CSC446: Numerical Methods for PDEs

Robert Almgren

March 27, 2002

Throughout this course, we have discussed Fourier modes in the context of stability. Generally, we assume that a PDE or a discrete scheme has a solution of the form $u(x,t) = U(t)\exp(i\xi x)$ where ξ is a real number. Then we look for the time dependence of $U(t)$, and if it grows, we conclude that the method is unstable. An implicit assumption in this approach is that *any* initial data $u_0(x)$ may be written as a combination of these modes, so these special solutions are stable.

Now we turn this point of view into a full numerical method. It is the fourth of our major approaches to constructing a finite dimensional approxmation to a continuous problem. *Spectral methods* (or *pseudo-spectral*, more correctly) are very well suited for problems with smooth solutions, on bounded rectangular domains.

Fourier Transforms

Forward transform

Suppose $u(x)$ is a real or complex function defined on a finite interval $[0, L]$; it will be convenient to think of u as extended periodically outside this interval. The *Fourier coefficients* are the doubly infinite sequence of complex numbers

$$\hat{u}_k \;=\; \frac{1}{L}\int_0^L e^{-2\pi i k x/L} u(x)\,dx, \qquad \text{for } k = 0, \pm 1, \pm 2, \ldots$$

(the "hat" is commonly used to denote the Fourier transform). There are several different conventions for the prefactor (here $1/L$) but this choice has advantages as I hope I will convince you below.

We use the integer k as the index of the mode. The physical wave number is $\xi_k = 2\pi k/L$, so the wavelength is $\lambda_k = 2\pi/|\xi_k| = L/|k|$. Thus the kth mode has k waves in the period interval.

The symmetries of this sequence mirror those of the function. First, if u is real, then $\hat{u}_{-k} = \overline{\hat{u}_k}$, where the overline denotes complex conjugate. Thus in that case, the values of $\hat{u}_k$ with $k < 0$ duplicate the ones with $k > 0$.

Suppose u is *even:* $u(-x) = u(x)$. Since we have extended it outside the interval $[0, L]$ by assuming periodicity, this really means $u(L - x) = u(x)$. Then it is easy to see that $\hat{u}_{-k} = \hat{u}_k$ (the sequence is even), and again the negative indices contain duplicate information. If u is real *and* even, then the coefficients are real and even. If u is smooth when extended periodicially, then evenness requires zero derivative at $x = 0$; this case is relevant for insulating Neumann boundary conditions on a finite interval.

If u is *odd,* so $u(-x) = -u(x)$ or $u(L - x) = -u(x)$, then $\hat{u}_{-k} = -\hat{u}_k$. In particular, $\hat{u}_0 = 0$, since it is the average value of $u(x)$. If u is real as well, then all the Fourier coefficients are pure imaginary. This corresponds to zero Dirichlet conditions at the end of the interval.

Next, we compute some special "transform pairs," that follow from the *orthogonality relations*

$$\frac{1}{L}\int_0^L e^{-2\pi i k x/L} e^{2\pi i \ell x/L} dx \;=\; \frac{1}{L}\int_0^L e^{2\pi i(\ell-k)x/L} dx \;=\; \delta_{k\ell} \;=\; \begin{cases} 1, & \text{if } k = \ell \\ 0, & \text{else.} \end{cases}$$

Thus if $u(x)$ is a single complex exponential,

$$u(x) \;=\; e^{2\pi i \ell x/L} \qquad \Longrightarrow \qquad \hat{u}_k \;=\; \delta_{k\ell}.$$

The infinite sequence $\{\hat{u}_k\}$ is all zeros, with a single value equal to one at position ℓ. (We do not consider the case of non-integer values of ℓ.) If $u(x)$ is a sum of exponentials, then by linearity of the transform, we have

$$u(x) \;=\; a_1 e^{2\pi i \ell_1 x/L} + \cdots + a_n e^{2\pi i \ell_n x/L} \qquad \Longrightarrow \qquad \hat{u}_k \;=\; a_1 \delta_{k\ell_1} + \cdots + a_n \delta_{k\ell_n}$$

The sequence is all zeros, with values $a_1, \ldots, a_n$ in positions $\ell_1, \ldots, \ell_n$.

In particular, using the de Moivre relation $\cos x = (e^{ix} - e^{-ix})/2$, we see

$$u(x) \;=\; \cos\left(\frac{2\pi}{L}(\ell x - \alpha L)\right) \qquad \Longrightarrow \qquad \begin{cases} \hat{u}_\ell \;=\; \dfrac{1}{2}\, e^{-2\pi i \alpha} \\ \hat{u}_{-\ell} \;=\; \dfrac{1}{2}\, e^{2\pi i \alpha} \end{cases}$$

We have introduced α as a *phase shift.* A pure cosine wave has $\alpha = 0$, giving $\hat{u}_{\pm\ell} = 1/2$. and a sine has $\alpha = 1/4$, giving $\hat{u}_{\pm\ell} = \pm 1/2i$; you may check that the symmetry relationships above are respected. In general, a shift in *position* of the wave in physical space corresponds to a *rotation* of the Fourier coefficients.

