

Mat1062: Computational Methods for PDE

Mary Pugh

March 20, 2008

Ownership

These notes are built upon those of Rob Almgren who taught an analogous course in 2003. Whatever you learn of value from them is due to him. All mistakes and sources of confusion are to be blamed on me.

An Example

Spectral methods are similar to finite element methods in that they involve representing a solution as a linear combination of basis functions and using an inner product to find the coefficients of the linear combination. The difference is that spectral methods use the eigenfunctions of the linear operator as the basis functions.

Consider the operator $\mathcal{L} = \partial_{xx}$ on $[0, L]$. As we observed earlier, the eigenfunctions depend on the boundary conditions. For example,

$$\begin{aligned} u(0) = u(L) = 0 &\implies \sin(j\pi x/L) \quad j \in \mathbb{N} \\ u_x(0) = u_x(L) &\implies 1, \quad \cos(j\pi x/L) \quad j \in \mathbb{N} \\ u(0) = u(L) &\implies \exp(ij\pi x/L) \quad j \in \mathbb{Z} \\ u(0) = u_x(L) = 0 &\implies \sin((j + 1/2)\pi x/L) \quad j \in \mathbb{N} \end{aligned}$$

Here is an example of how we would use the basis to find a solution of an initial boundary value problem. Take $u_t = u_{xx} + f$ on $0 \leq x \leq L$, with $u = 0$ at $x = 0, L$. We write $u(x, t)$ in terms of the eigenfunctions

$$\phi_j(x) = \sin\left(\frac{j\pi x}{L}\right).$$

Then any approximation of the form

$$u_N(x, t) = \sum_{j=1}^N c_j(t) \phi_j(x)$$

satisfies the boundary conditions.

We define the inner product to be

$$\langle u, v \rangle = \frac{2}{L} \int_0^L u(x) v(x) dx.$$

Then

$$\langle \phi_j, \phi_j \rangle = \frac{2}{L} \int_0^L \sin^2\left(\frac{j\pi x}{L}\right) dx = \frac{1}{L} \int_0^L \left(1 - \cos\left(\frac{j2\pi x}{L}\right)\right) dx = 1$$

and for $j \neq k$,

$$\begin{aligned} \langle \phi_j, \phi_k \rangle &= \frac{2}{L} \int_0^L \sin\left(\frac{j\pi x}{L}\right) \sin\left(\frac{k\pi x}{L}\right) dx \\ &= \frac{1}{L} \int_0^L \left(\cos\left(\frac{(j-k)\pi x}{L}\right) - \cos\left(\frac{(j+k)\pi x}{L}\right)\right) dx = 0. \end{aligned}$$

Thus $\langle \phi_j, \phi_k \rangle = \delta_{jk}$ as we hoped. Furthermore,

$$\mathcal{L}\phi_j = -\left(\frac{j\pi}{L}\right)^2 \phi_j$$

and so the ODE system for the c_j is

$$\frac{dc_k}{dt} = -\left(\frac{k\pi}{L}\right)^2 c_k + f_k(t)$$

with

$$f_k(t) = \frac{2}{L} \int_0^L f(x, t) \sin\left(\frac{k\pi x}{L}\right) dx.$$

In this way, solving the PDE has been reduced to solving infinitely many ODE.

The above is the exact same thing that you would have found via separation of variables. When you first learnt separation of variables, you were learning to work with the eigenfunctions of the Laplacian and were taking advantage of the fact that, for reasonable boundary conditions, the eigenfunctions are orthogonal to one another. It may not have been presented to you in this language.

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

$$\hat{u}_k = \langle u, \phi_k \rangle = \frac{1}{L} \int_0^L u(x) e^{-ik2\pi x/L} dx, \quad \text{for } k = 0, \pm 1, \pm 2, \dots \quad (1)$$

(the “hat” is commonly used to denote the Fourier transform). Above, we use the inner product for complex-valued functions:

$$\langle f, g \rangle := \frac{1}{L} \int_0^L f(x) \overline{g(x)} dx$$

where the overline denotes complex conjugate. 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 k th mode has k waves in the period interval.

The symmetries of the sequence $\{\hat{u}_k\}$ mirror those of the function u . First, if u is real-valued, then $\hat{u}_{-k} = \overline{\hat{u}_k}$. 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 periodically, 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*

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

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

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

The infinite sequence $\{\hat{u}_k\}$ is all zeros, with a single value equal to one at index 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

$$u(x) = a_1 e^{i\ell_1 2\pi x/L} + \dots + a_n e^{i\ell_n 2\pi x/L} \implies \hat{u}_k = a_1 \delta_{k\ell_1} + \dots + a_n \delta_{k\ell_n} \quad (2)$$

The sequence is all zeros, with values a_1, \dots, a_n at indices ℓ_1, \dots, ℓ_n .

