \ldots, N_{coll}, \quad j = 1, 2, \ldots, N, \tag{33}$$

$$f_i = f(x_i), \quad j = 1, 2, \ldots N_{coll}, \tag{34}$$

where the collocation points are uniformly distributed over the positive half of the physical interval, $x \in [0, \chi]$,

$$x_i \equiv \frac{(i-1)\chi}{(N_{coll}-1)}, \quad i = 1, 2, \ldots, N_{coll}, \tag{35}$$

where $N_{coll} \geq N$ is the number of collocation points.

If $N_{coll} > N$, then the matrix problem is overdetermined and must be solved in a least-squares sense by singular value decomposition (SVD) or something similar. Even when the matrix $\vec{\vec{M}}$ is square, it unfortunately is ill conditioned and nearly singular.

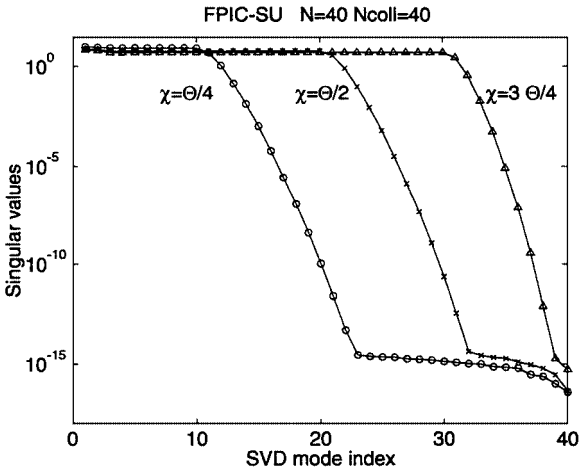

FIG. 8. Singular values for three different values of χ. $N = 40$ and $\Theta = \pi$, but the qualitative pattern of singular values is insensitive to N and Θ.

Because the SVD decomposition is well discussed in standard linear algebra textbooks such as [23], we have banished the SVD formulas used here to an appendix. The crucial point is that by neglecting SVD modes whose singular values are less than some cutoff ϵ, we can filter much of the ill conditioning from the matrix equation [21, 23]. The best choice of ϵ is problem-dependent (other than the general criterion that ϵ must be *small* if the SVD-filtered solution is to be a meaningful approximate solution of the matrix equation).

Figure 8 shows the singular values for the FPIC-SU. Roughly $(\chi/\Theta)N$ singular values are $O(1)$; the corresponding SVD modes are well conditioned. The remaining singular values decrease exponentially with mode number. The plateaus that appear for all three cases for large N are obviously due to roundoff in the SVD factorization. Thus, there is no point in including these modes in an SVD computation of $\vec{a}$ because these modes have roundoff-corrupted singular values that are orders of magnitude too large. Thus, the SVD-filtering makes sense.

In practice, though, it is not only the modes whose singular values are $O(10^{-15})$, and therefore close to machine epsilon, that are suspect. Accuracy is improved by filtering some modes of larger but still tiny singular value.

Another helpful trick, long popular in numerical analysis for ill-conditioned linear systems, is to apply iterative refinement. This is described in Appendix A. It adds only slightly to the cost because the expense of the SVD factorization (typically $O(N^3)$ operations [23]) is large compared to the cost of refinement, which is only $O(2N^2 + 2N_{coll}^2)$ per iteration. The error in solving the linear system is often reduced by an order of magnitude or more, as we illustrate through our examples.

Figure 9 shows how the error for f_{pole} varies with both N and the SVD cutoff. When N is sufficiently large, the maximum pointwise difference on the physical interval between the trigonometric polynomial, whose coefficients $\vec{a}$ are calculated by the SVD method, and the function $f_{pole}(x)$ is as much as 10 billion times smaller than the maximum of $f_{pole}(x)$ on the physical domain. This remarkably small error is achieved in spite of the fact that f_{pole} has two simple poles in the extension zone!

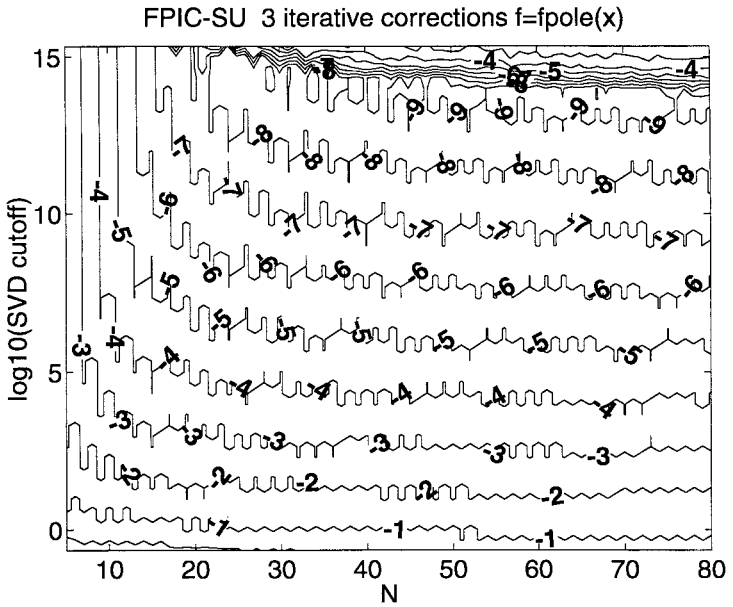

FIG. 9. Logarithm (base 10) of the maximum pointwise error on the *physical* interval only for various N, where N is the degree of the Fourier (cosine) approximation (also the number of collocation points since here $N_{coll} = N$) and the other coordinate is the SVD cutoff. (All modes whose singular values are less than the cutoff are filtered from the computation of the N-term cosine series.) The function being approximated is $f_{pole} = (3/4)\pi/\{(3/4)\pi - x\} + (3/4)\pi/\{(3/4)\pi + x\}$. Minimum errors after zero, one, two, and three iterative corrections are 9.4×10^{-8}, 2.1×10^{-9}, 4.2×10^{-10}, and 8.4×10^{-11} respectively; the maximum of $f_{pole}(x)$ on the physical interval is 5.

The bad news is that the error does not decrease with increasing N to near machine epsilon, which is $O(10^{-16})$, but instead reaches a plateau which is many orders of magnitude larger. Clearly, the SVD filtering has not completely eradicated the ill conditioning of the FPIC scheme, but it has nevertheless made a small error possible.

The densely packed contours next to the upper edge of the diagram show that when the SVD cutoff goes to zero, the error rises steeply. In other words, the SVD filtering is absolutely essential to obtain accurate results. As the cutoff ϵ increases, however the error slowly rises from the minimum.

Figure 10 shows how error varies with the SVD cutoff for a fixed number of collocation points. The broad minimum shows that a cutoff in the range 10^{-12}–10^{-14} gives the best results (not only for the illustrated function but also for all the others that we tried). Thus, although the SVD cutoff, and the need to choose it, is an additional complication in Fourier extension, it is not a very serious complication. We used 10^{-13} in this section and 10^{-12} for the comparisons with the bell-imbricate methods in Section 6.

Figure 11 shows the exact $f(x)$ (dashed line) and the periodic, extended function $\tilde{f}$ (solid line). Instead of rising to a singularity at $x = 3\pi/4$, the extended function smoothly levels off and then decays to a smooth oscillation. Although $\tilde{f}$ has no resemblance to $f(x)$ over most of the extension zone, the maximum pointwise error on the physical interval is only 5.4×10^{-10}!

Figures 12 and 13 show the miminum errors as a function of N and the SVD cutoff, and compare f with $\tilde{f}$, for a different function which is smooth everywhere on the extension

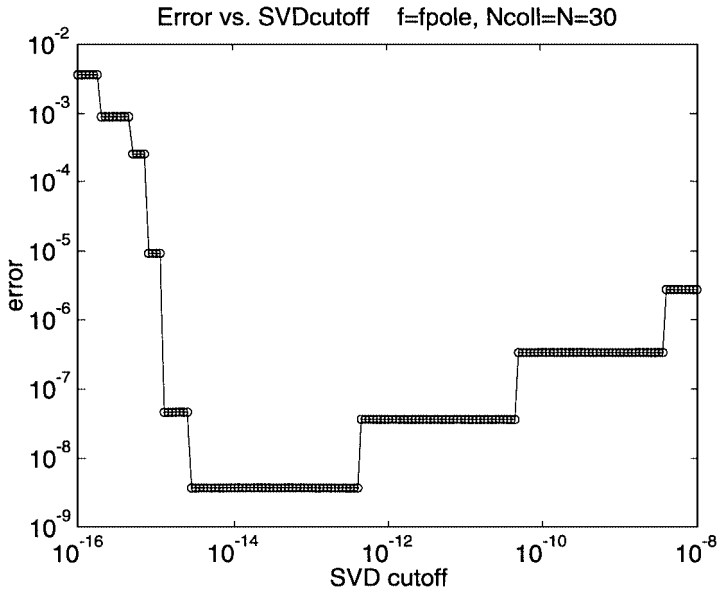

FIG. 10. The error in the L_∞ norm on the physical interval versus the SVD cutoff ϵ for $f(x) = f_{pole}(x)$ (as in the previous figure) and $N = N_{coll} = 50$. The error jumps in discrete units because the number of filtered SVD modes, which is of course always an integer, jumps discontinously as ϵ varies.

interval: $f(x) \equiv \cos(40.5x)$. In contrast to the previous function, here there are rather large errors until the basis for approximating $\cos(40.5x)$ has functions of the same or larger wavenumber, i.e., until $N > 40$. Thereafter, the error decreases rapidly with increasing N—provided the SVD cutoff is in the range 10^{-12}–10^{-14}—to even smaller values than attainable for the previous example.

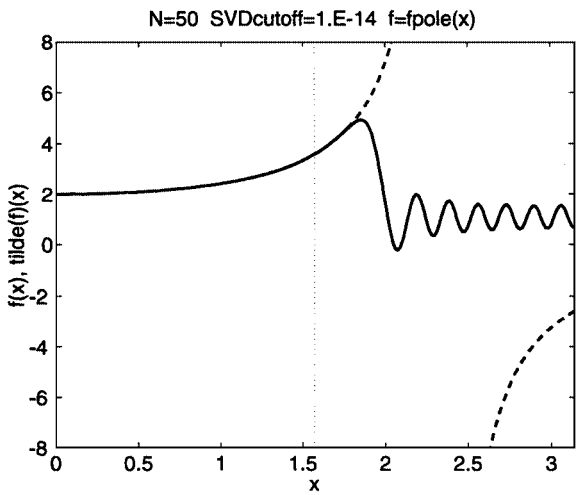

FIG. 11. Fourier cosine interpolation of degree 50 on $x \in [0, \pi]$ with the 50 collocation points restricted to the physical domain $x \in [0, \pi/2]$, where $f = f_{pole}(x) = (3/4)\pi/([3/4]\pi - x) + (3/4)\pi/([3/4]\pi + x)$. The dashed curve is the analytic function $f(x)$, which has a pole at $x = 2.356$. The solid curve is the C^∞ extension, $\tilde{f}$. The approximant that is $\tilde{f}$ includes the basis functions $\{1, \cos(x), \cos(2x), \ldots, \cos(49x)\}$.