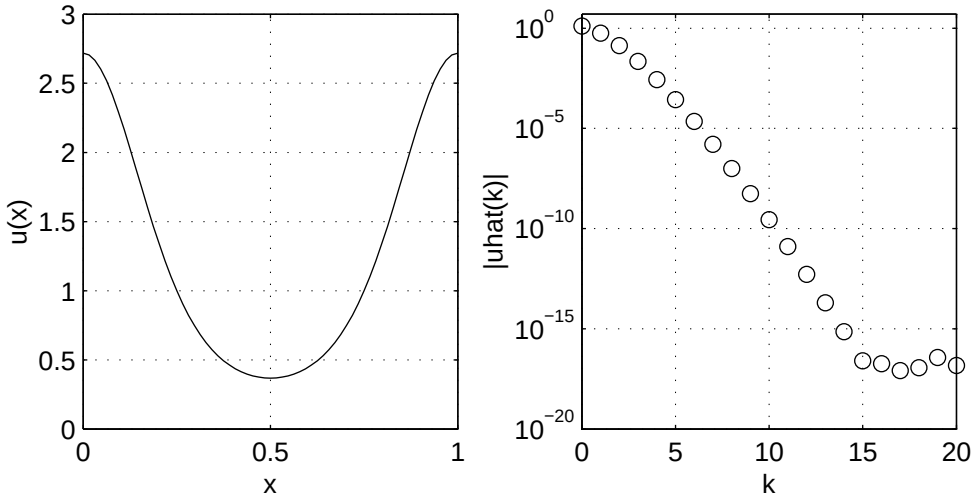

Figure 1: The very smooth function $u(x) = \exp(\cos(2\pi x))$ (left picture) and its numerically computed Fourier coefficients (right). Since u is analytic in the entire complex plane, the coefficients decay faster than any exponential in k. Thus in this log-linear plot, the coefficients go to zero faster than any line, until they bottom out at machine precision around $k = 15$.

We now know what the Fourier transform is, when $u(x)$ is any finite sum of complex exponentials. You may be willing to believe that if $u(x)$ is an *infinite* sum of complex exponentials, then the same relationships hold:

$$u(x) \;=\; \sum_{\ell=-\infty}^{\infty} a_\ell \, e^{2\pi i \ell x/L} \qquad \Longrightarrow \qquad \hat{u}_k \;=\; a_k$$

assuming the numbers $\{a_\ell\}_{\ell=-\infty}^{\infty}$ are such that the infinite sum exists (the a_ℓ decay fast enough as $\ell \to \pm\infty$). But this is very far from saying that *every* function $u(x)$ can be represented as a sum of complex exponentials.

Decay behavior

The decay of the coefficients $\{\hat{u}_k\}$ as $k \to \pm\infty$ is closely related to the smoothness of $u(x)$. The best possible behavior, and the reason that spectral methods work well for smooth problems, comes when u is very smooth. The smoothest that a function of a real variable can be is *real analytic,* meaning that it is the restriction to the real axis of a complex analytic function.

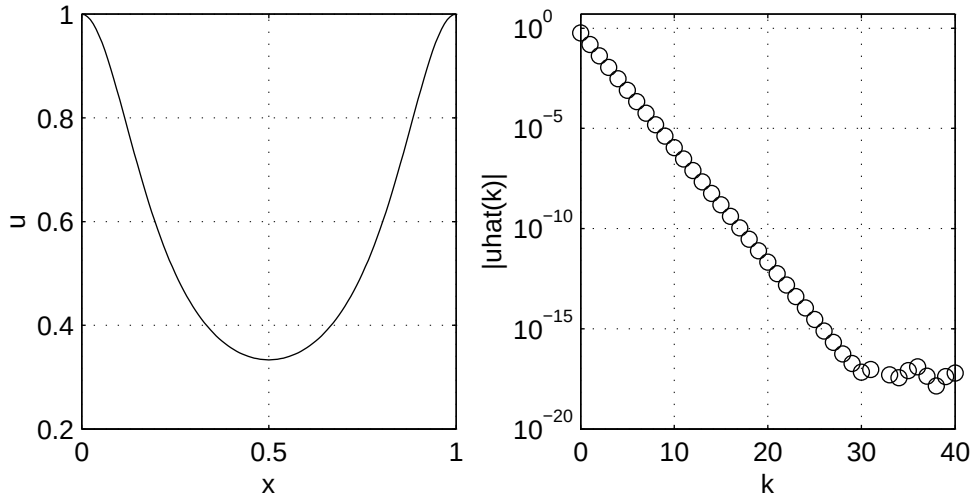

Figure 2: The real analytic function $u(x) = 1/(2 - \cos(2\pi x))$, together with its Fourier coefficients. Although this function is real analytic and looks about the same as the example in Figure 1, it has singularities in the complex plane at a finite distance from the real axis. Hence the magnitude of the Fourier coefficients decays only exponentially, rather than super-exponentially, and we need $k = 30$ to get down to machine precision.