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

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

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 the symmetry relationships above. In general, *translation* in physical space corresponds to a *rotation* of the Fourier coefficients.

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:

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

assuming the numbers $\{a_\ell\}_{\ell=-\infty}^{\infty}$ are such that the infinite sum exists.

Decay behavior

The decay of the coefficients $\{\hat{u}_k\}$ as $k \rightarrow \pm\infty$ 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.

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 $-\rho \leq \text{Im } x \leq \rho$. Then the magnitudes of the Fourier coefficients \hat{u}_k decay at least exponentially with parameter controlled by ρ :*

$$|\hat{u}_k| \sim \mathcal{O}(C e^{-c\rho|k|}), \quad |k| \rightarrow \infty.$$

where $c = 2\pi/L$ and C depends on ρ , L and u but not on k .




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 $k = 15$.

The trigonometric polynomials (2) are analytic in the entire complex plane, and the Fourier series are exactly zero beyond a finite index; this is certainly rapid decay. For more realistic examples, see Figures 1, 2.

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

$$\hat{u}_k = \frac{1}{L} \int_0^L u(x) e^{-ik2\pi x/L} dx = e^{2\pi k\rho/L} \cdot \frac{1}{L} \int_0^L u(x + i\rho) e^{-ik2\pi x/L} dx.$$

Now since u is analytic along the line $\text{Im } x = \rho$, it is continuous and therefore bounded. Thus 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 $\text{Im } x = -\rho$. This finishes the proof.

An alternate proof would be to change variables to $z = \exp(i2\pi x/L)$. 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




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.

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 $\text{Im } x = \pm\rho$. If u has singularities in the complex plane then the decay rate of $|\hat{u}_k|$ also has an algebraic term multiplying the exponential term that can be used to extract information about the nature of the closest singularities of $u(x)$ to the real axis.¹

As an application of the theorem, consider

$$u_\epsilon(x) = \frac{1}{\epsilon^2 + \sin^2(\pi(x - 1/2))} \quad (3)$$

on $[0, 1]$. This is a periodic function, by construction, see the left plot in Figure 4. It has a maximum value of $1/\epsilon^2$ achieved at $x = 1/2$. The smaller ϵ is, the taller and narrower the peak at $x = 1/2$ will be. If we were to integrate u_ϵ we see that we get a function which has a sharp transition at $x = 1/2$, and so the following discussion is relevant not just for localized