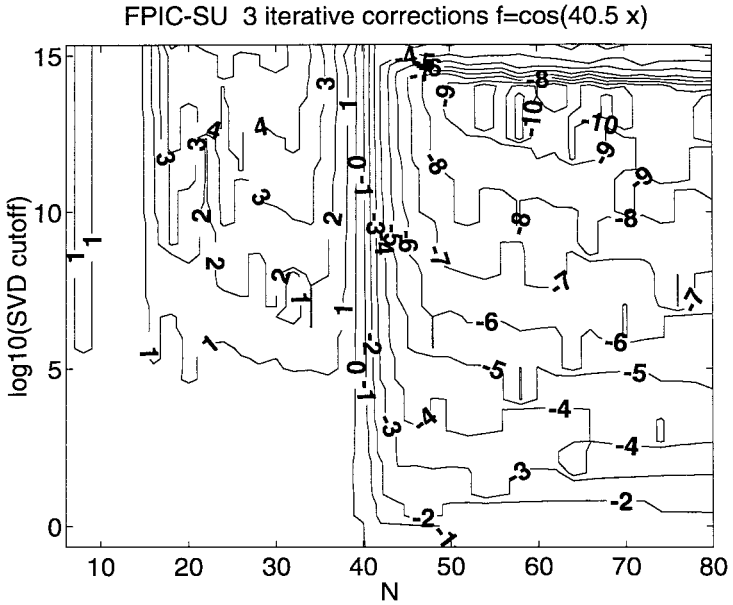

FIG. 12. The function being approximated is $f = \cos(40.5x)$. Logarithm (base 10) of the maximum pointwise error on the *physical* interval only for various N, where N is the degree of the Fourier (cosine) approximation (also the number of collocation points since here $N_{coll} = N$) and the other coordinate is the SVD cutoff. Minimum errors after zero, one, two, and three iterative corrections are 9.7×10^{-9}, 1.9×10^{-10}, 3.5×10^{-12}, and 3.8×10^{-12}, respectively; the maximum of f on the physical interval is 1.

Figure 14 shows how the L_∞ error for $f = f_{pole}(x)$ for a fixed SVD cutoff and various N after zero, one, two, and three iterative corrections. Both panels of the graph show clearly that two or three refinements can reduce the error by several orders of magnitude. This in turn reminds us that the interpolation problem is genuinely ill conditioned,

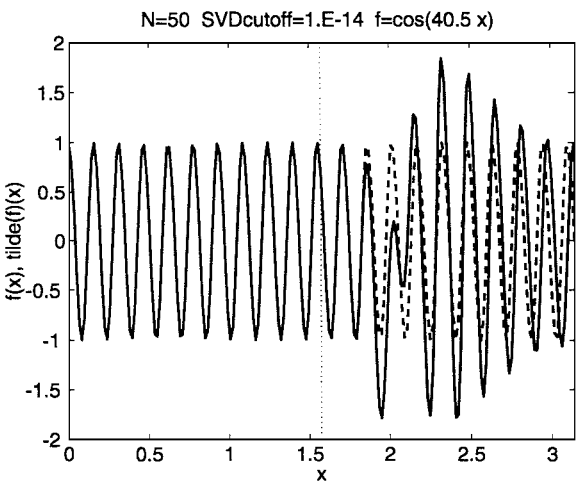

FIG. 13. Fourier cosine interpolation of degree 50 on $x \in [0, \pi]$ with the 50 collocation points restricted to the physical domain $x \in [0, \pi/2]$, where $f = \cos(40.5)$ and the SVD cutoff is 10^{-14}. (Solid line) $\tilde{f}$. (Dashed line) $f(x)$. The vertical dotted line is the boundary between the physical interval and the extension interval. Because of the symmetry of both f and $\tilde{f}$ with respect to $x = 0$, only the positive half of the periodic domain is illustrated.

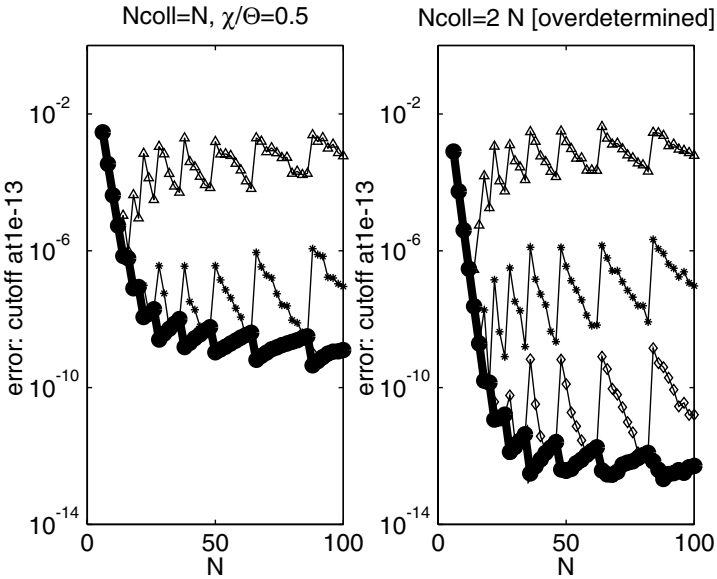

FIG. 14. Errors for the FPIC-SU method applied to $f = f_{pole}(x) = (3/4)\pi/([3/4]\pi - x) + (3/4)\pi/([3/4]\pi + x)$ for various numbers N of cosine functions. The unknowns of the matrix problem are the Fourier cosine coefficients. The left and right panels are identical except that the number of collocation points, N_{coll}, was equal to N in the left panel (interpolation) but *double* N in the right panel, leading to an overdetermined matrix problem for the spectral coefficients with a matrix $\vec{\vec{M}}$ that has twice as many rows as columns. The top curve (triangles) in each panel shows the errors when the matrix problem is solved by SVD factorization-with-cutoff with no iterative refinement; the second-from-the-top curve (asterisks) is the result of one iterative refinement, and the lower two curves are the result of a second and third iterative refinement, respectively. In the left panel, there is no improvement after the second iteration, so the errors after two and three refinements are indistinguishable. For the overdetermined system, however, the third refinement (thick curve with disks) is a significant improvement over the errors after two refinements (diamonds). The SVD cutoff is 10^{-13}.

but with the use of the SVD factorization-with-cutoff and iterative refinement, it is still possible to obtain as many as 10 decimal places of accuracy over the whole physical domain.

Another strategy for making an ill-conditioned problem better is to collocate with more points than spectral coefficients. The matrix $\vec{\vec{M}}$ is then rectangular and the matrix problem is overdetermined. The right panel of Figure 14, compared with the left, shows that dramatic increases in accuracy may also be generated by what we call "overcollocation." Later, we offer a detailed analysis of why overcollocation helps.

Figure 15 is the same as the previous graph except that the singular-on-the-extension-interval function has been replaced by $f(x) = \cos(40.5x)$. Again, overcollocation adds another four decimal places of accuracy, and two iterative refinements—three with overcollocation—reduce error dramatically. For $N > 60$ with $N_{coll} = 2N$, the effect of three refinements is to diminish the error by a factor of 10^9!

4.2. Nonmatrix Ill Conditioning or Why Overcollocation Is Good

Much of the roundoff error in Fourier extensions of the third kind arises from matrix operations, but there is a nonmatrix sort of ill conditioning which is intrinsic to Fourier

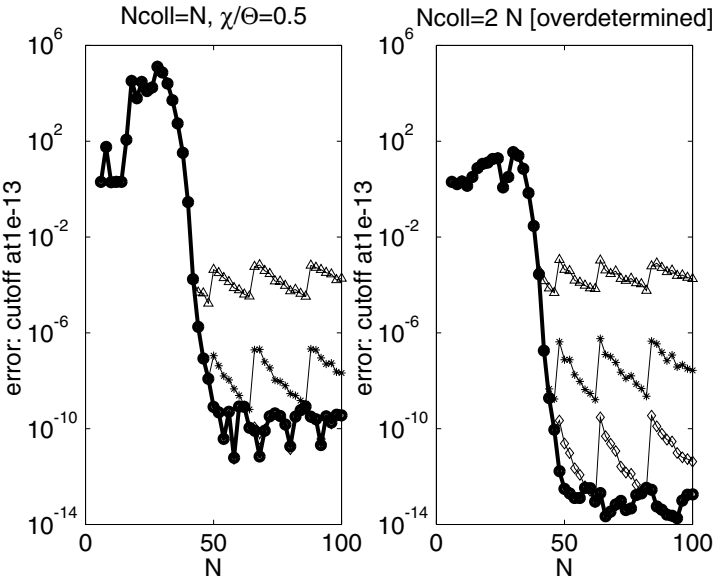

FIG. 15. Same as the previous figure except that $f = \cos(40.5x)$. $\Theta = 2\chi = \pi$ so that the extension zone is equal in size to the physical interval.

interpolation when the collocation points are confined to only part of the periodicity interval.

The fundamental problem is the existence of a class of functions which are *small* on the collocation grid but *large* in the gaps between the interpolation points. We demonstrate the existence of these in two ways. First, we explicitly construct such a function. Second, we solve an optimization problem which generates the worst case: the function which has the largest norm relative to its magnitude on the grid.

The explicit small-on-the-grid/large-in-the-gaps example is

$$\Omega(x; m, \alpha, N) = \{\exp(-\alpha^2(x - [7/8]\pi)^2) + \exp(-\alpha^2(x - [9/8]\pi)^2)\} \times \prod_{k=1}^{m} \{\cos(x) - \cos(x_{N+1-k})\}, \tag{36}$$

where $m \ll N$ and the x_{N+1-k} are the rightmost collocation points on the physical interval. A typical example of this three-parameter family is illustrated in Fig. 16. Because the product factor in Ω has roots at the m points on the collocation grid where it would otherwise be largest, the maximum value of Ω on the collocation grid is very small. In between the roots, however, Ω rises to rather large values so that its maximum on the collocation grid is at the right edge of the physical interval. The function then rises still further to achieve its global maximum on the extension zone.

It follows that a tiny alteration in the values of $f(x)$ on the grid—the only input to an interpolation scheme—may produce large changes in the interpolant on the physical interval.

The insert in Fig. 16 shows that the Fourier coefficients of $\Omega(x)$ converge very rapidly; indeed, the rate is faster than geometric because Ω is an entire function. Consequently, filtering high Fourier components will not purge a perturbation like Ω.