Theorem: Suppose that $u(x)$, periodic in x with period L, can be extended to complex values of x so that it is analytic in a strip $-\rho \le \operatorname{Im} x \le \rho$. Then the magnitudes of the Fourier coeffients $\hat{u}_k$ decay at least exponentially with parameter controlled by ρ :

$$|\hat{u}_k| \;\sim\; \mathcal{O}\big(\, C(\rho)\, e^{-c\rho|k|} \,\big), \qquad |k| \to \infty.$$

where $C(\rho)$ and c do not depend on k; $C(\rho)$ may depend on ρ but not c.

Proof: If $u(x)$ is analytic in $0 \le \operatorname{Im} x \le \rho$, then the path of integration may be shifted upwards to $\operatorname{Im} x = \rho$ without changing the value of the integral, since the contributions along the vertical end segments cancel. Thus

$$\hat{u}_k \;=\; \frac{1}{L}\int_0^L e^{-2\pi i k x/L} u(x)\, dx \;=\; e^{2\pi k\rho/L} \cdot \frac{1}{L}\int_0^L e^{-2\pi i k x/L} u(x + i\rho)\, dx.$$

Now since u is analytic along the line $\operatorname{Im} x = \rho$, it is continuous, and is bounded. Then the value of the integral is bounded independently of k, and this gives the result for $k > 0$, with $c = 2\pi/L$. The case $k < 0$ follows by shifting the path down to $\operatorname{Im} x = -\rho$.

```
function specplot( n )

L = 1;

x = ((1:n)-1)*L/n;
u = exp(cos(2*pi*x/L));

k = [ 0:((n/2)-1) 0 -((n/2)-1):-1 ];
uhat = fft(u)/n;

figure(1)
plot( [ x L ], [ u u(1) ] );
xlabel('x');  ylabel('u(x)');  grid on;

figure(2);
semilogy( 0:((n/2)-1), abs(uhat(1:(n/2))), 'o' );
xlabel('k');  ylabel('|uhat(k)|');  grid on
axis( [ 0 n/2 1e-20 5 ] );
```

Figure 3: The Matlab code that produced the plots in Figures 1, 2, and (with some modifications) 4. `fft` is used to compute the DFT, and we plot only the first half of the resulting sequence; since u is real (and even) the second half contains redundant information.

Alternatively, you can change variables to $z = \exp(2\pi i x/L)$. Then $u(x)$ being analytic in a strip corresponds to $v(z) = u((L/2\pi i)\log z)$ being analytic in an annulus, and the Fourier series for u is the Laurent series for v, whose coefficients decay rapidly.

In the theorem we included the edges of the strip, so there cannot be any singularities right at $\operatorname{Im} x = \pm\rho$. The actual decay of the coefficients has an algebraic contribution (superimposed on the exponential) that can be used to extract information about the nature of the closest singularities of $u(x)$ to the real axis.[1]

Furthermore, if $u(x)$ is not real analytic, then the coefficients have algebraic decay, with an exponent controlled by the differentiability of u. If u is infinitely differentiable but not real analytic, the coefficients decay faster than any power of k, but not exponentially (this is a rare case).

The trigonometric polynomials above are analytic in the entire complex plane, and the Fourier series are exactly zero beyond a finite index; this is

[1] *e.g.,* M. J. Shelley, "A study of singularity formation in vortex-sheet motion by a spectrally accurate vortex method," *J. Fluid Mech.* 244 (1992) 493–526.

certainly rapid decay. For more realistic examples, see Figures 1, 2.

Note that we are referring to the smoothness of u extended to a periodic function. If, for example, u is defined so $u(0) \neq u(L)$, then the extension will be discontinuous and the decay of the coefficients will be slow.

It is the *super-algebraic* decay of the Fourier coefficients for smooth functions that makes spectral methods so powerful. In the language of finite-difference methods with errors $\sim h^p$, these have order $p = \infty$ (when everything works right). This is called *spectral accuracy.*

Inverse tranform

Now let us consider the opposite of the above operation: Suppose we start with a function $u(x)$, compute the infinite sequence of coefficients $\{\hat{u}_k\}_{k=-\infty}^{\infty}$, and we form the infinite sum

$$u(x) \;=\; \sum_{k=-\infty}^{\infty} \hat{u}_k \, e^{2\pi i k x/L}.$$

As described above, if $u(x)$ is smooth, the magnitudes of the $\hat{u}_k$ decay rapidly as $|k| \to \pm\infty$, and so this sum exists for each x. Furthermore, we have taken the liberty of naming the sum $u(x)$ since indeed the infinite sum is equal to the original function if it is smooth.

Thus we have two representations of the *same* information. In the "physical space representation," we have the function $u(x)$, consisting of its values (real or complex) at all the points $x \in [0, L]$. In the "Fourier representation," we have an infinite sequence of complex numbers. You may think this is very strange, since the first form contains uncountably many different numbers, while the second has only countably many. In fact, the assumption of smoothness reduces the amount of information in the function u to countably many degrees of freedom.

