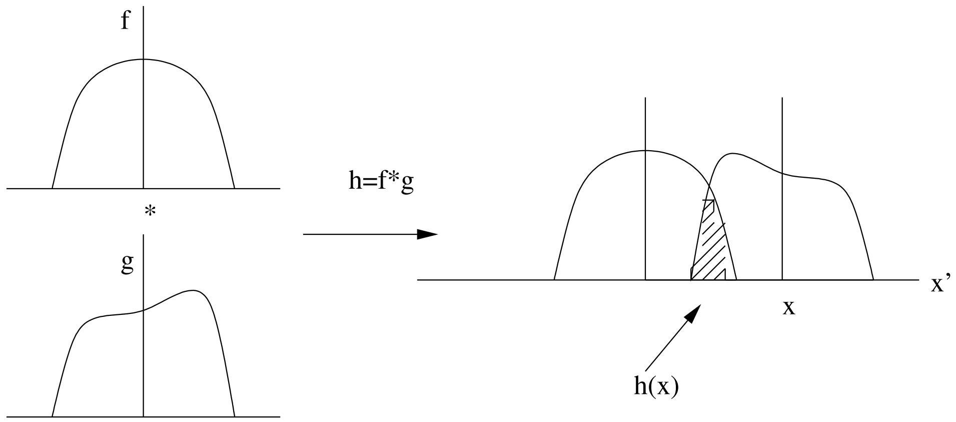

Figure 6.6: Illustration of the convolution of two functions, viewed as the area of the overlap resulting from a relative shift of x.

FOURIER ANALYSIS: LECTURE 9

6 Convolution

Convolution combines two (or more) functions in a way that is useful for describing physical systems. Convolution describes, for example, how optical systems respond to an image: it gives a mathematical description of the process of blurring. We will also see how Fourier solutions to differential equations can often be expressed as a convolution. The FT of the convolution is easy to calculate, so Fourier methods are ideally suited for solving problems that involve convolution.

First, the definition. The convolution of two functions $f(x)$ and $g(x)$ is defined to be

$$f(x) * g(x) = \int_{-\infty}^{\infty} dx'\, f(x')g(x-x')\ , \qquad (6.59)$$

The result is also a function of x, meaning that we get a different number for the convolution for each possible x value. Note the positions of the dummy variable x', especially that the argument of g is $x - x'$ and not $x' - x$ (a common mistake in exams).

There are a number of ways of viewing the process of convolution. Most directly, the definition here is a measure of *overlap*: the functions f and g are shifted relative to one another by a distance x, and we integrate to find the product. This viewpoint is illustrated in Fig. 6.6.

But this is not the best way of thinking about convolution. The real significance of the operation is that it represents a *blurring* of a function. Here, it may be helpful to think of $f(x)$ as a signal, and $g(x)$ as a blurring function. As written, the integral definition of convolution instructs us to take the signal at x', $f(x')$, and replace it by something proportional to $f(x')g(x - x')$: i.e. spread out over a range of x around x'. This turns a sharp feature in the signal into something fuzzy centred at the same location. This is exactly what is achieved e.g. by an out-of-focus camera.

Alternatively, we can think about convolution as a form of averaging. Take the above definition of convolution and put $y = x - x'$. Inside the integral, x is constant, so $dy = -dx'$. But now we are integrating from $y = \infty$ to $-\infty$, so we can lose the minus sign by re-inverting the limits:

$$f(x) * g(x) = \int_{-\infty}^{\infty} dy\, f(x-y)g(y)\ . \qquad (6.60)$$

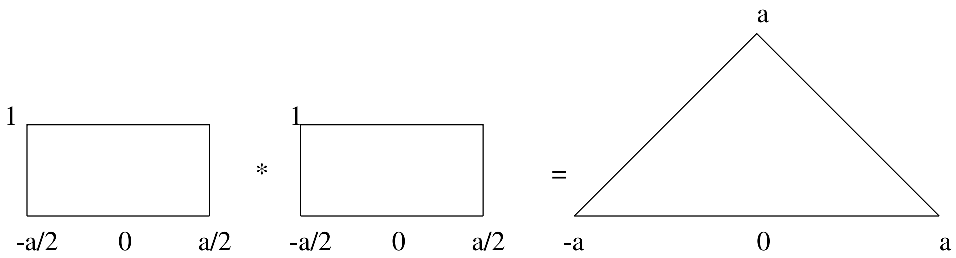

Figure 6.7: Convolution of two top hat functions.

This says that we replace the value of the signal at x, $f(x)$ by an average of all the values around x, displaced from x by an amount y and weighted by the function $g(y)$. This is an equivalent view of the process of blurring. Since it doesn't matter what we call the dummy integration variable, this rewriting of the integral shows that convolution is commutative: you can think of g blurring f or f blurring g:

$$f(x) * g(x) = \int_{-\infty}^{\infty} dz\ f(z)g(x-z) = \int_{-\infty}^{\infty} dz\ f(x-z)g(z) = g(x) * f(x). \tag{6.61}$$

6.1 Examples of convolution

1. Let $\Pi(x)$ be the top-hat function of width a.
   - $\Pi(x) * \Pi(x)$ is the triangular function of base width $2a$.
   - This is much easier to do by sketching than by working it out formally: see Figure 6.7.
2. Convolution of a general function $g(x)$ with a delta function $\delta(x-a)$.

   $$\delta(x-a) * g(x) = \int_{-\infty}^{\infty} dx'\ \delta(x'-a)g(x-x') = g(x-a). \tag{6.62}$$

   using the sifting property of the delta function. This is a clear example of the blurring effect of convolution: starting with a spike at $x = a$, we end up with a copy of the whole function $g(x)$, but now shifted to be centred around $x = a$. So here the 'sifting' property of a delta-function has become a 'shifting' property. Alternatively, we may speak of the delta-function becoming 'dressed' by a copy of the function g.

   The response of the system to a delta function input (i.e. the function $g(x)$ here) is sometimes called the *Impulse Response Function* or, in an optical system, the *Point Spread Function.*
3. Making double slits: to form double slits of width a separated by distance $2d$ between centres:

   $$[\delta(x+d) + \delta(x-d)] * \Pi(x)\ . \tag{6.63}$$

   We can form diffraction gratings with more slits by adding in more delta functions.

6.2 The convolution theorem

States that the Fourier transform of a *convolution* is a *product* of the individual Fourier transforms:

$$FT[f(x) * g(x)] = \tilde{f}(k)\, \tilde{g}(k) \tag{6.64}$$

$$FT[f(x)\, g(x)] = \frac{1}{2\pi} \tilde{f}(k) * \tilde{g}(k) \tag{6.65}$$

where $\tilde{f}(k)$, $\tilde{g}(k)$ are the FTs of $f(x)$, $g(x)$ respectively.