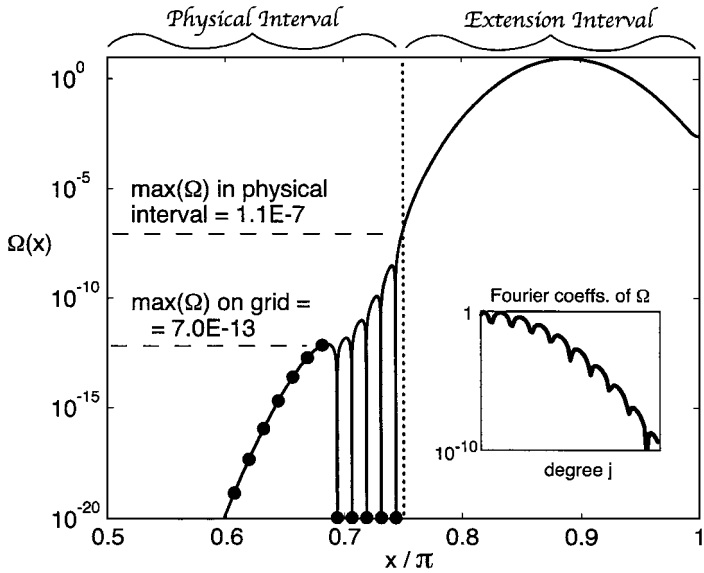

FIG. 16. $\Omega(x)$. The vertical dotted line is the boundary between the physical region and extension zone [at $x = \chi = (3/4)\pi$]. The black disks show the value of $\Omega(x)$ on the collocation grid on the physical interval only. The thick curve shows the value of Ω at all points including the extension zone. The downward spikes point to the roots at five of the collocation points; the zeros have been displayed as disks at $1.E-20$ since their true location at zero is infinitely off the axis on the logarithmic scale. The largest value of Ω on the grid points on the physical interval is only 7.0×10^{-13}, as marked by the lower horizontal dashed line. However, $\Omega(x)$ reaches much larger values *between* the grid points, peaking in magnitude on the physical interval at 1.1×10^{-7}, as marked by the upper horizontal line. Thus, $\Omega(x)$ is more than 100,000 times larger on the physical interval than it is on the collocation points. Note that half of the positive part of the extended interval is shown; $\Omega(x)$ decays rapidly as $|x| \to 0$ for $|x| < \pi/2$, the part of the interval which is not displayed. Because $\Omega(x) = \Omega(-x)$, only positive x are illustrated. $N = 20$ in the extension region with 60 collocation points on $x \in [0, (3/4)\pi]$; the other parameters of Ω are $m = 5$, $\alpha = 8$. The inset graph in the lower right shows the first 80 cosine coefficients of $\Omega(x)$ on a log-linear plot versus degree j.

However, if the number of collocation points N_{coll} is larger than N, the number of Fourier coefficients, then the scheme will detect the big-amplitude oscillations of a function like Ω. Indeed, if $N_{coll} = 2N$, as in our examples, all the local maxima of Ω are at or very near some of the added collocation points. The result is that the minimum error drops from that of pure interpolation (i.e., $N_{coll} = N$) by a factor of 100,000 or more. Overcollocation is an effective way to eliminate perturbations such as our exemplary function, Ω, and thereby obtain a more accurate approximation to $f(x)$ over the physical interval.

4.3. Optimization: The Most Dangerous Perturbation

To systematically calculate even nastier examples of such small-on-the-grid functions, let the function $\omega(x)$ have the expansion

$$\omega(x) \equiv \sum_j \mu_j \phi_j(x), \tag{37}$$

where the $\phi_j(x)$ are the basis functions and μ_j are the spectral coefficients. Let brackets $\langle , \rangle$ denote the usual integral inner product over the physical interval (which depends obviously on the behavior of the functions between the grid points) and let parentheses denote the discrete inner product which depends on function values on the collocation grid only:

$$\langle f(x), g(x)\rangle \equiv \int_{-\chi}^{\chi} f(x)g(x)\,dx; \quad (f(x), g(x)) \equiv \frac{\pi}{N}\sum_k f(x_k)g(x_k). \tag{38}$$

The optimization problem is to maximize the continuous norm on the physical interval while constraining the discrete norm to be equal to one. Using a Lagrange multiplier λ to enforce the constraint, this optimization problem is equivalent to maximizing the "cost function"

$$J \equiv \langle \omega, \omega\rangle + \lambda\{1 - (\omega, \omega)\}. \tag{39}$$

The gradient of the cost function with respect to the spectral coefficients then gives the generalized linear eigenproblem

$$\vec{\vec{G}}\vec{\mu} = \lambda\vec{\vec{K}}\vec{\mu}, \tag{40}$$

where $\vec{\mu}$ is a vector containing the spectral coefficients μ_j and

$$G_{ij} \equiv \langle \phi_i, \phi_j\rangle, \quad K_{ij} \equiv (\phi_i, \phi_j). \tag{41}$$

The Lagrange multiplier λ is the eigenvalue. To interpret it, multiply both sides of the eigenproblem by $\vec{\mu}^T$. By substituting the spectral series for $\omega(x)$ and recalling the definitions of the matrix elements G_{ij} and K_{ij}, one finds that $\mu^T\vec{\vec{G}}\vec{\mu} = \langle \omega, \omega\rangle$ and $\mu^T\vec{\vec{K}}\vec{\mu} = (\omega, \omega)$. Recalling that (ω, ω) is constrained to equal one, we find that at the maximum,

$$\langle \omega, \omega\rangle = \lambda. \tag{42}$$

Thus, the square root of the largest eigenvalue gives the L_2 norm of that function $\omega(x)$ which has the largest norm (on the physical interval only) of all functions which have unit discrete norm. The other eigenmodes are less interesting; the second largest eigenvalue describes that mode which has the biggest norm of any function that is orthogonal to the nastiest mode, and so forth.

Figure 17 shows how λ varies with χ/Θ. In the limit that $\chi = \Theta$, which is the classical Fourier interpolation without extension, all the eigenvalues are one, and the interpolation is very well conditioned. However, as χ/Θ decreases, the largest eigenvalue λ rises exponentially fast even for the rather small N illustrated.

Figure 18 compares $\omega(x)$ (solid, unlabeled line) with the values of $\omega(x)$ on the collocation grid (circles). The maximum of ω on the physical interval is more than 2600 times larger than the largest $\omega(x_i)$—and N is only 12. As N increases, the ill conditioning increases exponentially fast for a given χ/Θ.

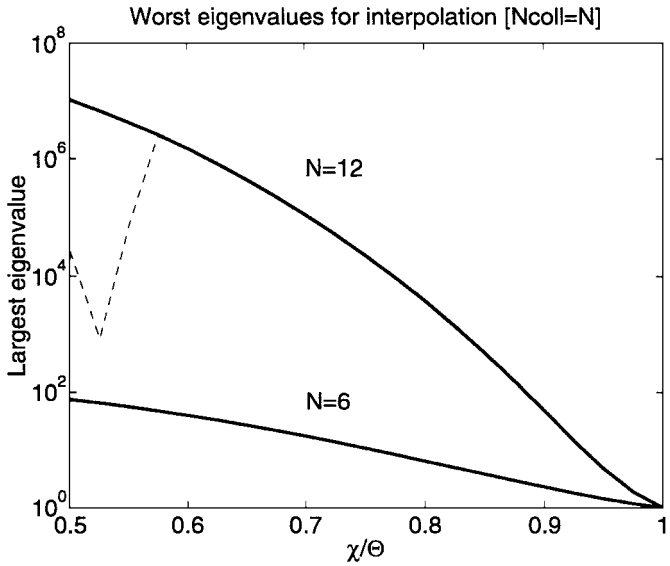

FIG. 17. The largest eigenvalue of the optimization problem to determine that function which has the largest physical interval norm of all functions which have unit discrete norm on the interpolation grid. Note that the upper solid curve was computed using 30-digit accuracy in Maple; Matlab generated the dashed curve, which is inaccurate for small χ/Θ; the optimization eigenproblem is itself very ill conditioned even for small N. ($N_{coll} = N$.)

Artifices like iterative refinement and SVD factorization can reduce roundoff error. However, these are impotent against the nonmatrix, intrinsic ill conditioning shown here. However, because these nasty perturbation functions are largest between the grid points, one can reduce their effect, and obtain much better accuracy, by increasing N_{coll} while the number of Fourier coefficients N is kept fixed.

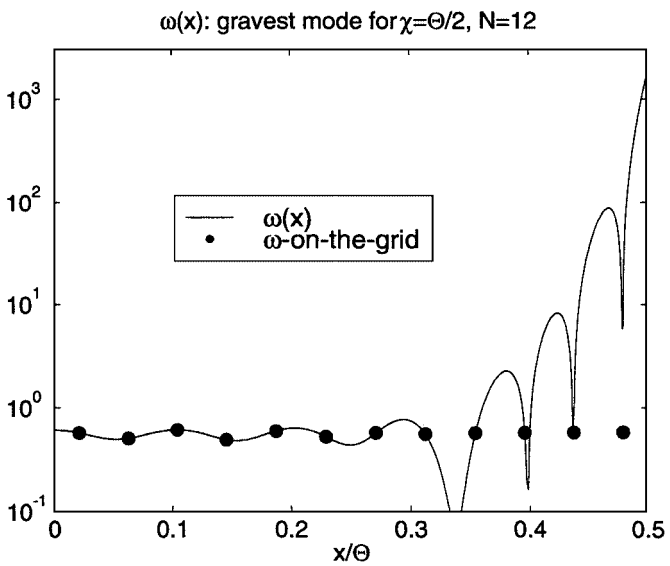

FIG. 18. The eigenmode $\omega(x)$ of largest eigenvalue for $N = 12$ interpolation points with $\chi/\Theta = 1/2$. The solid curve shows ω; the dots show the values of ω at the interpolation points. $\max(\omega) = 1645$, but the maximum of ω *restricted to the interpolation points* is only 0.61.

5. TWO-POINT TAYLOR SERIES/HERMITE INTERPOLATION

5.1. Description

Garbey and Tromeur-Dervout have proposed a radically different remedy, which is to define $\tilde{f}$ on the extension region by that fifth-degree polynomial $\mathcal{P}$ which agrees exactly with $f(x)$ and its first two derivatives at the ends of the extension interval [12, 13]. When a polynomial interpolates one or more derivatives as well as the function itself, the polynomial is usually called a "Hermite interpolant." The extended function is defined by

$$\tilde{f}_{Hermite} \equiv \begin{cases} f(x), & x \in [-\chi, \chi], \\ \mathcal{P}(x), & x \in [\chi, -\chi]. \end{cases} \tag{43}$$

The limits of the second interval, which run from positive to negative, may seem like a misprint, but a periodic interval is effectively a circle with both extension zones joined into a single subinterval that is spanned by the polynomial. The upper limit can alternatively (and more conventionally) be specified as $x = 2\Theta - \chi$ since this is the same point, on a cyclic interval of period 2Θ, as $x = -\chi$. Note that $\mathcal{P}$ is an *ordinary* polynomial in powers of x, as opposed to a trigonometric polynomial.