A more formal way to say the above is to form the approximations to u consisting of only finitely many Fourier modes:

$$u_n(x) \;=\; \sum_{k=-n}^{n} \hat{u}_k \, e^{2\pi i k x/L}.$$

Certainly, for any finite n, this sum exists, and the question is whether $u_n(x) \to u(x)$ in some reasonable sense as $n \to \infty$. If u is real analytic, then this convergence is true in *every* sense. In general, the exact nature of the convergence depends on precisely how smooth u is:[2]

[2]See Walter A. Strauss, *Partial Differential Equations: An Introduction*, John Wiley & Sons 1992.

- If u has continuous derivatives of order 2, then $u_n \to u$ uniformly.
- If u is piecewise continuous with piecewise continuous first derivative, then $u_n \to u$ at points where u is continuous. At points of discontinuity, the series converges to the average of the left and right values.
- If u is square-integrable, then $\int |u_n - u|^2 \to 0$. This is often the most suitable notion, since this guarantees the existence of the associated inner product. In addition, we have *Parseval's identity*

$$\frac{1}{L} \int_0^L |u(x)|^2 \, dx \;=\; \sum_{k=-\infty}^{\infty} |\hat{u}_k|^2.$$

  The transform is just a rotation in some infinite-dimensional space.

These matters were very controversial throughout the 19th century; Joseph Fourier is credited with proposing that a series of *continuous* trigonometric polynomials could actually converge to a *discontinuous* function.

The situtation is somewhat reminiscent of finite elements, since there also, we projected our function into a space spanned by a discrete set of basis functions, and we reconstructed the original as well as we could. Here, we can recover the original function *exactly*, assuming some smoothness, by taking more and more basis functions. That was also true for finite elements, but the geometric complexity increased as the space got larger.

Convergence subtleties appear even for smooth $u(x)$. Consider

$$u(x) \;=\; \sum_{m=-\infty}^{\infty} \left(\tanh \frac{x - x_1 - mL}{\epsilon} \;-\; \tanh \frac{x - x_2 - mL}{\epsilon} \right) \;-\; 1$$

This function has steps of thickness ϵ at $x = x_1$ and $x = x_2$; the infinite sum makes it exactly periodic although the terms with $m \neq 0$ contribute an extremely small amount. Figure 4 shows this function with $L = 1$, $x_1 = 0.3179$, $x_2 = 0.8751$ (randomly chosen to break symmetries), and $\epsilon = 0.02$. Because ϵ is small, the decay of the Fourier modes is extremely slow although exponential, and they do not reach roundoff error until $k = 200$.

If we try to resolve this function with insufficiently many Fourier modes, we see the results shown in Figure 5, which shows $u_n(x)$ for $n = 8$. As n increases, the convergence is extremely rapid. The critical value for n is propotional to L/ϵ, so that the wavelength of the nth mode, L/n matches the length scale of the function being approximated, ϵ. In practice, if the spectrum does not drop to roundoff levels, you don't have enough modes.

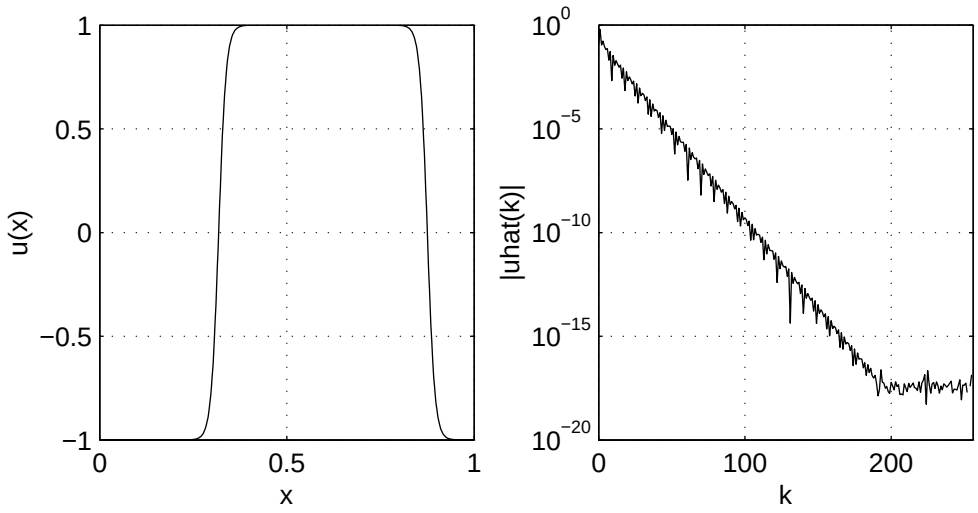

Figure 4: An even more realistic example—see text for explanation.

This example is artificial, but similar thin layers arise naturally in many physical problems such as fluid turbulence.

Two final comments:

- We shall not consider the extension to functions defined on all $\mathbb{R}$, with suitable decay assumptions at $\pm\infty$; then $\hat{u}$ becomes a function of a continuous variable k. Periodicity in x translates to discreteness in k.

- Everything we have said can be extended to functions of several variables, defined on a rectangle $[0, L_1] \times \cdots \times [0, L_d]$ instead of the interval $[0, L]$; the wave index k becomes a multi-index $(k_1, \ldots, k_d)$.

Spectral Methods

We have now argued that we may look at a particular function either in the original form $u(x)$, or in the Fourier representation $\hat{u}_k$. What can we do with the latter that we couldn't do with the former?