Note that:

$$\tilde{f}(k) * \tilde{g}(k) \equiv \int_{-\infty}^{\infty} dq\; \tilde{f}(q)\; \tilde{g}(k-q)\ . \tag{6.66}$$

We'll do one of these, and we will use the Dirac delta function.

The convolution $h = f * g$ is

$$h(x) = \int_{-\infty}^{\infty} f(x')g(x-x')\, dx'. \tag{6.67}$$

We substitute for $f(x')$ and $g(x-x')$ their FTs, noting the argument of g is not x':

$$\begin{aligned} f(x') &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \tilde{f}(k) e^{ikx'}\, dk \\ g(x-x') &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \tilde{g}(k) e^{ik(x-x')}\, dk \end{aligned}$$

Hence (relabelling the k to k' in g, so we don't have two k integrals)

$$h(x) = \frac{1}{(2\pi)^2} \int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} \tilde{f}(k) e^{ikx'}\, dk \int_{-\infty}^{\infty} \tilde{g}(k') e^{ik'(x-x')}\, dk' \right) dx'. \tag{6.68}$$

Now, as is very common with these multiple integrals, we do the integrations in a different order. Notice that the only terms which depend on x' are the two exponentials, indeed only part of the second one. We do this one first, using the fact that the integral gives 2π times a Dirac delta function:

$$\begin{aligned} h(x) &= \frac{1}{(2\pi)^2} \int_{-\infty}^{\infty} \tilde{f}(k) \int_{-\infty}^{\infty} \tilde{g}(k') e^{ik'x} \left(\int_{-\infty}^{\infty} e^{i(k-k')x'} dx' \right) dk' dk \\ &= \frac{1}{(2\pi)^2} \int_{-\infty}^{\infty} \tilde{f}(k) \int_{-\infty}^{\infty} \tilde{g}(k') e^{ik'x} \left[2\pi\delta(k-k') \right] dk' dk \end{aligned}$$

Having a delta function simplifies the integration enormously. We can do either the k or the k' integration immediately (it doesn't matter which you do – let us do k'):

$$\begin{aligned} h(x) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \tilde{f}(k) \left[\int_{-\infty}^{\infty} \tilde{g}(k') e^{ik'x} \delta(k-k')\, dk' \right] dk \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} \tilde{f}(k) \tilde{g}(k)\, e^{ikx}\, dk \end{aligned}$$

Since

$$h(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \tilde{h}(k)\, e^{ikx}\, dk \tag{6.69}$$

we see that

$$\tilde{h}(k) = \tilde{f}(k)\tilde{g}(k). \tag{6.70}$$

Note that we can apply the convolution theorem in reverse, going from Fourier space to real space, so we get the most important key result to remember about the convolution theorem:

$$\begin{aligned} \text{Convolution in real space} &\Leftrightarrow \text{Multiplication in Fourier space} \\ \text{Multiplication in real space} &\Leftrightarrow \text{Convolution in Fourier space} \end{aligned} \tag{6.71}$$

This is an important result. Note that if one has a convolution to do, it is often most efficient to do it with Fourier Transforms, not least because a very efficient way of doing them on computers exists – the *Fast Fourier Transform*, or FFT.

CONVENTION ALERT! Note that if we had chosen a different convention for the 2π factors in the original definitions of the FTs, the convolution theorem would look differently. Make sure you use the right one for the conventions you are using!

Note that convolution commutes, $f(x) * g(x) = g(x) * f(x)$, which is easily seen (e.g. since the FT is $\tilde{f}(k)\tilde{g}(k) = \tilde{g}(k)\tilde{f}(k)$.)

Example application: Fourier transform of the triangular function of base width $2a$. We know that a triangle is a convolution of top hats:

$$\Delta(x) = \Pi(x) * \Pi(x) \ . \tag{6.72}$$

Hence by the convolution theorem:

$$FT[\Delta] = (FT[\Pi(x)])^2 = \left(\text{sinc}\,\frac{ka}{2}\right)^2 \tag{6.73}$$

FOURIER ANALYSIS: LECTURE 10

7 Parseval's theorem for FTs (Plancherel's theorem)

For FTs, there is a similar relationship between the average of the square of the function and the FT coefficients as there is with Fourier Series. For FTs it is strictly called *Plancherel's theorem*, but is often called the same as FS, i.e. Parseval's theorem; we will stick with Parseval. The theorem says

$$\int_{-\infty}^{\infty} |f(x)|^2 \, dx = \frac{1}{2\pi} \int_{-\infty}^{\infty} |\tilde{f}(k)|^2 \, dk. \tag{7.74}$$

It is useful to compare different ways of proving this:

(1) The first is to go back to Fourier series for a periodic $f(x)$: $f(x) = \sum_n c_n \exp(ik_n x)$, and $|f|^2$ requires us to multiply the series by itself, which gives lots of cross terms. But when we integrate over one fundamental period, all oscillating terms average to zero. Therefore the only terms that survive are ones where $c_n \exp(ik_n x)$ pairs with $c_n^* \exp(-ik_n x)$. This gives us Parseval's theorem for Fourier series:

$$\frac{1}{\ell} \int_{-\ell/2}^{\ell/2} |f(x)|^2 \, dx = \sum_n |c_n|^2 \Rightarrow \int_{-\ell/2}^{\ell/2} |f(x)|^2 \, dx = \ell \sum_n |c_n|^2 = \frac{1}{\ell} \sum_n |\tilde{f}|^2, \tag{7.75}$$

using the definition $\tilde{f} = \ell c_n$. But the mode spacing is $dk = 2\pi/\ell$, so $1/\ell$ is $dk/2\pi$. Now we take the continuum limit of $\ell \to \infty$ and $dk \sum$ becomes $\int dk$.