A disadvantage of the Garbey/Tromeur-Dervout procedure is that the composite function generally has a discontinous third derivative at the breakpoints, $x = \pm\chi$, between the physical and extension intervals. Because of this, the Fourier series of $\tilde{f}_{Hermite}$ has coefficients with a fourth-order rate of convergence; that is, the jth coefficient a_j decreases asympotically as $O(j^{-4})$. However, their method is very simple and parameter-free (except for the width of the extension). For many applications, a fourth-order rate of convergence in the Fourier series will not compromise the overall effectiveness of the computation.

Still, it is useful to ask the question: What happens if the method is generalized to interpolate $f(x)$ and its derivatives up to and including order $(K-2)$ at the breakpoints?

THEOREM 5.1 (Fourier Convergence of Hermite Extension). *If the Hermite interpolant extension is a polynomial $\mathcal{P}$ of degree $2K-3$, the Fourier coefficients of $\tilde{f}_{Hermite}$ decrease as $O(j^{-K})$.*

Proof. Since the $(K-1)$ derivatives are discontinuous at the endpoints, the theorem follows from the theorem on p. 42 of [6]. ■

To distinguish this more precisely from general Hermite interpolation, which uses an arbitrary number of points, one may also refer to $\mathcal{P}$ as a "two-point Taylor series." (One-point Hermite interpolation is the usual Taylor series approximation.)

The two-point Taylor approach of Garbey and Tromeur-Dervout has the advantage that it can in principle be applied to third-kind extensions where the function $f(x)$ is unknown in the extension regions. Only the values of the derivatives of f at endpoints of the physical domain are needed. However, if these derivatives must be calculated by one-sided finite differences, rather than analytically, then the perils of high-order numerical differentiation make the Hermite interpolant much less attractive as a method for Fourier extension of *any* kind. High-order numerical differentiation is numerically unstable in the sense that it becomes more and more inaccurate (rapidly!) as the order of differentiation increases.

The two-point Taylor scheme is in any event bedeviled by convergence problems, described in the next section, which makes numerical differentiation irrelevant.

5.2. Convergence Theory of Two-Point Taylor Series (Hermite Interpolant)

The first theorem [9] describes the construction of the two-point Hermite interpolant of general degree where the points (a, b) of the theorem are $a = \chi$, $b = \chi + 2(\Theta - \chi)$ in our application.

THEOREM 5.2 (Two-Point Hermite Interpolant). *Let a and b denote distinct points. The polynomial of degree $2K - 3$ which interpolates $f(x)$ and its derivatives up to and including order $K - 2$ at the points $x = a$ and $x = b$ is*

$$\mathcal{P}_{2K-3}(x) \equiv (x-a)^{K-1} \sum_{j=0}^{K-2} \mathcal{A}_j (x-b)^j + (x-b)^{K-1} \sum_{j=0}^{K-2} \mathcal{B}_j (x-a)^j, \tag{44}$$

where

$$\mathcal{A}_j = \frac{1}{j!} \frac{d^j}{dx^j} \left\{ \frac{f(x)}{(x-a)^{K-1}} \right\} \Bigg|_{x=b}; \quad \mathcal{B}_j = \frac{1}{j!} \frac{d^j}{dx^j} \left\{ \frac{f(x)}{(x-b)^{K-1}} \right\} \Bigg|_{x=a}, \tag{45}$$

and where $\mathcal{P}_{2K-3}$ solves the interpolation problem

$$\mathcal{P}_{2K-3}(a) = f(a), \quad \frac{d}{dx}\mathcal{P}_{2K-3}(a) = \frac{df}{dx}(a) \ldots, \quad \frac{d^{K-2}}{dx^{K-2}}\mathcal{P}_{2K-3}(a) = \frac{d^{K-2}f}{dx^{K-2}}(a), \tag{46}$$

$$\mathcal{P}_{2K-3}(b) = f(b), \quad \frac{d}{dx}\mathcal{P}_{2K-3}(b) = \frac{df}{dx}(b) \ldots, \quad \frac{d^{K-2}}{dx^{K-2}}\mathcal{P}_{2K-3}(b) = \frac{d^{K-2}f}{dx^{K-2}}(b). \tag{47}$$

Proof. Example 3 of Section 2.5, p. 37, of [9]. ■

Unfortunately, this interpolation does not necessarily converge to $f(x)$ even when the function is analytic everywhere on the interval $x \in [a, b]$ as expressed by the following.

THEOREM 5.3 (Two-Point Hermite Convergence-on-Interval). *The two-point Hermite interpolation converges everywhere on the real interval between the two points $x \in [a, b]$ if and only if $f(x)$, the function being interpolated, is free of singularities everywhere within that domain in the complex plane which is within the lemniscate of Bernoulli whose two foci are $x = a$ and $x = b$.*

More generally, define the family of "ovals of Cassini" by

$$|(x-a)(x-b)| = \rho^2, \tag{48}$$

where $\rho > 0$ is a constant. Let $\mathcal{L}_\rho$ denote the region in the complex plane which is the interior of the oval with "radius" ρ. If $f(x)$ is analytic everywhere within $\mathcal{L}_{\rho_C}$, but not for any larger ρ, then the interpolant converges to $f(x)$ geometrically (with error at degree K proportional to $(\rho/\rho_C)^K$) everywhere in $\mathcal{L}_{\rho_C}$ but diverges outside this.

In particular, the interpolant converges to $f(x)$ everywhere on $x \in [a, b]$ if and only if $f(x)$ is free of singularities everywhere within $\rho = 1$, the heavy curve in Fig. 19. This limiting curve in the complex x-plane, in which the two ovals just touch at a single point, is the lemniscate of Bernoulli.

Proof. Example 7 on p. 86 plus Theorem 4.4.3 on pp. 90–91 of [9]. ■

The theorem has several amusing consequences. First, if $f(x)$ is singular on the extension zone, as can happen in an extension of the third kind, then we do not want to converge to it; rather, we want to converge to a different (necessarily only C^∞) function which is equal to $f(x)$ on the physical interval, but different on the extension zone. This is the nonapproximation property (NAP) that a useful extension must have, as discussed in Section 1.

Second, even for a first-kind extension where we have the freedom to specify an $f(x)$ that is analytic in the extension zone, the two-point Taylor series may fail to converge because of singularities off the real axis, but within the lemniscate of Bernoulli, shown as the thick curve in Fig. 19. This serious flaw of being wrecked by complex-plane singularities even when $f(x)$ is analytic, bounded, and smooth for real x is a liability shared with ordinary Taylor series, which are Hermite interpolations at a single point.

Third, the values of f and its derivatives at $x = -\chi$ are the values that the Hermite interpolant must match on the right side of the extension interval. If $f(x)$ is free of singularities on the extension interval, then the Hermite interpolant will converge to $f(x)$ everywhere except at the point halfway between the two interpolation points, $x = \pm\Theta$ in our application, where $\tilde{f}$ (or at least a sufficiently high derivative) will have a jump discontinuity if f is *nonperiodic* (Fig. 20). The Hermite interpolant unfortunately lacks the nonapproximation property. The price of this is that the Fourier coefficients of the extended function will converge only as $O(1/j)$ in the limit that the degree of the Hermite polynomial tends to infinity.

Figure 21 shows Hermite interpolants of different orders for a different $f(x)$ that is singular in the extension zone. The interpolant is clearly diverging and failing to have NAP in the vicinity of the pole as well as at $x = \pm\Theta$. However, the ninth-degree polynomial (solid)

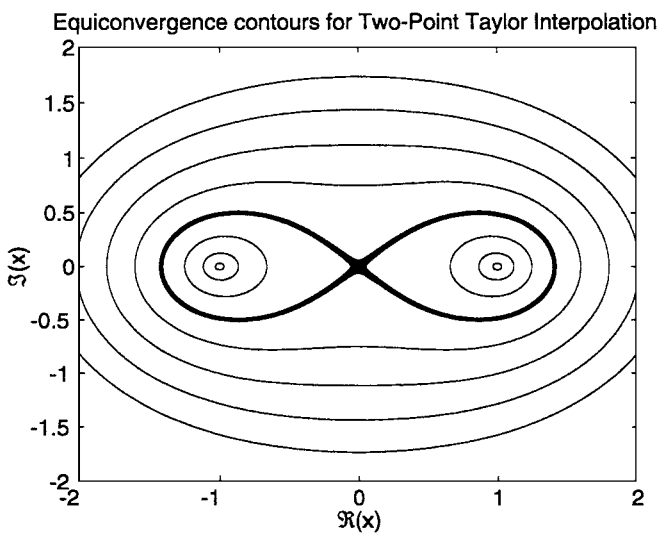

FIG. 19. Equiconvergence surfaces in the complex x-plane for two-point Taylor approximations. The two points where the Hermite interpolant conditions are applied are the centers of the "bulleyes," $x = \pm 1$; when $|a|, |b| \neq 1$, the shape of the contours are identical except for the simple translation and rescaling which takes the interval $[a, b]$ to $[-1, 1]$. The contours are those of the function $\rho \equiv \sqrt{|(x+1)(x-1)|}$, plotted at intervals of $1/4$ from $1/4$ to 2. The heavy contour is $\rho = 1$, which has the shape of the lemniscate of Bernoulli. Only if $f(x)$ is singularity free everywhere within this curve will the two-point Taylor approximation converge everywhere on the real interval $x \in [-1, 1]$.

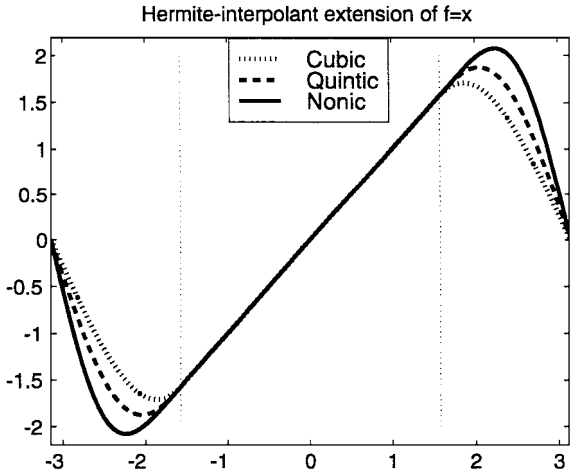

FIG. 20. Hermite interpolant extension of $f(x) = x$ from the physical domain $x \in [-\pi/2, \pi/2]$ (bounded by the vertical dotted lines) to $x \in [-\pi, \pi]$. Extension by means of cubic, quintic, and ninth-degree polynomials in the extension zone is shown. In the limit where the degree goes to infinity, the extended function becomes a straight line with jumps at $x = \pm\pi$ (i.e., $x = \pm\Theta$).

reaches a maximum only about 10 times the maximum of f on the physical interval—not a large enough ratio to cause roundoff error difficulties—and the Fourier series of the extended function will have sixth-order $[O(j^{-6})]$ convergence. Thus, the Hermite interpolant strategy is acceptable for *moderate* order even for functions singular on the extension interval, in spite of the fact that it must inevitably diverge (and fail) as the interpolant order tends to infinity.