The thing we do most often in numerical computations is to differentiate. If $v(x) = u'(x)$, then

$$\hat{v}_k \;=\; \frac{1}{L}\int_0^L e^{-2\pi i k x/L} v(x)\,dx \;=\; \frac{2\pi i k}{L}\,\frac{1}{L}\int_0^L e^{-2\pi i k x/L} u(x)\,dx \;=\; 2\pi i k\,\hat{u}_k$$

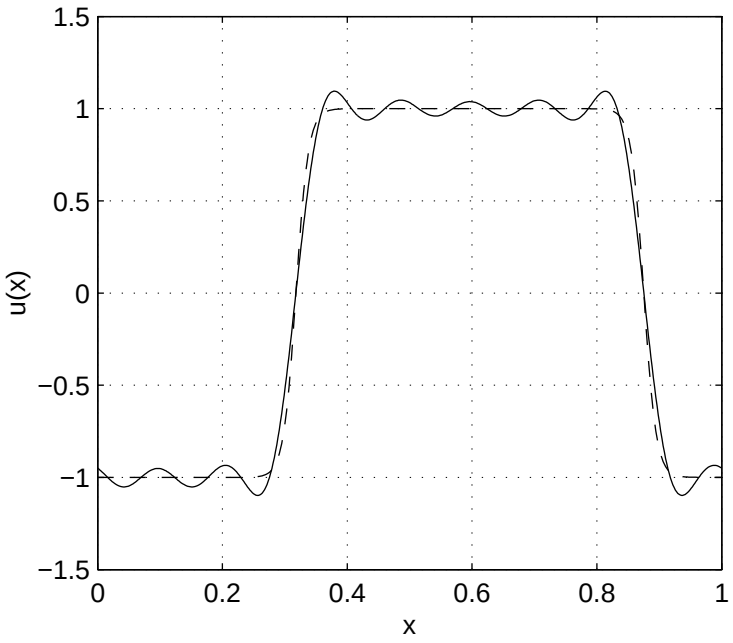

Figure 5: A reconstruction of the example in Figure 4 using only 8 modes and showing the Gibbs phenomenon (overshoot at discontinuities) is visible. Even though the function is smooth on fine scales, on intermediate scales you see the near-discontinuities.

in which we have integrated by parts, differentiating the exponential and changing sign (the boundary terms are zero because of periodicity). Equivalently, we can differentiate the inversion formula (we can take the derivative inside the sum if the $\hat{u}_k$ decay fast enough) to write

$$u'(x) \;=\; \frac{d}{dx} \sum_{k=-\infty}^{\infty} e^{2\pi i k x/L} \hat{u}_k \;=\; \sum_{k=-\infty}^{\infty} \frac{2\pi i k}{L} e^{2\pi i k x/L} \hat{u}_k.$$

The conclusion is the same: the Fourier series of the derivative is $2\pi ik/L$ times the Fourier series of the function. Whereas differentiation of a function involves a subtle limiting relationship between the functional values at "neighboring" points in space, differentiation of the Fourier represenation is a "diagonal" operation that acts on one mode at a time.

Similarly, convolution of two functions, an integral operator in physical space, is a mode-by-mode multiplication in Fourier space. And certain

other integral operators (Hilbert transform, for example, are trivial diagonal operations in the Fourier representation.

However, *nonlinear* operations such as multiplication in physical space generally correspond to infinite sums in Fourier space. Some operations are easier in one represenation, some in the other.

Discrete Fourier Transform

Suppose $u_0, \ldots, u_n$ are an array of n real or complex numbers. The forward *discrete Fourier transform* (FFT) is the transformation

$$\hat{u}_k \;=\; \frac{1}{n} \sum_{j=0}^{n-1} e^{-2\pi i j k/n} u_j, \qquad k = 0, \ldots, n-1.$$

As with the continuous transform, there are several conventions for choice of the prefactor. You may easily verify that the inverse transformation is

$$u_j \;=\; \sum_{k=0}^{n-1} e^{2\pi i j k/n} \hat{u}_k, \qquad j = 0, \ldots, n-1$$

using the orthogonality relation

$$\frac{1}{n} \sum_{j=0}^{n-1} e^{-2\pi i j k/n} e^{2\pi i \ell k/n} \;=\; \delta_{j\ell}.$$

Again, we have the conservation of discrete L^2 norm

$$\frac{1}{n} \sum_{j=0}^{n-1} |u_j|^2 \;=\; \sum_{k=0}^{n-1} |\hat{u}_k|^2.$$

This operation is discrete on both sides: it takes a set of n numbers into a different set of n numbers. Since it is linear, it is equivalent to multiplication by a unitary matrix Q; inversion is multiplication by $Q^{-1} = Q^*$. Probably the most significant numerical algorithm of all time was the discovery that these operations could be performed in time $n \log n$ rather than n^2.[3]

Let us point out one additional feature of the transform relationships. Although we have said that j, k vary only from 0 to $n-1$, in fact you could use the formulas to compute u_j or $\hat{u}_k$ for *any* integer values. The result will be a periodic sequence in both physical and Fourier space. Originally we had periodicity in x, corresponding to discreteness in k; now we have added discreteness in x which gives periodicity in k.