(2) Alternatively, we can give a direct proof using delta-functions:

$$|f(x)|^2 = f(x) f^*(x) = \left(\frac{1}{2\pi} \int \tilde{f}(k) \exp(ikx) \, dk \right) \times \left(\frac{1}{2\pi} \int \tilde{f}^*(k') \exp(-ik'x) \, dk' \right), \tag{7.76}$$

which is

$$\frac{1}{(2\pi)^2} \iint \tilde{f}(k) \tilde{f}^*(k') \exp[ix(k - k')] \, dk \, dk'. \tag{7.77}$$

If we now integrate over x, we generate a delta-function:

$$\int \exp[ix(k - k')] \, dx = (2\pi) \delta(k - k'). \tag{7.78}$$

So

$$\int |f(x)|^2 \, dx = \frac{1}{2\pi} \iint \tilde{f}(k) \tilde{f}^*(k') \, \delta(k - k') \, dk \, dk' = \frac{1}{2\pi} \int |\tilde{f}(k)|^2 \, dk. \tag{7.79}$$

7.1 Energy spectrum of decaying signal

As in the case of Fourier series, $|\tilde{f}(k)|^2$ is often called the *Power Spectrum* of the signal. If we have a field (such as an electric field) where the energy density is proportional to the *square* of the field, then we can interpret the square of the Fourier Transform coefficients as the *energy* associated with each frequency – i.e. total energy radiated is

$$\int_{-\infty}^{\infty} |f(t)|^2 \, dt. \tag{7.80}$$

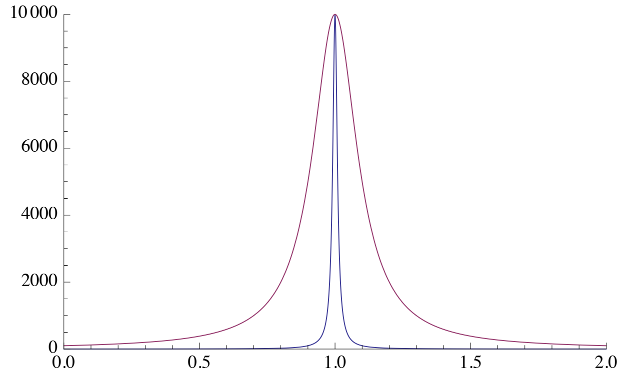

Figure 7.8: Frequency spectrum of two separate exponentially decaying systems with 2 different time constants τ. (x axis is frequency, y axis $\propto |\tilde{f}(\omega)|^2$ in arbitrary units).

By Parseval's theorem, this is equal to

$$\frac{1}{2\pi}\int_{-\infty}^{\infty} |\tilde{f}(\omega)|^2\, d\omega. \tag{7.81}$$

and we interpret $|\tilde{f}(\omega)|^2/(2\pi)$ as the energy radiated per unit (angular) frequency, at frequency ω.

If we have a quantum transition from an upper state to a lower state, which happens spontaneously, then the intensity of emission will decay exponentially. We can model this semi-classically as a field that oscillates with frequency ω_0, but with an amplitude that is damped with a timescale $\tau = 1/a$:

$$f(t) = e^{-at}\cos(\omega_0 t) \qquad (t > 0). \tag{7.82}$$

Algebraically it is easier to write this as the real part of a complex exponential, do the FT with the exponential, and take the real part at the end. So consider

$$f(t) = \frac{1}{2}e^{-at}(e^{i\omega_0 t} + e^{-i\omega_0 t}) \qquad (t > 0). \tag{7.83}$$

The Fourier transform is [1]

$$\begin{aligned} \tilde{f}(\omega) &= \frac{1}{2}\int_0^{\infty} (e^{-at-i\omega t+i\omega_0 t} + e^{-at-i\omega t-i\omega_0 t})\, dt & (7.84)\\ \Rightarrow 2\tilde{f}(\omega) &= \left[\frac{e^{-at-i\omega t+i\omega_0 t}}{-a-i\omega+i\omega_0} - \frac{e^{-at-i\omega t-i\omega_0 t}}{-a-i\omega-i\omega_0}\right]_0^{\infty} \\ &= \frac{1}{(a+i\omega-i\omega_0)} + \frac{1}{(a+i\omega+i\omega_0)} \end{aligned}$$

(7.85)

This is sharply peaked near $\omega = \omega_0$; near this frequency, we therefore ignore the second term, and the frequency spectrum is

$$|\tilde{f}(\omega)|^2 \simeq \frac{1}{4\,[a+i(\omega-\omega_0)]}\frac{1}{[a-i(\omega-\omega_0)]} = \frac{1}{4\,[a^2+(\omega-\omega_0)^2]}. \tag{7.86}$$

[1] Note that this integral is *similar* to one which leads to Delta functions, but it isn't, because of the e^{-at} term. For this reason, you can integrate it by normal methods. If $a = 0$, then the integral does indeed lead to Delta functions.

This is a Lorentzian curve with width $a = 1/\tau$. Note that the width of the line in frequency is inversely proportional to the decay timescale τ. This is an example of the Uncertainty Principle, and relates the *natural width* of a spectral line to the decay rate. See Fig. 7.8.

8 Correlations and cross-correlations

Correlations are defined in a similar way to convolutions, but look carefully, as they are slightly different. With correlations, we are concerned with how similar functions are when one is displaced by a certain amount. If the functions are different, the quantity is called the *cross-correlation*; if it is the same function, it is called the *auto-correlation*, or simply *correlation*.

The cross-correlation of two functions is defined by

$$c(X) \equiv \langle f^*(x) g(x+X) \rangle \equiv \int_{-\infty}^{\infty} f^*(x) g(x+X)\, dx. \qquad (8.87)$$

Compare this with convolution (equation 6.59). X is sometimes called the *lag*. Note that cross-correlation does not commute, unlike convolution. The most interesting special case is when f and g are the same function: then we have the *auto-correlation function*.

The meaning of these functions is easy to visualise if the functions are real: at zero lag, the auto-correlation function is then proportional to the variance in the function (it would be equal if we divided the integral by a length ℓ, where the functions are zero outside that range). So then the *correlation coefficient* of the function is

$$r(X) = \frac{\langle f(x) f(x+X) \rangle}{\langle f^2 \rangle}. \qquad (8.88)$$

If r is small, then the values of f at widely separated points are unrelated to each other: the point at which r falls to 1/2 defines a characteristic width of a function. This concept is used particularly in random processes.

The FT of a cross-correlation is

$$\tilde{c}(k) = \tilde{f}^*(k)\, \tilde{g}(k). \qquad (8.89)$$

This looks rather similar to the convolution theorem, which is is hardly surprising given the similarity of the definitions of cross-correlation and convolution. Indeed, the result can be proved directly from the convolution theorem, by writing the cross-correlation as a convolution.

A final consequence of this is that the FT of an auto-correlation is just the power spectrum; or, to give the inverse relation:

$$\langle f^*(x) f(x+X) \rangle = \frac{1}{2\pi} \int |\tilde{f}|^2 \exp(ikX)\, dk. \qquad (8.90)$$

This is known as the *Wiener-Khinchin theorem*, and it generalises Parseval's theorem (to which it reduces when $X = 0$). It is straightforward to prove directly, by writing the Fourier integral for f twice and using a delta-function; we will do this in the workshops.