Figure 6 above reiterates this point by showing that high-order Hermite interpolation, which requires accurate knowledge of rather high-order derivatives of $f(x)$ at the ends of

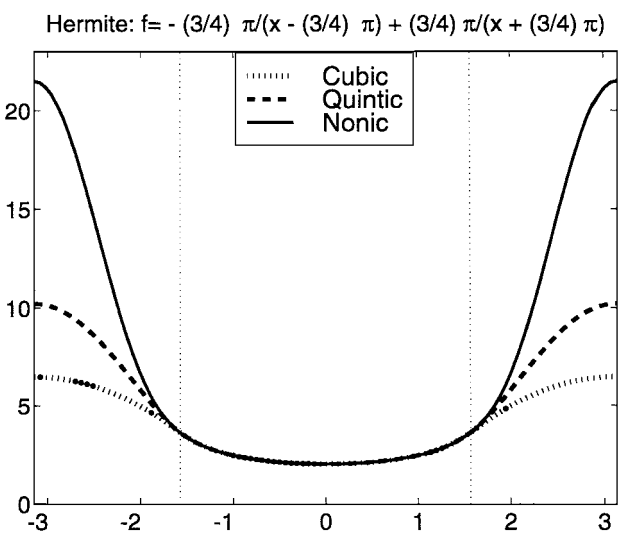

FIG. 21. Same as previous figure except for a different function: $f = f_{pole} = -(3/4)\pi/(x - (3/4)\pi) + (3/4)\pi(x + (3/4)\pi)$.

the physical interval, has an inferior rate of Fourier series convergence compared to the bell-based methods of first-kind extensions.

Nevertheless, for difficult flow problems, such as that attacked by Garbey [12] and Garbey and Tromeur-Dervout [13], the fourth-order accurate/quintic polynomial extension that they employed may be satisfactory. The point is that this moderate order scheme cannot be reliably extended to higher order and thus is not really a "spectral" method.

6. COMPARISON OF FIRST-KIND AND THIRD-KIND METHODS FOR A SUITE OF TEST FUNCTIONS

Any method that works for extensions of the third kind can also be applied to extensions of the first kind merely by ignoring the analytic continuation of $f(x)$ into the extension interval. An obvious question is: Is FPIC-SU better than the bell–imbricate methods for first-kind extensions?

Unfortunately, there is no theorem to show that the FPIC-SU method even converges! We have therefore resorted to head-to-head comparisons between the best first-kind method (the overlapped bell–imbricate method) and the best third-kind scheme, FPIC-SU.

The theory of spectral series [6] shows that Fourier expansions should converge much more rapidly (asymptotically!) for entire functions, such as the cosine and the powers of x than for singular functions. Our tests therefore employed a couple of singular functions in addition to a mix of entire functions. Theory shows that the *type* of singularity—first-order pole versus second-order pole versus logarithm—has only a secondary role in Fourier convergence, modifying the asymptotic coefficients a_j as $j \to \infty$ by perhaps a power of j, but not altering the *exponential* factor of j. It is therefore sufficient to use test functions with second-order poles.

Theory [6] shows that an arbitrary function can always be decomposed into its symmetric and antisymmetric parts, $S(x)$ and $A(x)$, through $S \equiv [f(x) + f(-x)]/2$ and $A \equiv [f(x) - f(-x)]/2$, and each can be separately approximated by a cosine and sine series, respectively. Further, there is no qualitative difference between functions of different symmetries as long as they have singularities in the same locations. Therefore, we lose no generality whatsoever in restricting our examples to symmetric functions in this section.

Figures 22–26 show the errors versus the truncation N of the Fourier series for five representative test functions. The pattern is very consistent.

1. FPIC-SU is roughly equal in accuracy to the bell–imbricate (first kind) method when both iterative refinement and overcollocation are employed.

2. Without overcollocation, the third-kind FPIC-SU method is significantly worse the the bell–imbricate scheme.

3. Without iterative refinement, FPIC-SU is much worse than the bell–imbricate scheme; the error plateaus at $O(10^{-5})$ versus $O(10^{-13})$ for the bell–imbricate algorithm, and the error does not decrease with increasing N.

The conclusion is that the bell–imbricate method, which explicitly uses the analytic continuation of $f(x)$ into the extension zones, is no better in accuracy than the best variant of FPIC-SU, which requires only grid point values of f on the physical interval. (Indeed, for the entire functions, the bell–imbricate is actually a little *worse* than FPIC-SU, especially for x^4, but the first-kind method is equal or a little better for *singular* functions.) However, the graphs show that overcollocation and iterative refinement are both essential to producing

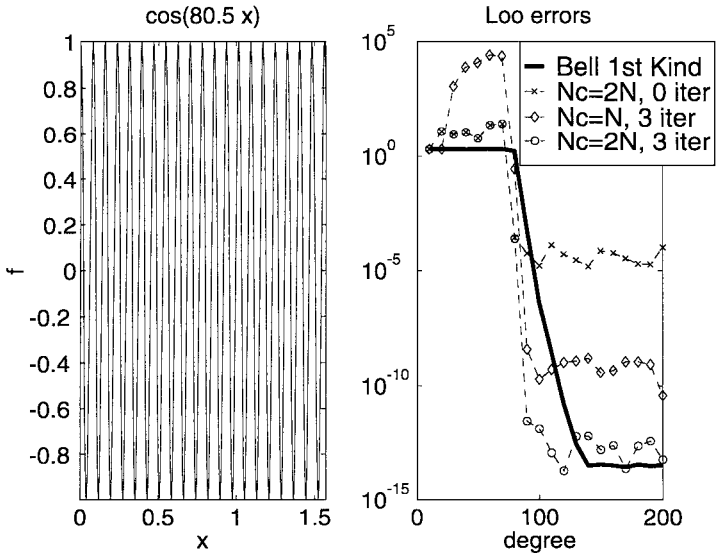

FIG. 22. (Left) a plot of $f(x) = \cos(80.5x)$. (Right) A comparison of errors for this f in the L_∞ norm, plotted versus the truncation N, for the overlapped bell–imbricate method (thick curve) and three variants of FPIC-SU: SVD with overcollocation ($N_{coll} = 2N$) but no iterative corrections (top thin curve with x's), SVD with iterative refinement but $N_{coll} = N$ (middle thin curve with diamonds), and SVD with both iterative refinement and overcollocation (bottom thin curve with circles). The SVD cutoff $\epsilon = 10^{-12} \cdot \chi = \pi/2$, $\Theta = \pi$.

this algorithmic draw between the competing methods. This is unfortunate because SVD with overcollocation requires an expensive $O(N^3)$ factorization of a rectangular matrix with many more rows than columns. The bell–imbricate scheme requires only pointwise operations (the multiplication of f by the bell $\mathcal{T}$) followed by a standard Fourier transform at a cost of $O(N \log_2(N))$ operations. Furthermore, there is a rigorous convergence theory for the bell–imbricate method but not for the third-kind method.

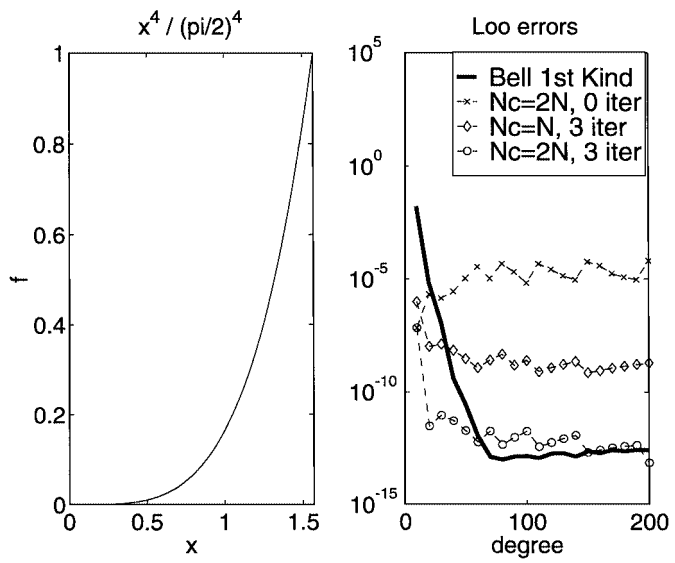

FIG. 23. Same as previous figure but for $f(x) = x^4/(\pi/2)^4$.

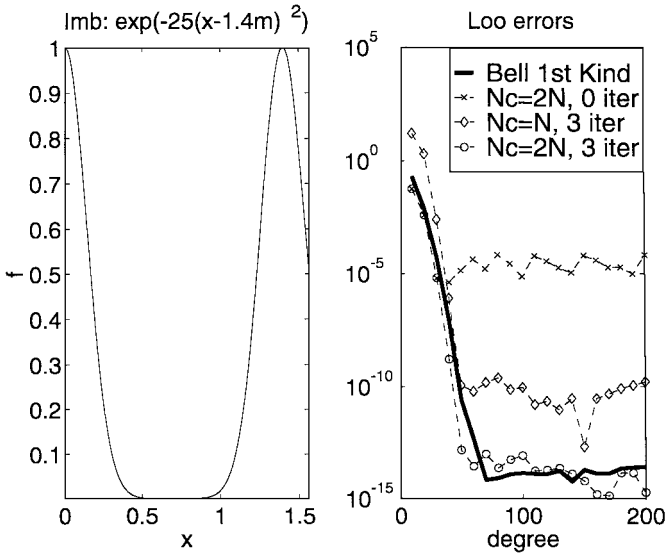

FIG. 24. Same as previous figure but for $f(x) = \sum_{m=-\infty}^{\infty} \exp(-25(x - 1.4m)^2)$, which is an entire function with narrow peaks.