[3] If you're looking for a code, the current fastest one is `FFTW`; this is what's in Matlab.

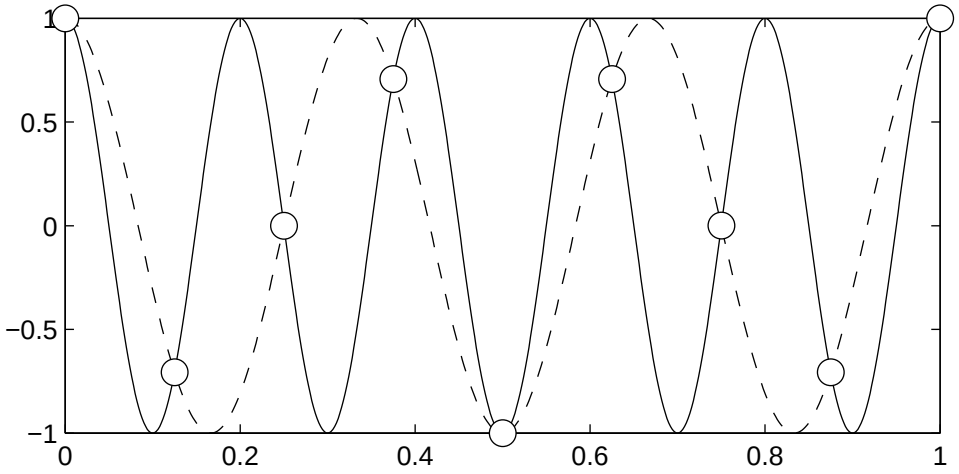

Figure 6: Aliasing error. The grid has $n = 8$, and the mode $k_1 = 3$ is plotted with dashed line and circles. The mode $k_2 = 5 = n - k_1$ looks exactly the same on this grid. The highest resolvable frequence would be $k = n/2 = 4$.

Relation between discrete and continuous transforms

The interesting question is what happens if the original numbers u_j are samples of a smooth function $u(x)$: what is the relationship between the discrete coefficients $\hat{u}^{\mathrm{d}}_k$ and the continuous coefficients $\hat{u}^{\mathrm{c}}_k$.

Thus, assume that $u_j = u(x_j)$, where $x_j = jL/n$ are equally spaced grid points. Then using the representation formula we readily calculate

$$\begin{aligned}\hat{u}^{\mathrm{d}}_k &= \frac{1}{n}\sum_{j=0}^{n-1} e^{-2\pi ijk/n} u(jL/n) = \frac{1}{n}\sum_{j=0}^{n-1} e^{-2\pi ijk/n} \sum_{\ell=-\infty}^{\infty} e^{2\pi i\ell(jL/n)/L}\hat{u}^{\mathrm{c}}_\ell \\ &= \sum_{\ell=-\infty}^{\infty} \left(\frac{1}{n}\sum_{j=0}^{n-1} e^{2\pi i(\ell-k)j/n} \right) \hat{u}^{\mathrm{c}}_\ell .\end{aligned}$$

The sum in parentheses has the value 1 when $k - \ell$ is *any integer multiple of* n. Thus

$$\hat{u}^{\mathrm{d}}_k = \sum_{m=-\infty}^{\infty} \hat{u}^{\mathrm{c}}_{k+mn}.$$

The discrete mode k collects all the energy from all the periodic images of mode k in the continuous spectrum.

This is the reason it is so important that the wave amplitudes decay rapidly, and that we have a large enough value of n: if both of these are true, then the contribution from the other periodic modes is extremely small. However, if $u(x)$ is not a smooth enough function, or n is not large enough, then modes for small values of k will contain energy from much higher modes. This is called *aliasing error.*

The highest wave number that can be represented on a grid of length n has $k = n/2$, or $\xi = \pi n/L$. The wavelength is $2\pi/\xi = 2L/n = 2h$; these modes are ± 1 at the grid points.

Computation of derivatives

Now, suppose again that we want to compute approximate values $v_j \approx u'(x_j)$ of the derivative at the grid points x_j, given only the values u_j of the function at the grid points. This is precisely the problem we considered extensively in finite differences, where we generally settled on the centered difference formula $v_j = (u_{j+1} - u_{j-1})/2h$. This expression has a second-order error $\mathcal{O}(h^2)$. In special cases, we had reasons to take a one-sided difference that was only first-order accurate. Conversely, if we wanted higher-order accuracy we could achieve it by using more neighbors on either side. Now we want to see if we can come up with a better formula using Fourier transformations.

Here is the algorithm:

1. Take the discrete Fourier transform of the given discrete data. This is equivalent to interpolating the function by a trigometric polynomial of degree n that passes exactly through the given points, and is very smooth in between.

2. Differentiate that polynomial, which is easily done by multiplying each mode by its wave number k.

3. Take the inverse transform to get values of the derivative at the grid points in physical space.

There is one tricky issue: Since each element of the discrete tranform is a sum of an infinite sequence of wave components of the original function, what wave number k should we use in step 2?