Finally, note that much of this discussion applies also to periodic functions defined as Fourier series, where the proof is even easier.

$$f(x) = \sum_n c_n \exp(ik_n x) \Rightarrow f^*(x) f(x+X) = \sum_{n,m} c_n^* c_m \exp[i(k_m - k_n)x] \exp(ik_m X). \qquad (8.91)$$

If we now interpret the averaging $\langle \dots \rangle$ as integrating in x over one period and dividing by the period, the $\exp[i(k_m - k_n)x]$ term yields just δ_{mn}. Hence

$$\langle f^*(x) f(x+X) \rangle = \sum_n |c_n|^2 \exp(ik_n X). \tag{8.92}$$

9 Fourier analysis in multiple dimensions

We have now completed all the major tools of Fourier analysis, in one spatial dimension. In many cases, we want to consider more than one dimension, and the extension is relatively straightforward. Start with the fundamental Fourier series, $f(x) = \sum_n c_n \exp(i2\pi nx/\ell_x)$. $f(x)$ can be thought of as $F(x, y)$ at constant y; if we change y, the effective $f(x)$ changes, so the c_n must depend on y. Hence we can Fourier expand these as a series in y:

$$c_n(y) = \sum_m d_{nm} \exp(i2\pi my/\ell_y), \tag{9.93}$$

where we assume that the function is periodic in x, with period ℓ_x, and y, with period ℓ_y. The overall series is than

$$F(x, y) = \sum_{n,m} d_{nm} \exp[2\pi i(nx/\ell_x + my/\ell_y)] = \sum_{n,m} d_{nm} \exp[i(k_x x + k_y y)] = \sum_{n,m} d_{nm} \exp[i(\boldsymbol{k} \cdot \boldsymbol{x})]. \tag{9.94}$$

This is really just the same as the 1D form, and the extension to D dimensions should be obvious. In the end, we just replace the usual kx term with the dot product between the position vector and the wave vector.

The Fourier transform in D dimensions just involves taking the limit of $\ell_x \to \infty$, $\ell_y \to \infty$ etc. The Fourier coefficients become a continuous function of $\boldsymbol{k}$, in which case we can sum over bins in k space, each containing $N_{\rm modes}(\boldsymbol{k})$ modes:

$$F(\boldsymbol{x}) = \sum_{\rm bin} d(\boldsymbol{k}) \exp[i(\boldsymbol{k} \cdot \boldsymbol{x})] \, N_{\rm modes}. \tag{9.95}$$

The number of modes in a given k-space bin is set by the period in each direction: allowed modes lie on a grid of points in the space of k_x, k_y etc. as shown in Figure 9.9. If for simplicity the period is the same in all directions, the *density of states* is $\ell^D/(2\pi)^D$:

$$N_{\rm modes} = \frac{\ell^D}{(2\pi)^D} \, d^D k. \tag{9.96}$$

This is an important concept which is used in many areas of physics.

The Fourier expression of a function is therefore

$$F(\boldsymbol{x}) = \frac{1}{(2\pi)^D} \int \tilde{F}(\boldsymbol{k}) \exp[i(\boldsymbol{k} \cdot \boldsymbol{x}) \, d^D k], \tag{9.97}$$

Where we have defined $\tilde{F}(\boldsymbol{k}) \equiv \ell^D d(\boldsymbol{k})$. The inverse relation would be obtained as in 1D, by appealing to orthogonality of the modes:

$$\tilde{F}(\boldsymbol{k}) = \int F(\boldsymbol{x}) \exp[-i(\boldsymbol{k} \cdot \boldsymbol{x})] \, d^D x. \tag{9.98}$$

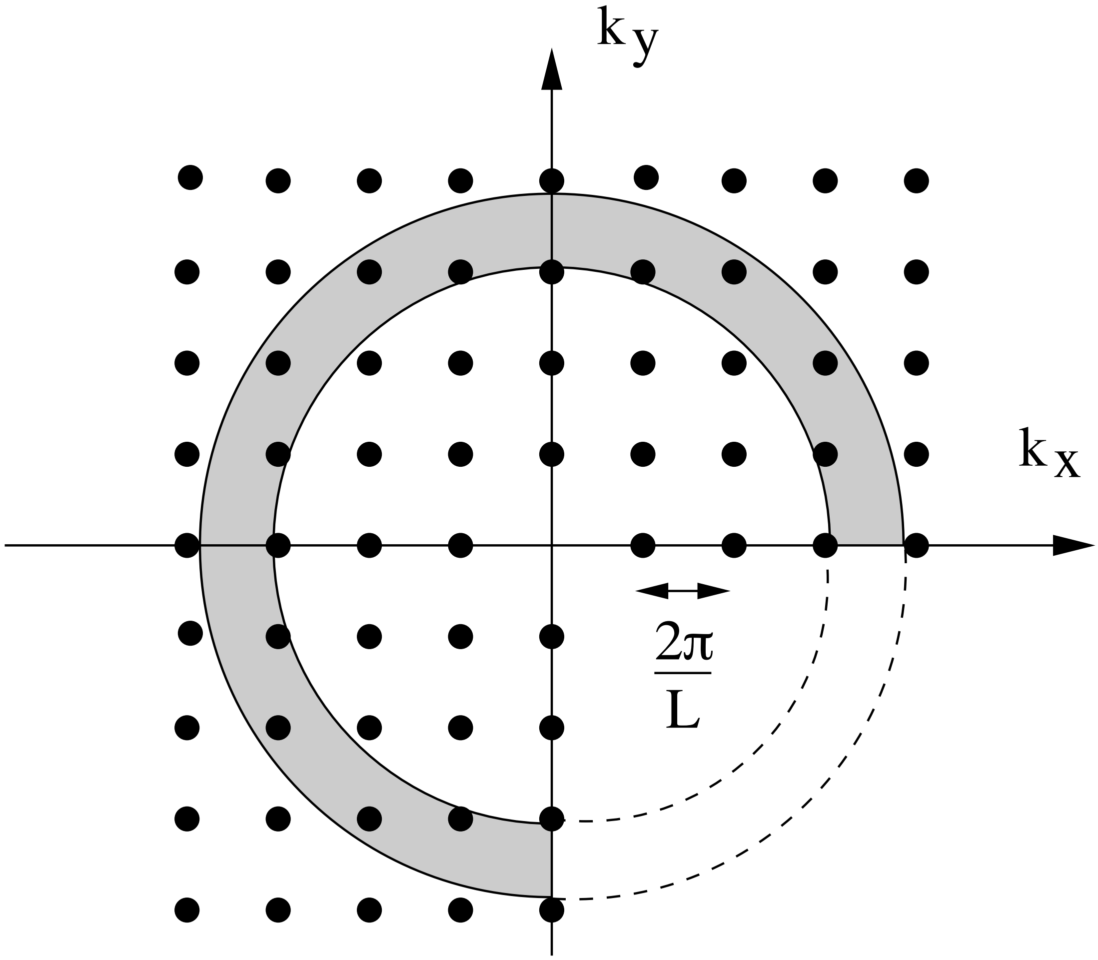

Figure 9.9: Illustrating the origin of the density of states in 2D. The allowed modes are shown as points, with a separation in k_x and k_y of $2\pi/\ell$, where ℓ is the periodicity. The number of modes between $|k|$ and $|k| + d|k|$ (i.e. inside the shaded annulus) is well approximated by $(\ell/2\pi)^2$ times the area of the annulus, as $\ell \to \infty$, and the mode spacing tends to zero. Clearly, in D dimensions, the mode density is just $(\ell/2\pi)^D$.