Figure 27 compares the Fourier coefficients as computed by the two algorithms. The spectral coefficients for the bell–imbricate method (left) behave as expected, with rapid decrease with increasing degree; the exception is those for $\cos(80.5x)$, which unsurprisingly rise to a peak at about $j = 80$ and then began to fall monotonically. In contrast, the Fourier coefficients for FPIC-SU yield very flat graphs until near the truncation, with most of the coefficients being of exactly the same magnitude. Although we have not systematically

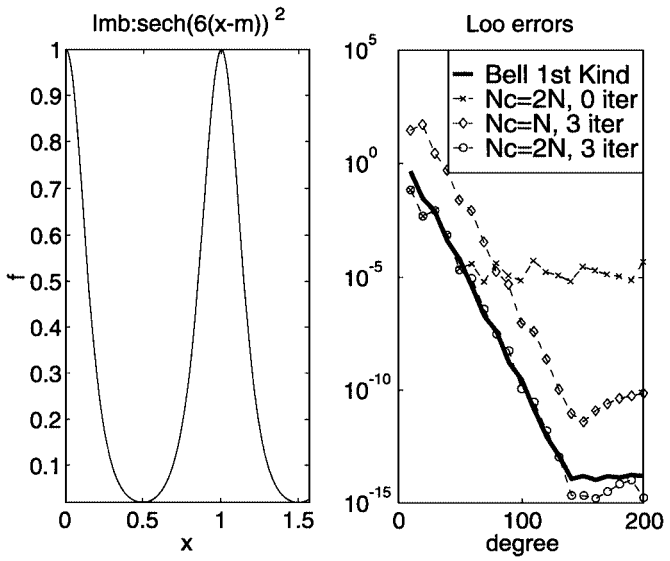

FIG. 25. Same as previous figure but for $f(x) = \sum_{m=-\infty}^{\infty} \mathrm{sech}^2(6(x - m)^2)$; this has simple poles at $x_p = m \pm i(\pi/12)$, where m is an integer.

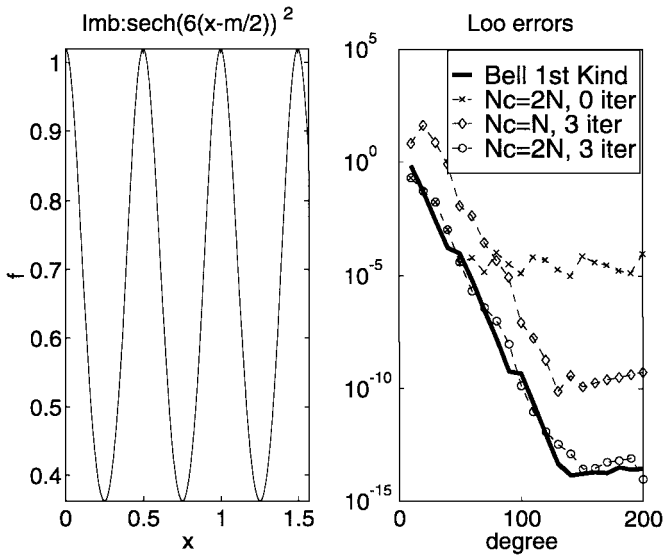

FIG. 26. Same as previous figure but for $f(x) = \sum_{m=-\infty}^{\infty} \mathrm{sech}^2(6(x - m/2)^2)$; this is similar to the previous function except that the singularities and peaks are twice as dense in $\Re(x)$.

investigated derivative errors, the flat Fourier spectra of FPIC-SU suggest that *derivatives* of the approximations generated by this algorithm would be much noiser (and likely less accurate) than the bell–imbricate approximation.

For all these reasons, we advocate the bell–imbricate overlapped method for first-kind problems and FPIC-SU only for third-kind problems. It is gratifying, though, that there is no discernable difference between them in terms of either (i) accuracy for fixed N or (ii) the smallest attainable error as $N \to \infty$, which is close to machine epsilon for both.

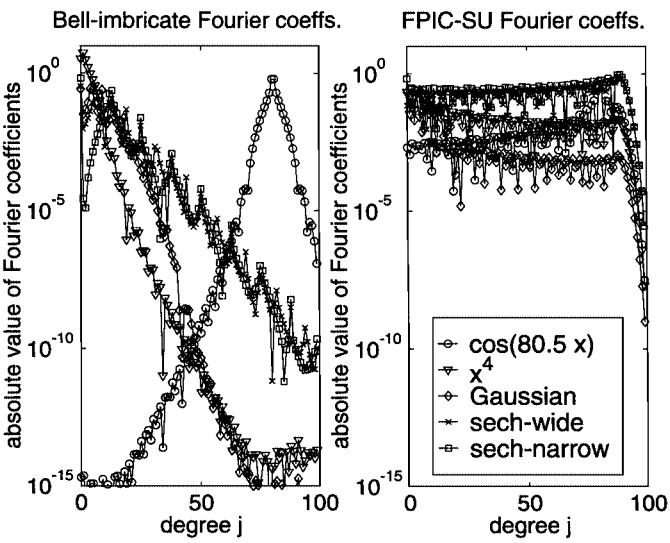

FIG. 27. Absolute values of Fourier coefficients for the bell–imbricate overlapped first-kind method (left) and the third- and first-kind FPIC-SU methods with both overcollocation $N_{coll} = 2N$ and three SVD iterative refinements (right). The legend applies to both panels.

7. GENERALIZATIONS: DOMAIN EMBEDDING (FICTITIOUS DOMAIN) IN MULTIPLE DIMENSIONS AND CHEBYSHEV EXTENSION

7.1. Extension in More Than One Coordinate

A multidimensional generalization of Fourier extension is the embedding of an irregularly shaped domain into a larger, regular shape, followed by spectral expansion on the larger domain. For example, to compute the two-dimensional flow inside a triangle, one could expand the flow as a spectral series on a rectangle that completely encloses the triangle (Fig. 28). Special tricks like conformal mapping and orthogonal curvilinear coefficients are unnecessary; the sole complication of the geometry is that the boundary conditions must be imposed on the triangular surface instead of the walls of the circumscribed rectangle. This strategy is known as the method of "fictitious domain" or "embedded domain."

To apply a multidimensional generalization of Fourier extension, it is necessary to embed the circumscribed rectangle in an even larger rectangle so that a nonperiodic $f(x, y)$ can be deformed into a periodic extended function $\tilde{f}$ without modifying it on the physical domain. For this reason, it is common to use Chebyshev polynomials as the basis in one or both coordinates ("Chebyshev extension"); the embedding domain needs to be just large enough to enclose the physical domain.

Elghaoui and Pasquetti [10, 11], Szumbarski and Floryan [22], and Mason [17] have successfully applied this multidimensional Chebyshev extension, but there is a danger. A solution which is analytic within the irregularly shaped physical domain may have singularities within the larger enclosing domain: the mountain-in-fog problem.

7.2. The Mountain-in-Fog Difficulty

The peril to a pilot in thick fog, radarless, is that of flying blindly into high ground. The peril to the applied mathematician is that of extending a function whose unique analytic continuation has a "mountain," that is, a singularity, within the extension zones, as shown schematically in Fig. 29.

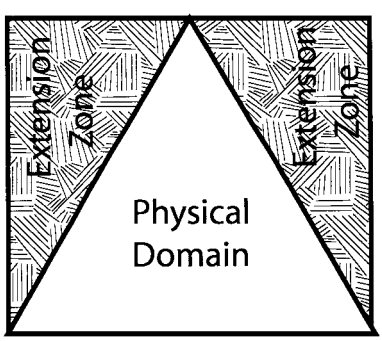

FIG. 28. Schematic of two-dimensional spectral extension. To compute a solution on the interior of the triangle, the unknown is expanded as a two-dimensional tensor product spectral series on the rectangle which contains the triangle. The values of the spectral series on the shaded extension zones have no physical significance. The spectral series must be constrained to satisfy the boundary conditions on the walls of the triangle.

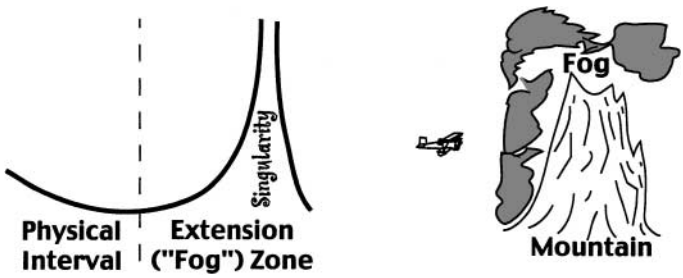

FIG. 29. In the days before radar, many aircraft, such as the little biplane on the right, came to grief by flying into hills or mountain that were hidden by clouds or fog. Fourier extension of the third kind is tricky for the same reason: because $f(x)$ is known only on the physical interval, the extension zones are, metaphorically speaking, regions of dense mathematical fog. They, too, offer opportunities for what aviators call "controlled flight into terrain" in that a function which is smooth and analytic everywhere on the physical region may have poles, branch points, or other singularities lurking in the extension interval, as shown schematically on the left.

DEFINITION 7.1 (Mountain-in-Fog Difficulty). A function f which is well behaved and smooth on the physical domain may have singularities within the larger, extended domain. The peril for spectral extension of the third kind is that by definition one has no knowledge of $f(x)$ outside the physical domain and, thus, has no *a priori* information about such singularities. Nevertheless, poles, branch points, and discontinuities within the extended domain will wreck the convergence of the spectral series of all extensions $\tilde{f}$ which are proportional to the analytic continuation of f or otherwise constructed by schemes that closely approximate the product of f with a bell $\mathcal{T}$.

For example, in one dimension, the function

$$f_{pole}(x) \equiv \frac{1}{x - x_s} \tag{49}$$

has a simple pole at $x = x_s$. If the location x_s of the pole is outside the physical interval, $x \in [-\chi, \chi]$, then $f_{pole}(x)$ is analytic, infinitely differentiable, and very smooth everywhere on the physical interval. However, if the pole is in the extension zone, $|\chi| < |x_s| \leq |\Theta|$, then blindly approximating f using a spectral series must fail.

Chebyshev extension is simpler than Fourier because it is irrelevant whether $\tilde{f}(\Theta) = \tilde{f}(-\Theta)$. However, if there is a singularity on the extension domain, then a successful Chebyshev extension must have the nonapproximation property (NAP) in the vicinity of each such singularity, or the Chebyshev series of the extended function will have an awful rate of convergence, if indeed it converges at all. This in turn implies that the mountain-in-fog problem excludes all classes of extension methods that approximate the analytic continuation of $f(x)$ on the extension interval.

Brazier-Smith [7] illustrated this landmine with computations on the interior of a domain bounded by an ellipsoid which was embedded in a sphere. Using spherical harmonics, it is trivial to solve Laplace's equation on the interior of the sphere and simultaneously fit the boundary conditions on the ellipsoid. The snag is that there may be singularities between the ellipsoid and the circumscribed sphere.

Nevertheless, Elghaoui and Pasquetti and Szumbarski and Floryan proceeded as if Brazier-Smith's paper had never been published and obtained good results. Mason proved that his solutions were nonsingular not only within his physical domain, a curvilinear

polygon, but also within the square which contained the polynomial, thus rigorously justifying his two-dimensional Chebyshev expansions on the enclosing square. Is the mountain-in-fog only a theoretical peril, unlikely in the real world?