Answer: Use the *smallest* value of k, meaning closest to zero, since that is the one corresponding to the largest energy. For $k = 0, \ldots, (n/2) - 1$, this is simply k itself. But for $k = (n/2) + 1, \ldots, n - 1$, the image $k - n$ is closer to zero than k itself. For $k = n/2$, we don't know which one is closer. Since

we hope there is no energy there anyway, it is usually safe to set that mode to zero. Thus the algorithm sets

$$\hat{v}_k = \begin{cases} \dfrac{2\pi i k}{L}\,\hat{u}_k, & k = 0, \ldots, \dfrac{n}{2} - 1 \\ 0, & k = n/2 \\ -\dfrac{2\pi i (n-k)}{L}\,\hat{u}_k, & k = \dfrac{n}{2} + 1, \ldots, n-1. \end{cases}$$

You do a forward transform before, and an inverse transform after. Note that the multiplier is not the index k, but the *physical* wave number $\xi = 2\pi k/L$.

This is a *global* algorithm, since each output value v_j depends on each input value u_j. This is the logical limit of taking derivatives using more and more neighbors. Its accuracy as a function of h is better than any power of h, as long as the grid is fine enough to resolve all features of the solution. If the solution is underresolved, it will give dramatically *bad* answers, like any high-order method. To summarize, the *good* things about spectral methods are

- Spectacular accuracy on smooth periodic problems, and
- Ability to handle some nonlocal integral effects.

The *bad* things are

- Terrible accuracy if the solution is not smooth, or cannot be smoothly extended to a periodic functions (if the boundary conditions are not homogeneous Neumann or Dirichlet).
- The grid must be uniform (equally spaced).
- It is very "finicky:" if the code isn't completely correct it will not do anything reasonable, and it will be hard to see what is wrong.

The Nonlinear Schrödinger Equation

This is a paradigm example of a nonlinear wave equation. It arises physically in a variety of contexts, ranging from nonlinear optics (index of refraction depends on light intensity) to dynamics of superfluid helium. It illustrates a lot of interesting behavior,[4] and is a stellar example of spectral methods.

[4]See Catherine Sulem and Pierre-Louis Sulem, *The Nonlinear Schrödinger Equation: Self-Focusing and Wave Collapse*, Applied Mathematical Sciences Vol. 139, Springer 1999.

Loosely speaking, the equation comes in two varieties: *defocusing* or with a *repulsive* potential, and *focusing*, or with an *attractive* potential. The difference in the equation is just a change of sign, but the behavior of solutions is very different. We start with the defocusing case.

Defocusing The PDE for the complex function $u(x,t)$ is

$$i\,u_t \;=\; u_{xx} \;-\; |u|^2 u.$$

In higher dimensions, the derivative just becomes the Laplacian. We shall suppose that the equation is to be solved on a periodic domain, $[0,L]$ in one dimension, or $[0,L]^d$ in d dimensions. This called the *cubic nonlinearity;* another popular choice is $|u|^{2\sigma}u$ with $\sigma = 1, 2, \ldots$.

One way to understand this equation (as always) is to look at conserved quantities, of which two are of special importance. The *particle number* is

$$N(t) \;=\; \int |u(x,t)|^2\,dx$$

(you may normalize it however you like). If the domain were unbounded (no periodicity), then we would impose decay conditions at ∞ so this quantity would be finite. To calculate the derivative, we recall that $|u|^2 = u\,\bar{u}$, so

$$\begin{aligned}
\frac{dN}{dt} \;&=\; \int \left(u\,\bar{u}_t \;+\; u_t\,\bar{u} \right) dx \\
&=\; \int_0^L \Big(u\left(i\bar{u}_{xx} - i|u|^2\bar{u}\right) \;+\; \bar{u}\left(-iu_{xx} + i|u|^2 u\right) \Big)\,dx \\
&=\; \int_0^L \Big(i\big(-u_x\bar{u}_x + u_x\bar{u}_x\big) \;+\; i|u|^2\big(-u\bar{u} + \bar{u}u\big) \Big)\,dx \\
&=\; 0,
\end{aligned}$$

where we have integrated by parts once on the derivative terms. Thus N is a constant of the motion.

The other conserved quantity is the energy, or *Hamiltonian*

$$H(t) \;=\; \int \left(|\nabla u|^2 \;+\; \frac{1}{2}|u|^4 \right) dx$$

(for general σ the second term is $|u|^{2\sigma+2}/(\sigma+1)$). Similar calculations as above show that $H(t)$ is also conserved.

Note that this energy is very similar to the one we used in the Allen-Cahn equation, except that there the time dynamics was *gradient descent:*

it decreased the energy as rapidly as possible. Here the dynamics is that of a *Hamiltonian system*, that moves "across" the energy.

For the 1-D NLS with cubic nonlinearity, there is in fact an infinite family of conserved quantities, which leads to the study of "integrable systems." In general, though (multi-dimensions or $\sigma > 1$), these are the only two and so we will focus on them.

Existence of the conserved energy H is of immense importance for understanding the behavior of solutions. Since it is the sum of two positive terms, neither of them can become larger than the initial value of H. This provides a bound on u and on its first derivative, which guarantees that the solution remains well-behaved for all time.