FOURIER ANALYSIS: LECTURE 11

10 Digital analysis and sampling

Imagine we have a continuous signal (e.g. pressure of air during music) which we sample by making measurements at a few particular times. Any practical storage of information must involve this step of *analogue-to-digital conversion*. This means we are converting a continuous function into one that is only known at discrete points – i.e. we are throwing away information. We would feel a lot more comfortable doing this if we knew that the missing information can be recovered, by some form of interpolation between the sampled points. Intuitively, this seems reasonable if the sampling interval is very fine: by the definition of continuity, the function between two sampled points should be arbitrarily close to the average of the sample values as the locations of the samples gets closer together. But the sampling interval has to be finite, so this raises the question of how coarse it can be; clearly we would prefer to use as few samples as possible consistent with not losing any information. This question does have a well-posed answer, which we can derive using Fourier methods.

The first issue is how to represent the process of converting a function $f(x)$ into a set of values $\{f(x_i)\}$. We can do this by using some delta functions:

$$f(x) \rightarrow f_s(x) \equiv f(x) \sum_i \delta(x - x_i). \tag{10.99}$$

This replaces our function by a sum of spikes at the locations x_i, each with a weight $f(x_i)$. This representation of the sampled function holds the information of the sample values and locations. So, for example, if we try to average the sampled function over some range, we automatically get something proportional to just adding up the sample values that lie in the range:

$$\int_{x_1}^{x_2} f_s(x)\, dx = \sum_{\text{in range}} f(x_i). \tag{10.100}$$

10.1 The infinite comb

If we sample regularly with a spacing Δx, then we have an 'infinite comb' – an infinite series of delta functions. The comb is (see Fig. 10.10):

$$g(x) = \sum_{j=-\infty}^{\infty} \delta(x - j\Delta x) \tag{10.101}$$

This is also known as the *Shah function.*

To compute the FT of the Shah function, we will write it in another way. This is derived from the fact that the function is periodic, and therefore suitable to be written as a Fourier *series* with $\ell = \Delta x$:

$$g(x) = \sum_n c_n \exp(2\pi n i x/\Delta x). \tag{10.102}$$

The coefficients c_n are just

$$c_n = \frac{1}{\Delta x} \int_{-\Delta x/2}^{\Delta x/2} \delta(x)\, dx = \frac{1}{\Delta x}, \tag{10.103}$$

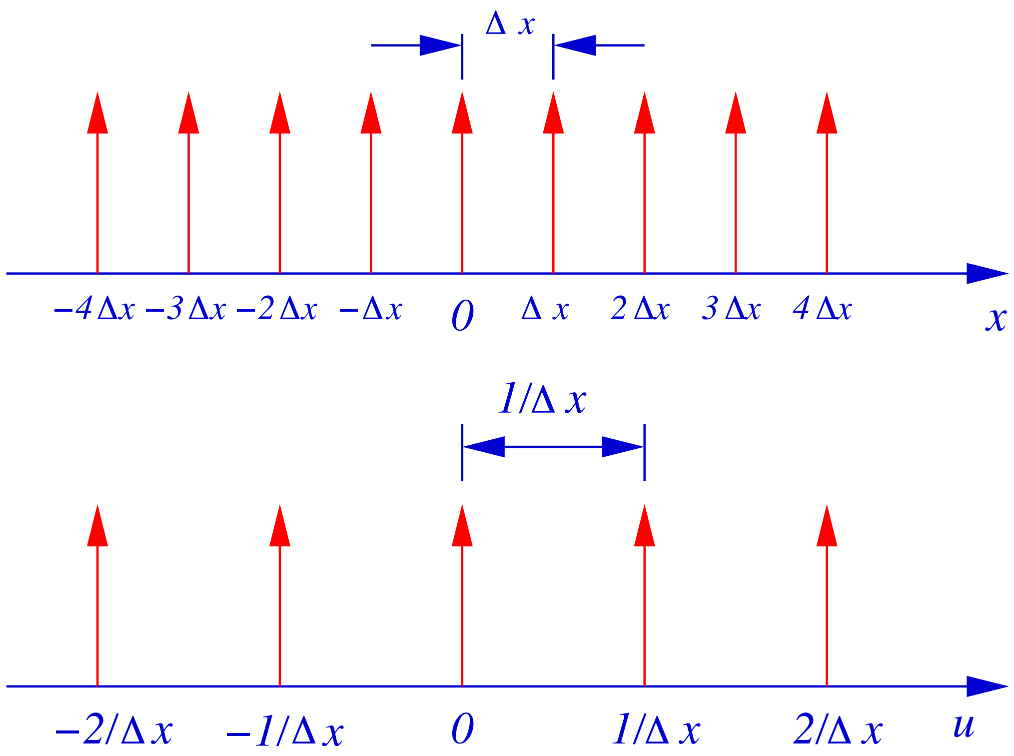

Figure 10.10: Top: An infinite comb in real space. This represents the sampling pattern of a function which is sampled regularly every Δx. Bottom: The FT of the infinite comb, which is also an infinite comb. Note that u here is $k/(2\pi)$.

so that

$$g(x) = \frac{1}{\Delta x}\sum_n \exp(2\pi n i x/\Delta x) = \frac{1}{2\pi}\int \tilde{g}(k)\exp(ikx)\,dx. \tag{10.104}$$

From this, we can see that $\tilde{g}(k)$ must involve a sum of delta-functions in k space, since $g(x)$ has ended up as a sum of $\exp(ik_n x)$ terms, each of which could be sifted out of the Fourier integral by a contribution to $\tilde{g}(k)$ that is $\propto \delta(k - k_n)$. More formally, we could take the FT of our new expression for $g(x)$, which would yield a series of delta-functions. In any case,

$$\tilde{g}(k) = \frac{2\pi}{\Delta x}\sum_{n=-\infty}^{\infty}\delta(k - 2\pi n/\Delta x). \tag{10.105}$$

which is an infinite comb in Fourier space, with spacing $2\pi/\Delta x$.