7.3. Two-Dimensional Example of Singularities in the Extension Domain

Brazier-Smith gave a resounding No! Mason, years earlier, agreed with him by proving his extension zone was free of "mountains." McIver and Peregrine [19] pointed out the perils of solving water wave problems using analytic continuation from the wavy surface; Bateman *et al*. [2] noted that "so far no problems of this type have been encountered" but modified their procedure as a precaution. Here, we agree with all these by solving a flow between wavy walls, just like Szumbarski and Floryan [22], except our flow is inviscid. With periodic boundary conditions in x, Szumbarski and Floryan expanded their flow as a Fourier series in x with a Chebyshev series in the cross-channel coordinate y over the entire square that encloses one period of the duct. However, the expansion-on-the-square fails for the potential flow in Fig. 30.

The exact flow is two rows of contrarotating potential vortices with the streamfunction

$$\psi(x, y) = \log\{\cosh(y - \pi) - \cos(x)\} - \log\{\cosh(y + \pi) - \cos(x)\}. \tag{50}$$

The walls were artfully chosen to follow two contours of $\psi(x, y)$ so that the usual condition that the flow is parallel to the walls is satisfied. The walls oscillate as far as $y = \pm 2\pi$, so to use an expansion within a rectangle, the embedding domain must be at least that big. Unfortunately, the exact analytic solution (50) shows explicitly that there are branch points at $x = 0$, $y = \pm\pi$ on the interior of the embedding domain.

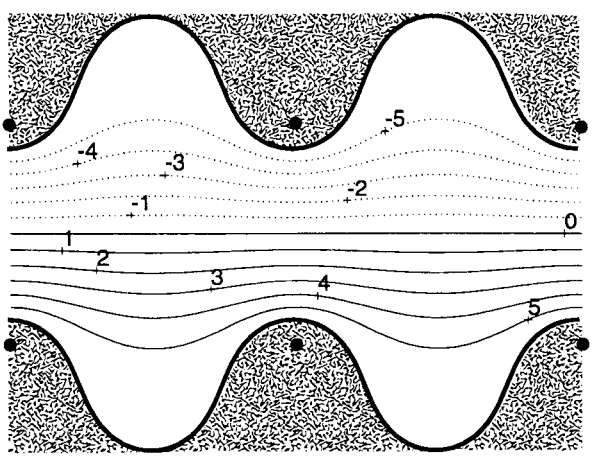

FIG. 30. Two-dimensional inviscid (potential) flow between wavy walls where the walls are shown as the thick oscillating curves. The illustration shows the domain $[-2\pi, 2\pi] \otimes [-2\pi, 2\pi]$; the portion of the square which is outside the fluid domain is shaded. The solid and dashed lines are the streamlines; the flow velocity is everywhere tangent to the streamlines and proceeds from left to right. Although the flow is analytic and smooth everywhere within the flow domain, the streamfunction has logarithmic singularities in the walls, as marked by the black disks. It is not possible to expand $\psi(x, y)$ as a Chebyshev–Fourier series over the illustrated square because of the branch points at $x = 0, \pm 2\pi$, $y = \pm\pi$ within the walls.

Thus, naive embedding must *sometimes fail*. A robust embedding procedure must replace $\psi(x, y)$ by a modified function $\hat{\psi}(x, y)$ which differs from $\psi(x, y)$—and is nonsingular—inside the walls (shaded region in the figure.) A Chebyshev series in the cross-channel direction must be a nonapproximation in the vicinity of all such singularities.

7.4. Strategies for Multidimensional Extension

It is easy to generalize the bell–imbricate method if, through conformal mapping, one can devise an orthogonal boundary-fitting coordinate system. A good multidimensional bell can then be obtained by applying a one-dimensional bell $\mathcal{T}$ that varies with the coordinate normal to the curving boundary.

However, if one has such an orthogonal coordinate system with new coordinates $\xi(x, y)$, $\eta(x, y)$, a simpler alternative is to apply a standard spectral method on the rectangle in the $\xi - \eta$ plane which is the image of the physical domain. Embedded domain/spectral extension methods are then *unnecessary*.

The generalization of FPIC-SU to multiple dimensions is straightforward. However, if N is the number of spectral coefficients in one dimension, then the cost of FPIC-SU is $O(N^6)$ in two dimensions and $O(N^9)$ in three dimensions!

Elghaoui and Pasquetti [10, 11] and Szumbarski and Floryan [22] ignored all these complications and successfully solved problems without conformal mapping or application of a multidimensional generalization of the extension schemes described here. It is clear, however, that they were lucky: By the grace of God, their solutions were free of poles and branch points in the extension zone.

What to do when one is not lucky in two or more dimensions remains an open problem.

8. CLOSING REMARKS

Most engineers and mathematicians take a course in complex variables and never recover from it. The emphasis on analytic functions can give the misleading impression that analytic functions are all there is. Fourier extension requires thinking outside this conceptual box. The whole art of extension is based on creating periodic functions that are *not* analytic. Furthermore, these have singularities which are poles or logarithms not such that the function or a derivative of finite order is infinite at a point but such that the extension $\tilde{f}$ is infinitely differentiable even at the singular points. Such not-analytic-but-C^∞ functions are the cloth from which the extended, periodic functions are tailored.

A good extension must satisfy the conflicting properties of being a good approximation to $f(x)$ in a neighborhood of the endpoints of the physical domain (HIP) and at the same time being very different from $f(x)$ (NAP) in the vicinity of the $x = \pm\Theta$, the boundaries of the extension zone, and also in the neighborhood of any singularities of f in the extension interval. When $f(x)$ is known outside the physical interval, an extension of the first kind, we show that it is easy to construct a robust, exponentially convergent extension by multiplying f by a C^∞ bell (bell–imbricate) method.

When f is not known outside the physical interval, one must use different methods for extensions of the third kind. The Fourier physical interval collocation with spectral unknowns (FPICS-SU) proved to be very reliable and accurate. Indeed, it performed as well or better as the bell–imbricate first-kind method in head-to-head comparisons. Nonetheless, there is a price for ignorance (in this case, ignorance or disregard of the analytic continuation

of $f(x)$ onto the extension zone). FPIC-SU requires singular value decomposition (SVD) with iterative refinement and the use of more collocation points than unknowns. If N is the number of Fourier coefficients to be computed, the bell–imbricate method is an $O(N \log_2(N))$ algorithm whereas FPIC-SU is $O(N^3)$ with a large proportionality constant.

We prove that the bell–imbricate method must yield infinite-order convergence; no similar theorem is known for FPIC-SU. Furthermore, the SVD-based approximation has Fourier coefficients which are very flat (i.e., do not decrease with degree j) until very close to the truncation limit. In contrast, the bell–imbricate coefficients decay geometrically, as is typical of Fourier series. There is clearly some interesting theoretical work undone for the FPIC-SU algorithm.

Still, the empirical performance of both the bell–imbricate and FPIC-SU methods on a suite of test functions is encouraging. Fourier extension is a way to greatly extend the range of Fourier pseudospectral methods.

In Section 7, we showed that multidimensional problems in complex geometry can in principle be solved by extension, too. However, singularities in the unphysical extension zones are a peril dubbed here the mountain-in-fog problem. In principle, such difficulties could be overcome by generalizations of the techniques described here. However, there are many practical difficulties. It is nontrivial, without embracing complications that extension is supposed to avoid, to construct a good, smooth bell function to approximate the characteristic function of an irregular multidimensional domain. Our best third-kind method, FPIC-SU, generalizes to two or more dimensions in a straightforward way, but cost becomes a major issue, and there have been no numerical experiments. Multidimensional Fourier and Chebyshev extensions are hard problems that must be left for the future.

In this work, we have limited ourselves to extending a function $f(x)$ that is known explicitly on the physical interval at least. This can be enough to allow efficient Fourier pseudospectral solutions to some differential equations, as explained in the Introduction. However, another open problem is to devise a general-purpose differential equation solver using Fourier extension where the entire solution, and not just the inhomogeneous term and the particular solution, is expanded as a trigonometric series.

APPENDIX A

Singular Value Decomposition Theorems and Definitions

THEOREM A.1 (SVD Decomposition: Square or Rectangular Matrices). *For any matrix $\vec{\vec{M}} \in R^{m\times n}$, there exist orthogonal matrices $\vec{\vec{U}} \in R^{m\times m}$ and $\vec{\vec{V}} \in R^{n\times n}$ and scalars $\sigma_1, \ldots, \sigma_{n_s}$ such that*

$$\vec{\vec{M}} = \vec{\vec{U}}\,\vec{\vec{S}}\,\vec{\vec{V}}^T, \tag{A.1}$$

where $\vec{\vec{S}} = diag(\sigma_1, \ldots, \sigma_{n_s})$ and where $n_s = \min(m, n)$, that is, $\vec{\vec{S}}$ is a diagonal matrix with the σ_j on the main diagonal. The columns of $\vec{\vec{V}}$ are orthogonal to each other; the same is true of the columns of $\vec{\vec{U}}$. If all the singular values are positive, then $\vec{\vec{A}}$ has an ordinary inverse which may be written in terms of the singular value decomposition as

$$\vec{\vec{M}}^{-1} = \vec{\vec{V}}\,\vec{\vec{S}}^{-1}\vec{\vec{U}}^T. \tag{A.2}$$

DEFINITION A.1 (SVD Inverse/Pseudoinverse). If one or more singular values are less than a user-specified (but *small*) cutoff ϵ, then for any scalar define

$$\sigma^+ \equiv \begin{cases} 1/\sigma & \text{if } \sigma \geq \epsilon \\ 0 & \text{if } \sigma < \epsilon. \end{cases} \tag{A.3}$$

The ϵ-pseudoinverse of a diagonal matrix is defined by $\vec{\vec{S}}^+ \equiv \text{diag}(\sigma_1^+, \ldots, \sigma_n^+)$. The ϵ-pseudoinverse of $\vec{\vec{M}}$ is then

$$\vec{\vec{M}}^+ = \vec{\vec{V}}\,\vec{\vec{S}}^+\vec{\vec{U}}^T. \tag{A.4}$$

THEOREM A.2 (Pseudoinverse Approximate Solution to a Matrix Equation). *The matrix problem*

$$\vec{\vec{M}}\vec{a} = \vec{f} \tag{A.5}$$

has the ϵ-pseudoinverse approximate solution

$$\vec{a} = \vec{\vec{M}}^+ \vec{f} = \vec{\vec{V}}\,\vec{\vec{S}}^+\vec{\vec{U}}^T \vec{f}. \tag{A.6}$$

When the number of rows m is greater than the number of columns n and the SVD cutoff $\epsilon = 0$ or whenever the number of nonzero singular values is less than n, $\vec{a}$ is a least-squares solution to the matrix equation. That is, it minimizes the L_2 vector norm of the residual $\vec{r} \equiv \vec{\vec{M}}\vec{a} - \vec{f}$.