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.

```

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 and 2. `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.

“blow-ups” but also for localized sharp fronts. We can use the Fourier coefficients to determine how many meshpoints we need to properly resolve u_ϵ for a given value of ϵ . Consider the related function $1/(\epsilon^2 + x^2)$. This function is real analytic and has two poles in the complex plane at $\pm i\epsilon$. Similarly, u_ϵ has poles at $.5 \pm i\epsilon$. This means that by the above theorem, the Fourier coefficients will decay like

$$|\widehat{u}_{\epsilon k}| \approx e^{-|k|\epsilon}.$$

If we seek the wave number at which the coefficient is of order of round-off we find

$$e^{-|k_c|\epsilon} \approx 10^{-14} \implies k_c \approx 1/\epsilon$$

And so we need wave numbers up to $\mathcal{O}(1/\epsilon)$ to fully resolve u_ϵ . This points to one of the weaknesses of spectral methods. As we will see shortly, to resolve the N th wave number we need $2N$ meshpoints, uniformly distributed. And so we have to use lots of points everywhere because of this fine structure localised in one area. This is unavoidable because spectral methods are




Figure 4: The function u_ϵ as given by (3) with $\epsilon = 1/40$. The function is sampled on a uniform mesh with $dx = .0005$. Left plot: the graph of u_ϵ as a function of x . Right plot: Because there are 2000 sample points, 1000 frequencies can be resolved. This is a plot of $\log_{10}(|\hat{u}_{\epsilon k}|)$ as a function of k . The Fourier coefficients reach the level of round-off around $k = 800$.

based on eigenfunctions which are inherently non-localized objects. Using finite element or finite difference methods, we could concentrate meshpoints where they are most needed.² For the ϵ used in Figure 4 we need enough meshpoints to be able to find the $k = 800$ modes: \hat{u}_{800} and \hat{u}_{-800} . To see the effect of not having enough modes, consider Figure 5. Here, I took the Fourier Transform of u_ϵ and then constructed a function v by using $\hat{u}_{\epsilon-100}, \hat{u}_{\epsilon-99}, \dots, \hat{u}_{\epsilon 99}, \hat{u}_{\epsilon 100}$. (That is, v is what one would get by applying a low-pass filter to u_ϵ .) In the left plot of Figure 5 we see the effect of under-resolution on the graph near the local maximum. In the right plot of Figure 5 we see the effect of under-resolution on the graph near the flanking flat areas. Note that there are significant oscillations.

If $u(x)$ is not real analytic, then the Fourier coefficients have an algebraic rate of 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).

²This is, of course, more complicated than I'm making it sound. Finite elements and finite-difference methods have to work hard if the structure of interest is moving in time. And of course, if fine structures are appearing at multiple places in space, and moving in time (think of turbulent fluid flow) then you may as well resolve all points in space equally. Also, as we'll see, spectral methods are more accurate than finite element and finite difference methods. This means that you may end up having to use a lot more finite elements than you'd expected to, resulting in a scheme that's no faster than the spectral method.




Figure 5: We took the Fourier transform of u_ϵ , set $\hat{u}_{\epsilon k} = 0$ for all $|k| > 100$, and then applied an inverse Fourier transform resulting in a function v . Left plot: The solid line is u_ϵ and the dashed line is v . The graph has been chosen to focus in near the peak of u_ϵ . Right plot: The solid and dashed lines are as in the left plot. The graph has been chosen to focus in on the left flank of the peak.

Note that we are referring to the smoothness of u after it has been extended to a periodic function on \mathbb{R} . If $u(0) \neq u(L)$, then the extension will be discontinuous and the Fourier coefficients will decay like $1/|k|$.

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 transform

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

$$u(x) = \sum_{k=-\infty}^{\infty} \hat{u}_k e^{ik2\pi x/L}.$$

As described above, if $u(x)$ is smooth, then $|\hat{u}_k|$ decays rapidly as $|k| \rightarrow \pm\infty$, 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:

$$u_n(x) = \sum_{k=-n}^n \hat{u}_k e^{ik2\pi x/L}.$$

Certainly, for any finite n , this sum exists, and the question is whether $u_n(x) \rightarrow u(x)$ in some reasonable sense as $n \rightarrow \infty$. 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:³

- If u has continuous derivatives of order 2, then $u_n \rightarrow u$ uniformly.
- If u is piecewise continuous with piecewise continuous first derivative, then $u_n \rightarrow u$ 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 $\int |u_n - u|^2 \rightarrow 0$. This is often the most suitable notion, since this guarantees the existence of the associated inner product. In addition, we have *Parseval's identity*

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

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 situation 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.

Two final comments:

³See Walter A. Strauss, *Partial Differential Equations: An Introduction*, John Wiley & Sons 1992.