Focusing The *focusing* or *attractive* form of the equation is

$$i\,u_t \;=\; u_{xx} \;+\; |u|^2 u,$$

in which only the sign of the nonlinear term (relative to the derivative) has been changed. The particle number N is conserved as above. But now the conserved Hamiltonian is

$$H(t) \;=\; \int \left(|\nabla u|^2 \;-\; \frac{1}{2}|u|^4 \right) dx$$

with a sign change. Now the conserved energy is the *difference* of two positive terms. Nothing prevents both of them from becoming infinite together: the solution can and does develop very interesting singularities.

Splitting

A standard way to handle equations like NLS, in which u_t is the sum of two terms, is to look at the PDEs in which only one term appears at a time. Here the structure is particular simple, because each one is a type of rotatation, one in physical and one in Fourier space.

Rotation in Fourier space First, let us consider the linear wave equation

$$i\,u_t \;=\; u_{xx}.$$

Since this is a linear equation with constant coeffiecients, its solutions are described by looking at individual Fourier modes. With $u(x,t) = U(t)\exp(i\xi x)$, we immediately see that $iU'(t) = -\xi^2 U(t)$, so we can write the solution as

$$u(x,t) \;=\; U_0 e^{i\xi^2 t} e^{i\xi x} \;=\; U_0 e^{i\xi(x+\xi t)}.$$

Of course, in a periodic box of length L, ξ is restricted to the discrete values $2\pi k/L$ for k an integer.

The phase speed $c = -\xi$ of the mode depends on ξ. Thus this wave equation is *dispersive:* different wavelengths propagate at different speeds. This is in contrast to hyperbolic systems like $u_{tt} = c^2 u_{xx}$, in which each different wavelength propagates at the *same* speed c (except for errors introduced by the discretization). Only for hyperbolic systems can an initial disturbance move while preserving shape. For a dispersive system any solution inevitably breaks up into a combination of many waves.

It is trivial to compute this solution in the Fourier representation. If

$$u(x,0) \;=\; \sum_k \hat{u}_k(0)\, e^{2\pi i k x/L},$$

(whether the sum is finite or not), then the solution at later times is

$$u(x,t) \;=\; \sum_k \hat{u}_k(t)\, e^{2\pi i k x/L}, \qquad \hat{u}_k(t) \;=\; e^{i(2\pi k/L)^2 t}\, \hat{u}_k(0).$$

Each mode just sits there and spins independently of all the others. Note that the rotation rate increases quadratically with k. On a discrete grid of size n, the highest mode is $n/2$, so the fastest rotation rate is $(\pi n/L)^2$. The rotation period of the kth mode is $T_k = 2\pi/(2\pi k/L)^2 = L^2/2\pi k^2$, so the shortest period is $T_{n/2} = 2L^2/\pi n^2$.

If you wanted to do an explicit method, without using the special rotation structure, then you would need to take the time step τ smaller than this intrinsic time. It is just as for the diffusion equation, where we needed to resolve the decay time of the smallest modes. We could have avoided the problem there by doing a Fourier transform and evolving each mode separately, as we are doing here (but the FFTs would have been much slower than the finite differencing).

Rotation in physical space The other half of the problem is

$$i\, u_t \;=\; -|u|^2 u.$$

Since this equation has no space derivatives, it is just an ODE at each point, and the solution is immediately seen to be

$$u(x,t) \;=\; e^{i|u|^2 t}\, u(x,0).$$

Again, this is just a rotation, that keeps constant the value of $|u(x,t)|^2$. Each separate spatial point just sits and spins independently.

Pseudo-spectral algorithm

Now we can put all this together into a numerical algorithm. The idea is that we do each effect in alternation. But since one is easiest to do in physical space, the other in Fourier, we have to keep transforming back and forth. That is why it is called "pseudo-spectral:" a fully spectral algorithm would do all operations in Fourier space, but that works for very few problems since almost everything has nonlinearity. Such algorithms only became possible with the invention of the fast Fourier transform.

Here is the outline. We start with a list $(u_0^m, \ldots, u_{n-1}^m$ representing the solution at time level m (keeping n for the number of space grid points). Here's how we get to level $m+1$ with time step τ.

1. *Forward transform.* Do an FFT on u^m to compute the amplitudes $\hat{u}_k^m$ in the Fourier-space representation.

2. *Fourier rotation.* Spin each mode as follows:

$$\hat{u}_k^{m+1/2} \;=\; \exp\bigl(i\xi_k^2\tau\bigr)\,\hat{u}_k^m, \qquad \xi_k = \frac{2\pi}{L}\min\{k, n-k\}$$

   (it should be $k-n$ rather than $n-k$, but we square it anyway).

3. *Backward transform.* Do an inverse FFT on $\hat{u}^{m+1/2}$ to compute the elements $u_j^{m+1/2}$ in the physical-space representation.

4. *Physical rotation.* Rotate each grid point in place:

$$u_j^{m+1} \;=\; \exp\left(i\left|u_j^{m+1/2}\right|^2\tau\right)\,u_j^{m+1/2}$$