When $0 < \epsilon \ll 1$, the SVD approximation minimizes (or almost minimizes) the residual given the constraints of (i) roundoff error and (ii) whatever intrinsic ill conditioning is in $\vec{\vec{M}}$. Unfortunately, the best SVD cutoff, ϵ, must be determined experimentally. However, in our application, ϵ is not sensitive to $f(x)$ but is primarily determined by the rectangular matrix $\vec{\vec{M}}$, which depends only on the extension method and on N and N_{coll}, which are the size of the rectangular matrix.

DEFINITION A.2 (Iterative Refinement (of the Solution to a Matrix Problem)). Define the residual of the matrix equation $\vec{\vec{M}}\vec{a} = \vec{f}$ by $\vec{r} \equiv \vec{\vec{M}}\vec{a} - \vec{f}$. If the norm of the residual is large compared to the norm of $\vec{a}$ multiplied by machine epsilon, then the accuracy of the computed solution $\vec{a}$ can often be improved by iterative refinement. Each iteration requires the three steps

$$\vec{r}^{(k)} \equiv \vec{\vec{M}}\vec{a}^{(k)} - \vec{f}, \quad k = 0, 1, 2, \ldots, \tag{A.7}$$

$$\vec{\delta}^{(k)} = \vec{\vec{M}}^+\vec{r}^{(k)}, \tag{A.8}$$

$$\vec{a}^{(k+1)} = \vec{a}^{(k)} - \vec{\delta}^{(k)}, \tag{A.9}$$

where $\vec{a}^{(0)}$ is the pseudoinverse (or Gaussian elimination) approximation.

Iterative refinement is very cheap because the cost of the matrix–vector multiplications are $O(N^2)$ per iteration whereas the cost of the SVD factorization is $O(N^3)$.

Empirically, one to three refinements are usually sufficient in the sense that further corrections rarely reduce the error significantly.

All proofs and background can be found in the well-written text of Trefethen and Bau [23].

APPENDIX B

Fourier Interpolation with Collocation on Physical Interval Only (FPIC): Flavor One—Cardinal Function Approach (FPIC-CF)

The most direct way to perform interpolation is to use the Lagrange or cardinal basis. That is to say, we form trigonometric polynomials of degree $N-1$ which are zero at all of the interpolation points except for one. The interpolant is then the sum of all the cardinal functions, weighted by the known values of $f(x)$ at the interpolation points.

When the points are evenly distributed over the whole interval, the cardinal basis is well conditioned and efficient. In Fourier extension, alas, the interpolation points are clustered in the physical interval only. This causes disastrous troubles with roundoff error.

For simplicity, we restrict attention to symmetric functions, which allows us to construct the cardinal basis from cosine functions only. The general, unsymmetric case is similar but requires messier notation to index both sines and cosines. For any distribution of interpolation points, the N-point cosine cardinal functions are

$$C_j^{cos}(x;N) \equiv \prod_{k=1,k\neq j}^{N} \frac{\cos(x)-\cos(x_k)}{\cos(x_j)-\cos(x_k)}. \tag{B.1}$$

Unfortunately, Fig. 31 shows that with points distributed uniformly over the physical interval, a cardinal function blows up rapidly on the extension zone. Indeed, even for N as small as 20, the ratio of the largest maximum to the smallest local maximum is greater than 10^{20}. This implies that the oscillations on the physical interval would be lost in the roundoff

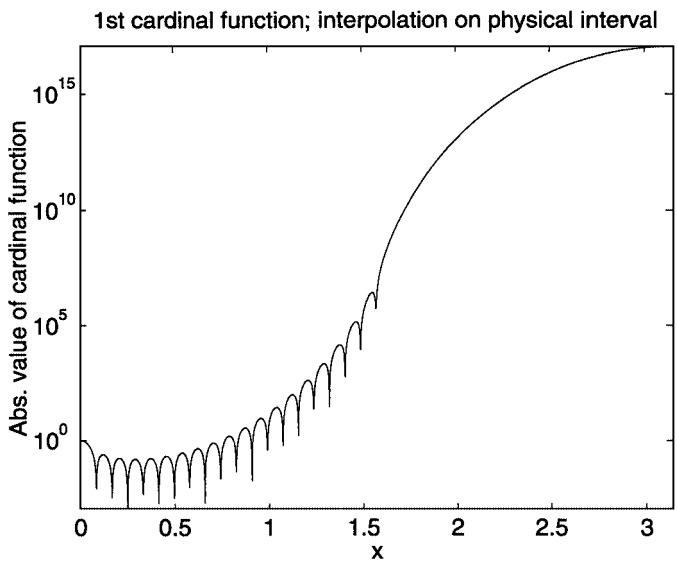

FIG. 31. The cosine cardinal function associated with $j=1$ for interpolation with 20 evenly spaced points on the physical interval only. There are 20 collocation points on $x \in [0, \pi/2]$; the extended domain is $x \in [-\pi, \pi]$.

error relative to the huge magnitude of the cardinal function on the extension zone. The cardinal function basis, so successful when the points fill the whole periodicity interval, is a disaster wrapped in a catastrophe when the points are confined to only part of the interval.

APPENDIX C

FPIC-EXGRU: Grid Point Values as Unknowns

It is possible to apply Fourier physical interval interpolation with the the grid point values on the extension zone as the unknowns. We abbreviate this as the FPIC-EXGRU scheme. When the extension zone is small, this method has the advantage that a small number of grid-point values on the extension zone are sufficient, together with the known values of $f(x)$ on the physical interval, to determine a trigonometric approximation of much larger degree.

The good news is that it is possible to obtain errors similar to those of FPIC-SU when the extension zone and physical interval are of the same size. The bad news is twofold. First, a small extension interval did not reduce the maximum errors on the physical interval but actually increased errors. The reason is that the larger physical interval allows more oscillations between the collocation points.

Overcollocation was successful in reducing such errors when the spectral coefficients were the unknowns. The second piece of bad news is that we have not been sufficiently clever to devise an overcollocation method with grid-point unknowns. We experimented by using an unbalanced grid with a higher density of points on the physical interval than on the uniform grid, but we were not successful in finding an overcollocation scheme that improved on the simple collocation scheme described above.

Because our experiments with the grid-point values as the unknowns produced no improvement and no overcollocation scheme, we shall not discuss FPIC-EXGRU further. However, we cannot yet discount the possibility that more ingenious schemes may yet exploit grid-point values as unknowns.

ACKNOWLEDGMENTS

This work was supported by NSF Grants OCE9986368 and OCE9986368. I thank the two reviewers and Associate Editor Bayliss for very detailed and helpful comments.

REFERENCES

1. A. Averbuch, L. Vozovoi, and M. Israeli, On a fast direct elliptic solver by a modified Fourier method, *Numer. Algorithms* **15**, 287 (1997).
2. W. J. D. Bateman, C. Swan, and P. H. Taylor, On the calculation of the water particle kinematics arising in a directionally spread wavefield, submitted for publication.
3. J. P. Boyd, New directions in solitons and nonlinear periodic waves: Polycnoidal waves, imbricated solitons, weakly non-local solitary waves and numerical boundary value algorithms, in *Advances in Applied Mechanics*, edited by T.-Y. Wu and J. W. Hutchinson (Academic Press, New York, 1989), Vol. 27, p. 1.
4. J. P. Boyd, Construction of Lighthill's unitary functions: The imbricate series of unity, *Appl. Math. Comput.* **86**, 1 (1997).
5. J. P. Boyd, Asymptotic Fourier coefficients for a C^∞ bell (smoothed-"top-hat" function) and the Fourier extension problem, submitted for publication.

6. J. P. Boyd, *Chebyshev and Fourier Spectral Methods* (Dover, New York, 2001), 2nd ed.
7. P. R. Brazier-Smith, On the limitations of spherical harmonics for the solution of Laplace's equation, *J. Comput. Phys.* **54**, 524 (1989).
8. G. F. Carrier, M. Krook, and C. E. Pearson, *Functions of a Complex Variable* (McGraw–Hill, New York, 1966).
9. P. J. Davis, *Interpolation and Approximation* (Dover, New York, 1975).
10. M. Elghaoui and R. Pasquetti, Mixed spectral-boundary element embedding algorithms for the Navier–Stokes equations in the vorticity-stream function formulation, *J. Comput. Phys.* **153**, 82 (1996).
11. M. Elghaoui and R. Pasquetti, A spectral embedding method applied to the advection-diffusion equation, *J. Comput. Phys.* **125**, 464 (1996).
12. M. Garbey, On some applications of the superposition principle with Fourier basis, *SIAM J. Sci. Comput.* **22**, 1087 (2000).
13. M. Garbey and D. Tromeur-Dervout, A new parallel solver for the nonperiodic incompressible Navier–Stokes equations with a Fourier method: Application to frontal polymerization, *J. Comput. Phys.* **145**, 316 (1998).
14. J. E. Haugen and B. Machenhauer, A spectral limited-area model formulation with time-dependent boundary conditions applied to the shallow-water equations, *Mon. Weather Rev.* **121**, 2618 (1993).
15. M. Högberg and D. S. Henningson, Secondary instability of cross-flow vortices in Falkner–Skan-Cooke boundary layers, *J. Fluid Mech.* **364**, 339 (1998).
16. M. Israeli, L. Vozovoi, and A. Averbuch, Spectral multidomain technique with local Fourier basis, *J. Sci. Comput.* **8**, 135 (1993).
17. J. C. Mason, Chebyshev polynomial approximations for the L-shaped membrane eigenvalue problem, *SIAM J. Appl. Math.* **15**, 172 (1967).
18. G. Matviyenko, Optimized local trigonometric bases, *Appl. Comput. Harmonic Anal.* **3**, 301 (1996).
19. P. McIver and D. H. Peregrine, *Motion of a Free Surface and Its Representation by Singularities*, Tech. Rep. (Bristol University, Bristol, 1981).
20. J. Nordström, N. Nordin, and D. Henningson, The fringe region technique and the Fourier method used in the direct numerical simulation of spatially evolving viscous flows, *SIAM J. Sci. Comput.* **20**, 1365 (1999).
21. W. H. Press, B. H. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes: The Art of Scientific Computing* (Cambridge Univ. Press, New York, 1986).
22. J. Szumbarski and J. M. Floryan, A direct spectral method for determination of flows over corrugated boundaries, *J. Comput. Phys.* **153**, 378 (1999).
23. L. N. Trefethen and D. Bau, III, *Numerical Linear Algebra* (Soc. for Industr. and Appl. Math., Philadelphia, 1997).
24. L. Vozovoi, M. Israeli, and A. Averbuch, Multi-domain local Fourier method for PDEs in complex geometries, *J. Comput. Appl. Math.* **66**, 543 (1996).
25. L. Vozovoi, A. Weill, and M. Israeli, Spectrally accurate solution of non-periodic differential equations by the Fourier–Gegenbauer method, *SIAM J. Numer. Anal.* **34**, 1451 (1997).