The FT of a function sampled with an infinite comb is therefore ($1/2\pi$ times) the convolution of this and the FT of the function:

$$\tilde{f}_s(k) = \frac{1}{2\pi}\tilde{f}(k) * \tilde{g}(k) = \frac{1}{\Delta x}\sum_{n=-\infty}^{\infty}\tilde{f}(k - 2\pi n/\Delta x). \tag{10.106}$$

In other words, each delta-function in the k-space comb becomes 'dressed' with a copy of the transform of the original function.

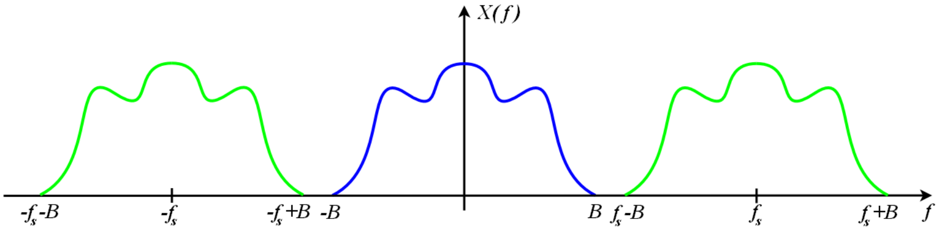

Figure 10.11: If the sampling is fine enough, then the original spectrum can be recovered from the sampled spectrum.

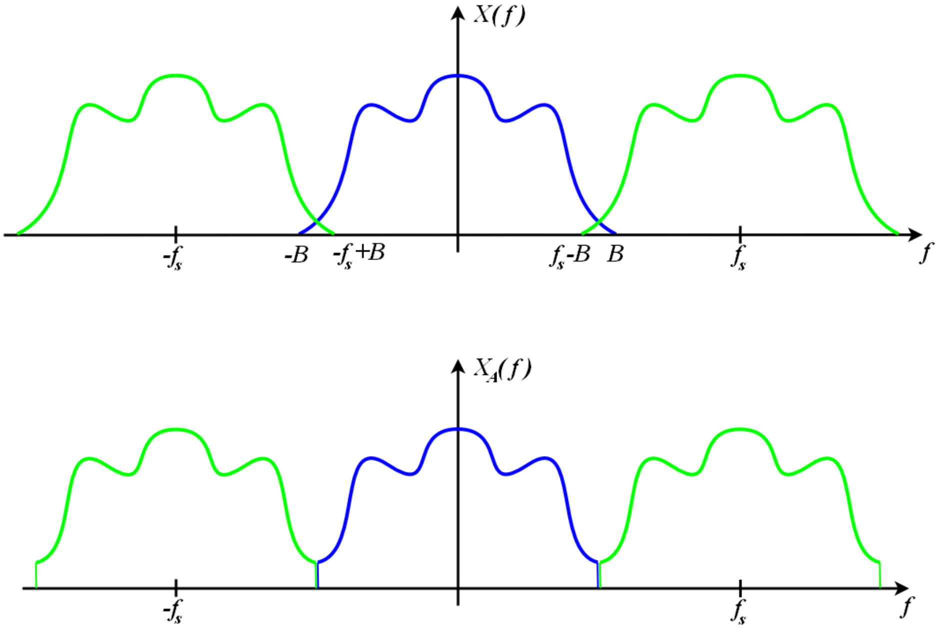

Figure 10.12: If the sampling is *not* fine enough, then the power at different frequencies gets mixed up, and the original spectrum cannot be recovered.

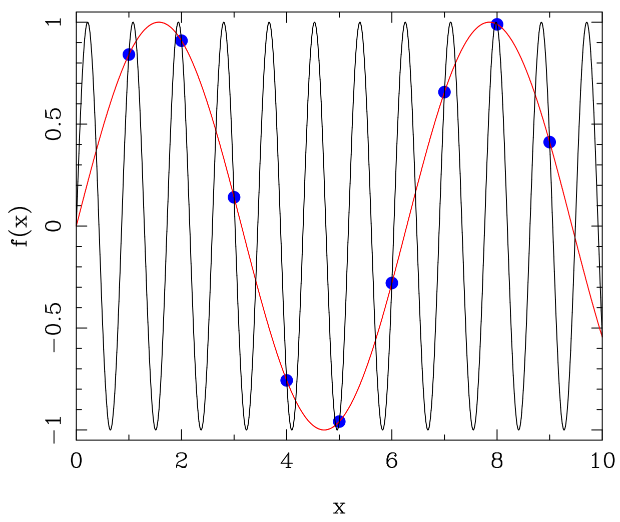

Figure 10.13: If $\sin t$ is sampled at unit values of t, then $\sin(t + 2\pi t)$ is indistinguishable at the sampling points. The sampling theorem says we can only reconstruct the function between the samples if we know that high-frequency components are absent.

10.2 Shannon sampling, aliasing and the Nyquist frequency

We can now go back to the original question: do the sampled values allow us to reconstruct the original function exactly? An equivalent question is whether the transform of the sampled function allows us to reconstruct the transform of the original function.

The answer is that this is possible (a) if the original spectrum is *bandlimited*, which means that the power is confined to a finite range of wavenumber (i.e. there is a maximum wavenumber k_{max} which has non-zero Fourier coefficients); and (b) if the sampling is fine enough. This is illustrated in Figs 10.11 and 10.12. If the sampling is not frequent enough, the power at different wavenumbers gets mixed up. This is called *aliasing*. The condition to be able to measure the spectrum accurately is to have a sample at least as often as the *Shannon Rate*

$$\Delta x = \frac{1}{\pi k_{\rm max}}. \qquad (10.107)$$

The Nyquist wavenumber is defined as

$$k_{\rm Nyquist} = \frac{\pi}{\Delta x} \qquad (10.108)$$

which needs to exceed the maximum wavenumber in order to avoid aliasing:

$$k_{\rm Nyquist} \geq k_{\rm max}. \qquad (10.109)$$

For time-sampled data (such as sound), the same applies, with wavenumber k replaced by frequency ω.

There is a simple way of seeing that this makes sense, as illustrated in Figure 10.13. Given samples of a Fourier mode at a certain interval, Δx, a mode with a frequency increased by any multiple of $2\pi/\Delta x$ clearly has the same result at the sample points.