- 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, \dots, 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

$$\hat{v}_k = \frac{1}{L} \int_0^L e^{-ik2\pi x/L} v(x) dx = \frac{ik2\pi}{L} \frac{1}{L} \int_0^L e^{-ik2\pi x/L} u(x) dx = 2\pi ik \hat{u}_k$$

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

$$u'(x) = \frac{d}{dx} \sum_{k=-\infty}^{\infty} \hat{u}_k e^{ik2\pi x/L} = \sum_{k=-\infty}^{\infty} \frac{ik2\pi}{L} \hat{u}_k e^{ik2\pi x/L}.$$

The conclusion is the same: the Fourier series of the derivative is $ik2\pi/L$ 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 representation 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 representation, some in the other.

Discrete Fourier Transform

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

$$\hat{u}_k^d = \frac{1}{n} \sum_{j=0}^{n-1} u_j e^{-ik2\pi j/n}, \quad k = 0, \dots, n-1. \quad (4)$$

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

$$u_j = \sum_{k=0}^{n-1} \hat{u}_k^d e^{ik2\pi j/n} \quad j = 0, \dots, n-1$$

using the orthogonality relation

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

To see why this orthogonality relation holds, first note that if $\ell = j$ then it's clearly true. If $\ell \neq j$ and $0 \leq \ell, j \leq n-1$ then we know that $1 - n \leq \ell - j \leq n-1$ and so

$$e^{i2\pi(\ell-j)/n} \neq 1 \implies \frac{1}{n} \sum_{k=0}^{n-1} \left(e^{i2\pi(\ell-j)/n} \right)^k = \frac{1}{n} \frac{1 - \left(e^{i2\pi(\ell-j)/n} \right)^n}{1 - e^{i2\pi(\ell-j)/n}} = 0$$

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

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

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 $Q^{-1} = Q^*$.

Naively, it takes $\mathcal{O}(n)$ operations to compute \hat{u}_k^d using (4). And so to compute $\hat{u}_0^d, \hat{u}_1^d, \dots, \hat{u}_{n-1}^d$ sequentially would take $\mathcal{O}(n^2)$ operations. 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 . This algorithm is the *Fast Fourier Transform*. If it weren't for the FFT, spectral methods wouldn't be so popular. Note that the FFT is for these particular basis functions. If you were working in a domain that had other eigenfunctions (Bessel functions in the radial direction and trig functions in the angular direction, say) then you would need a FBT for the radial stuff and a FFT for the angular stuff.




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

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 Fourier coefficients, \hat{u}_k^d , as given by (4) and the continuous Fourier coefficients, \hat{u}_k^c , as given by (1)?

The discrete Fourier transform has to do with a collection of numbers u_0, \dots, u_{n-1} . The continuous Fourier transform has to do with a function sampled at n points: $x_j = jL/n$. To relate the two, we assume that $u_j = u(x_j)$, where $x_j = jL/n$ are equally spaced grid points. Then using the discrete Fourier transform (4) we readily calculate

$$\begin{aligned} \hat{u}_k^d &= \frac{1}{n} \sum_{j=0}^{n-1} u(jL/n) e^{-ik2\pi j/n} = \frac{1}{n} \sum_{j=0}^{n-1} \left(\sum_{\ell=-\infty}^{\infty} \hat{u}_\ell^c e^{i\ell 2\pi(jL/n)/L} \right) e^{-ik2\pi j/n} \\ &= \sum_{\ell=-\infty}^{\infty} \hat{u}_\ell^c \left(\frac{1}{n} \sum_{j=0}^{n-1} e^{i(\ell-k)2\pi j/n} \right). \end{aligned}$$

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

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

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 we have a large enough value of n . We recall that the discrete Fourier transform (4) was defined for $k = 0, 1, \dots, n-1$. Consider k such that $0 \leq k \leq \lfloor n/2 \rfloor$. Then

$$\hat{u}_k^d = \hat{u}_k^c + \sum_{m=1}^{\infty} \hat{u}_{k-mn}^c + \sum_{m=1}^{\infty} \hat{u}_{k+mn}^c$$

and $|k+mn| > \lfloor n/2 \rfloor$ for all m . That is, if $0 \leq k \leq \lfloor n/2 \rfloor$ then \hat{u}_k^d equals \hat{u}_k^c plus energy from higher modes. If we have chosen n sufficiently large then all of these contributions from higher modes will be at the level of round-off and so we are comfortable treating \hat{u}_k^d as the same thing as \hat{u}_k^c .

Similarly, if $\lfloor n/2 \rfloor < k \leq n-1$ then

$$\hat{u}_k^d = \hat{u}_{k-n}^c + \sum_{m=2}^{\infty} \hat{u}_{k-mn}^c + \sum_{m=0}^{\infty} \hat{u}_{k+mn}^c$$

Here, we note that $|k-n| \leq \lfloor n/2 \rfloor$ and that for the first series, $|k-mn| > n$ and for the second series $|k+mn| > \lfloor n/2 \rfloor$. And so in this case \hat{u}_k^d equals \hat{u}_{k-n}^c plus energy from higher modes. Again, if we have chosen n sufficiently large then all of these contributions from higher modes will be at the level of round-off and so we are comfortable treating \hat{u}_k^d as the same thing as \hat{u}_{k-n}^c .

To sum up, the discrete Fourier transform (4) gives us

$$\hat{u}_0^d, \hat{u}_1^d, \dots, \hat{u}_{n-1}^d.$$

To interpret \hat{u}_k^d as a Fourier coefficients we find the n -shift that brings k closest to 0:

$$\hat{u}_k^d = \begin{cases} \hat{u}_k^c & 0 \leq k \leq \lfloor n/2 \rfloor \\ \hat{u}_{k-n}^c & \lfloor n/2 \rfloor < k \leq n-1 \end{cases}$$

if n has been taken large enough so that all the higher energy modes have Fourier coefficients at the level of round-off error.

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. This highest frequency is called the Nyquist frequency.