10.2.1 Interpolation of samples

The idea of having data that satisfy the sampling theorem is that we should be able to reconstruct the full function from the sampled values; how do we do this in practice? If the sampled function is the product of $f(x)$ and the Shah function, we have seen that the FT of the sampled function is the same as $\tilde{f}/\Delta x$, for $|k| < \pi/\Delta x$. If we now multiply by $\tilde{T}(k)$: a top-hat in k space, extending from $-\pi/\Delta x$ to $+\pi/\Delta x$, with height Δx, then we have exactly $\tilde{f}$ and can recover $f(x)$ by an inverse Fourier transform. This k-space multiplication amounts to convolving the sampled data with the inverse Fourier transform of $T(k)$, so we recover $f(x)$ in the form

$$f(x) = [f(x)g(x)]*T(x) = \int f(q)\sum_n \delta(q-n\Delta x)\,T(x-q)\;dq = \int \sum_n f(n\Delta x)\delta(q-n\Delta x)\,T(x-q)\;dq, \tag{10.110}$$

using $f(x)\delta(x-a) = f(a)\delta(x-a)$. The sum of delta-functions sifts to give

$$f(x) = \sum_n f(n\Delta x)T(x - n\Delta x), \tag{10.111}$$

i.e. the function $T(x) = \sin[\pi x/\Delta x]/(\pi x/\Delta x)$ is the interpolating function. This is known as 'sinc interpolation'.

FOURIER ANALYSIS: LECTURE 12

10.3 CDs and compression

Most human beings can hear frequencies in the range 20 Hz – 20 kHz. The sampling theorem means that the sampling frequency needs to be at least 40 kHz to capture the 20 kHz frequencies. The CD standard samples at 44.1 kHz. The data consist of stereo: two channels each encoded as 16-bit integers. Allowing one bit for sign, the largest number encoded is thus $2^{15} - 1 = 32767$. This allows signals of typical volume to be encoded with a fractional precision of around 0.01% – an undetectable level of distortion. This means that an hour of music uses about 700MB of information. But in practice, this requirement can be reduced by about a factor 10 without noticeable degradation in quality. The simplest approach would be to reduce the sampling rate, or to encode the signal with fewer bits. The former would require a reduction in the maximum frequency, making the music sound dull; but fewer bits would introduce distortion from the quantization of the signal. The solution implemented in the MP3 and similar algorithms is more sophisticated than this: the time series is split into 'frames' of 1152 samples (0.026 seconds at CD rates) and each is Fourier transformed. Compression is achieved by storing simply the amplitudes and phases of the strongest modes, as well as using fewer bits to encode the amplitudes of the weaker modes, according to a 'perceptual encoding' where the operation of the human ear is exploited – knowing how easily faint tones of a given frequency are masked by a loud one at a different frequency.

10.4 Prefiltering

If a signal does not obey the sampling theorem, it must be modified to do so before digitization. Analogue electronics can suppress high frequencies – although they are not completely removed. The sampling process itself almost inevitably performs this task to an extent, since it is unrealistic to imagine that one could make an instantaneous sample of a waveform. Rather, the sampled signal is probably an average of the true signal over some period.

This is easily analysed using the convolution theorem. Suppose each sample, taken at an interval τ, is the average of the signal over a time interval T, centred at the sample time. This is a convolution:

$$f_c(t) = \int f(t')g(t - t')\, dt', \tag{10.112}$$

where $g(t - t')$ is a top hat of width T centred on $t' = t$. We therefore know that

$$\tilde{f}_c(\omega) = \tilde{f}(\omega) \sin(\omega T/2)/(\omega T/2). \tag{10.113}$$

At the Nyquist frequency, π/τ, the Fourier signal in f is suppressed by a factor $\sin(\pi T/2\tau)/(\pi T/2\tau)$. The natural choice of T would be the same as τ (accumulate an average signal, store it, and start again). This gives $\sin(\pi/2)/(\pi/2) = 0.64$ at the Nyquist frequency, so aliasing is not strongly eliminated purely by 'binning' the data, and further prefiltering is required before the data can be sampled.

11 Discrete Fourier Transforms & the FFT

This section is added to the course notes as a non-examinable supplement, which may be of interest to those using numerical Fourier methods in project work. We have explored the properties of sampled data using the concept of an infinite array of delta functions, but this is not yet a practical form that can be implemented on a computer.

11.1 The DFT

Suppose that we have a function, $f(x)$, that is periodic with period ℓ, and which is known only at N equally spaced values $x_n = n(\ell/N)$. Suppose also that $f(x)$ is band-limited with a maximum wavenumber that satisfies $|k_{\rm max}| < \pi/(\ell/N)$, i.e. it obeys the sampling theorem. If we wanted to describe this function via a Fourier series, we would need the Fourier coefficients

$$f_k(k) = \frac{1}{\ell}\int_0^\ell f(x)\,\exp[-ikx]\,dx. \tag{11.114}$$

This integral can clearly be approximated by summing over the N sampled values:

$$f_k(k) = \frac{1}{\ell}\sum_n f(x_n)\,\exp[-ikx_n]\,\ell/N = \frac{1}{N}\sum_n f(x_n)\,\exp[-ikx_n]; \tag{11.115}$$

in fact, we show below that this expression yields the exact integral for data that obey the sampling theorem. The range of grid values is irrelevant because of the periodicity of f. Suppose we sum from $n = 1$ to N and then change to $n = 0$ to $N - 1$: the sum changes by $f(x_0)\exp[-ikx_0] - f(x_N)\exp[-ikx_N]$, but $f(x_0) = f(x_N)$ and $x_N - x_0 = \ell$. Since the allowed values of k are multiples of $2\pi/\ell$, the change in the sum vanishes. We can therefore write what can be regarded as the definition of the *discrete Fourier transform* of the data:

$$f_k(k_m) = \frac{1}{N}\sum_{n=0}^{N-1} f(x_n)\,\exp[-ik_m x_n], \tag{11.116}$$

where the allowed values of k are $k_m = m(2\pi/\ell)$ and the allowed values of x are $x_n = n(\ell/N)$. This expression has an inverse of very similar form:

$$f(x_j) = \sum_{m=0}^{N-1} f_k(k_m)\,\exp[ik_m x_j]. \tag{11.117}$$

To prove this, insert the first definition in the second, bearing in mind that $k_m x_n = 2\pi mn/N$. This gives the expression

$$\frac{1}{N}\sum_{m,n} f(x_n)\,\exp[2\pi i m(j-n)/N] = \frac{1}{N}\sum_{m,n} f(x_n)\,z^m, \tag{11.118}$$

where $z = \exp[2\pi i(j-n)/N]$. Consider $\sum_m z^m$: where $j = n$ we have $z = 1$ and the sum is N. But if $j \neq n$, the sum is zero. To show this, consider $z\sum_m z^m = \sum_m z^m + z^N - 1$. But $z_N = 1$, and we have $z\sum_m z^m = \sum_m z^m$, requiring the sum to vanish if $z \neq 1$. Hence $\sum_m z^m = N\delta_{jn}$, and this orthogonality relation proves that the inverse is exact.

An interesting aspect of the inverse as written is that it runs only over positive wavenumbers; don't we need k to be both positive and negative? The answer comes from the gridding in x and k: since $k_m x_n = 2\pi mn/N$, letting $m \to n + N$ has no effect. Thus a mode with $m = N - 1$ is equivalent to one with $m = -1$ etc. So the k-space array stores increasing positive wavenumbers in its first elements, jumping to large negative wavenumbers once k exceeds the Nyquist frequency. In a little more detail, the situation depends on whether N is odd or even. If it is odd, then $m = (N+1)/2$ is an integer equivalent to $-(N-1)/2$, so the successive elements $m = (N-1)/2$ and $m = (N+1)/2$ hold the symmetric information near to the Nyquist frequency, $|k| = (N-1)/N \times \pi/(\ell/N)$. On the other hand, if N is even, we have a single mode exactly at the Nyquist frequency: $m = N/2 \Rightarrow |k| = (N/2)(2\pi/\ell) = \pi/(\ell/N)$. This seems at first as though the lack of pairing of modes at positive and negative k may cause problems with enforcing the Hermitian symmetry needed for a real $f(x)$, this is clearly not the case, since we can start with a real f and generate the N Fourier coefficients as above.

Finally, we should prove how this connects with our experience with using Fourier series. Here we would say

$$f_k(k) = \frac{1}{\ell}\int_0^\ell f(x)\, \exp[-ikx]\, dx. \tag{11.119}$$

We have taken the integral over 0 to ℓ rather than symmetrically around zero, but the same Fourier coefficient arises as long as we integrate over one period. Now, we have seen that the exact function can be interpolated from the sampled values:

$$f(x) = \sum_{n=-\infty}^{\infty} f(nX)T(x-nX) = \sum_{m=-\infty}^{\infty}\sum_{n=0}^{N-1} f([n+mN]X)T(x-[n+mN]X), \tag{11.120}$$

where $X = \ell/N$, T is the sinc interpolation function, and the second form explicitly sums over all the periodic repetitions of f. Putting this into the integral for f_k gives

$$f_k = \frac{1}{\ell}\sum_{m,n} f([n+mN]X)\int_0^\ell T(x-[n+mN]X)\, \exp[-ix]\, dx \tag{11.121}$$

$$= \frac{1}{\ell}\sum_{m,n} f([n+mN]X)\exp[-ik(n+mN)X]\int_{y_1}^{y_2} T(y)\, \exp[-iky]\, dy \tag{11.122}$$

$$= \frac{1}{\ell}\sum_{n=0}^{N-1} f(nX)\exp[-iknX]\sum_{m=-\infty}^{\infty}\int_{y_1}^{y_2} T(y)\, \exp[-iky]\, dy \tag{11.123}$$

$$= \frac{1}{\ell}\sum_{n=0}^{N-1} f(nX)\exp[-iknX]\int_{-\infty}^{\infty} T(y)\, \exp[-iky]\, dy. \tag{11.124}$$

The successive simplifications use (a) the fact that f is periodic, so $f(x+NX) = f(x)$; (b) the fact that allowed wavenumbers are a multiple of $2\pi/\ell$, so kNX is a multiple of 2π; (c) recognising that the y limits are $y_1 = -(n+mN)X$ and $y_2 = \ell - (n+mN)X$, so that summing over m joins together segments of length ℓ into an integral over all values of y. But as we saw in section 10.2.1, the integral has the constant value X while $|k|$ is less than the Nyquist frequency. Thus, for the allowed values of k, $f_k = (1/N)\sum_{n=0}^{N-1} f(nX)\exp[-iknX]$, so that the DFT gives the exact integral for the Fourier coefficient.

11.2 The FFT

We have seen the advantages of the DFT in data compression, meaning that it is widely used in many pieces of contemporary consumer electronics. There is therefore a strong motivation to compute the DFT as rapidly as possible; the Fast Fourier Transform does exactly this.

At first sight, there may seem little scope for saving time. If we define the complex number $W \equiv \exp[-i2\pi/N]$, then the DFT involves us calculating the quantity

$$F_m \equiv \sum_{n=0}^{N-1} f_n W^{nm}. \tag{11.125}$$

The most time-consuming part of this calculation is the complex multiplications between f_n and powers of W. Even if all the powers of W are precomputed and stored, there are still $N-1$ complex multiplications to carry out for each of $N-1$ non-zero values of m, so apparently the time for DFT computation scales as N^2 for large N.

The way to evade this limit is to realise that many of the multiplications are the same, because $W^N = 1$ but W^{nm} has nm reaching large multiples of N – up to $(N-1)^2$. As an explicit example, consider $N = 4$:

$$F_0 = f_0 + f_1 + f_2 + f_3 \tag{11.126}$$
$$F_1 = f_0 + f_1 W + f_2 W^2 + f_3 W^3 \tag{11.127}$$
$$F_2 = f_0 + f_1 W^2 + f_2 W^4 + f_3 W^6 \tag{11.128}$$
$$F_3 = f_0 + f_1 W^3 + f_2 W^6 + f_3 W^9. \tag{11.129}$$

There are apparently 9 complex multiplications (plus a further 5 if we need to compute the powers W^2, W^3, W^4, W^6, W^9). But the only distinct powers needed are W^2 & W^3. The overall transform can then be rewritten saving four multiplications: removing a redundant multiplication by W^4; recognising that the same quantities appear in more than one F_i; and that some multiplications distribute over addition:

$$F_0 = f_0 + f_1 + f_2 + f_3 \tag{11.130}$$
$$F_1 = f_0 + f_2 W^2 + W(f_1 + f_3 W^2) \tag{11.131}$$
$$F_2 = f_0 + f_2 + W^2(f_1 + f_3) \tag{11.132}$$
$$F_3 = f_0 + f_2 W^2 + W^3(f_1 + f_3 W^2). \tag{11.133}$$

So now there are 5 multiplications, plus 2 for the powers: a reduction from 14 to 7.

It would take us too far afield to discuss how general algorithms for an FFT are constructed to achieve the above savings for any value of N. The book *Numerical Recipes* by Press et al. (CUP) has plenty of detail. The result is that the naive $\sim N^2$ time requirement can be reduced to $\sim N \ln N$, provided N has only a few small prime factors – most simply a power of 2.