

The representation of sea surfaces as sums of sinusoids suggests further analysis using Fourier transforms, which will be a fundamental mathematical tool for our subsequent study of sea surfaces. There are many texts on Fourier transforms. Bracewell (1986) is a standard reference, and excellent sets of lecture notes and videos of lectures can be found on the web. The results needed for the subsequent discussions will therefore be stated without proof.

Continuous Fourier Transforms

Given a real function $z(x)$ of a continuous variable x , the Fourier transform $\hat{z}(\nu)$ of $z(x)$ is defined as

$$\hat{z}(\nu) \equiv \mathcal{F}\{z(x)\} \equiv \int_{-\infty}^{\infty} z(x)e^{-i2\pi\nu x} dx. \quad (1)$$

The inverse Fourier transform is given by

$$z(x) \equiv \mathcal{F}^{-1}\{\hat{z}(\nu)\} \equiv \int_{-\infty}^{\infty} \hat{z}(\nu)e^{+i2\pi\nu x} d\nu. \quad (2)$$

It can be shown that if we insert the $\hat{z}(\nu)$ integral of Eq. (likesection1) into Eq. (likesection2), we recover $z(x)$. (This is known as Fourier's integral theorem, the proof of which is not trivial.) Equations (likesection1) and (likesection2) are termed a Fourier transform pair.

Understanding the units of a Fourier transform is important. In our case, $z(x)$ is sea surface elevation, and both z and x have units of meters. Equation (likesection1) shows that $\hat{z}(\nu)$ thus has units of m^2 . The variance of z also has units of m^2 , which gives the first hint at a profound connection between the variance of a physical quantity and its Fourier transform. The units of m^2 in the Fourier transform can be rewritten as $\text{m}/(1/\text{m})$, which is units of z divided by units of spatial frequency ν . Indeed, the transform $\hat{z}(\nu)$ can be interpreted as showing "how much of z there is per unit frequency interval." The inverse transform then has units of (z over frequency) times frequency, which returns the original units of z .

A Fourier transform is a *spectral density function*. The integral of a spectral density function over a given frequency interval gives the variance in the physical quantity contributed by the frequencies in the integration interval. Density functions are rather peculiar mathematical creatures compared to point functions, which simply give the value of something at a given value of the independent variable (e.g. the temperature at location x and time t). The blackbody radiation function is another example of a spectral density function. The blackbody function shows how much energy is emitted (at a given temperature) per unit frequency interval of the emitted electromagnetic radiation. The blackbody function is discussed on the page A Common Misconception. If you are not familiar with the distinction between point and density functions, especially regarding how to change variables in density functions, you should take a look at the page on blackbody spectra before continuing with the present discussion.

The Fourier transform definitions above with the 2π in the exponents are those of the "Stanford school" of Bracewell (1986) and Goodman (1996). You will see others in the literature. For example, if we use $k = 2\pi\nu$ as the frequency variable, then Eqs. (likesection1)

and (likesection2) become

$$\hat{z}(k) \equiv \mathcal{F}\{z(x)\} \equiv \int_{-\infty}^{\infty} z(x)e^{-ikx} dx \quad (3)$$

and

$$z(x) \equiv \mathcal{F}^{-1}\{\hat{z}(k)\} \equiv \frac{1}{2\pi} \int_{-\infty}^{\infty} \hat{z}(k)e^{+ikx} dk. \quad (4)$$

This reappearance of the 2π in the second equation is required so that the inverse transform of the transform gets you back to where you started. In the $e^{\pm ikx}$ version, some people put the $1/2\pi$ in front of the other integral, and some put a $1/\sqrt{2\pi}$ in front of each integral. Some authors, e.g. Press et al. (1992), put the $+i$ in Eq. (likesection1) and the $-i$ in Eq. (likesection2). The choice of which sign to use on the i and where to put the 2π is almost a religion—most people stay with what they first learned, are convinced of the superiority of their definition, and are willing to die rather than change. Fortunately, it doesn't matter which definitions you use, so long as you are consistent in how the transform pair is defined so that you get back to where you started if you inverse transform a transform, or vice versa.

In our work, $z(x)$ is the sea surface elevation, which is a real number. However, even though $z(x)$ is real, $\hat{z}(\nu)$ (or $\hat{z}(k)$) is complex. Expanding the complex exponential in Eq. (likesection1) as the sum of a cosine and a sine, it is easy to see that $\hat{z}^*(\nu) = \hat{z}(-\nu)$. Such functions are called *Hermitian*. *A real function has a Hermitian Fourier transform. Conversely, if a function is Hermitian, it has a real inverse Fourier transform.* The Hermitian property will be an important constraint on the next pages when we wish to generate random realizations of a sea surface by computing the inverse Fourier transform of a complex function: we will have to conjure up a Hermitian function \hat{z} so that we end up with a real sea surface.

Discrete Fourier Transforms

Now suppose that we have sampled the sea surface $z(x)$ at a set of N evenly spaced points $x_r, r = 0, 1, \dots, N - 1$, in a region of size L ; $x_r = r\Delta x = rL/N$. We want to describe this sampled sea surface $z(x_r)$ as a sum of sinusoids. In general, these N values can be represented as a sum of a constant term, $N/2$ cosine terms, and $N/2 - 1$ sine terms (recall that there is no two-point sine term):

$$z(x_r) = \frac{a_0}{2} + \sum_{u=1}^{N/2} [a_u \cos(k_u x_r) + b_u \sin(k_u x_r)], \quad (5)$$

with $b_{N/2} \equiv 0$. Note that the sum runs from the fundamental frequency ($u = 1, k_1 = 2\pi/L$) through the Nyquist frequency ($u = N/2, k_{N/2} = 2\pi/(2\Delta x)$), with only a cosine term for the two-point wave. This sum is equivalent to

$$z(x_r) = \sum_{u=-N/2+1}^{N/2} c_u e^{ik_u x_r}, \quad (6)$$

which also contains a total of N independent real and imaginary parts of the c_n coefficients. (Recall from Eq. (5) of the previous page that $c_{-k} = c_k^*$, so these coefficients are not independent for $+k$ and $-k$ pairs.) These equations are the discrete-variable forms of Eqs. (6) and (7) of the previous page.

Comment on notation: It is common to use the letters i, j, k for dummy summation indices. However, we've already used i for $\sqrt{-1}$ and k for angular wavenumber, so the preceding equations would be hopelessly confusing if we reused i and k for summation indices. We will therefore use r and s for indices on spatial variables, e.g., (x_r, y_s) , and u and v for indices on frequency variables, e.g., k_u or ν_v . n and m also will be used as needed for dummy indices.

We now have a finite number N of discrete samples $z(x_r)$ of the sea surface, so we need a discrete form of the Fourier transform. The discrete Fourier transform (DFT) of $z(x_r)$ is defined as

$$\hat{z}(\nu_u) \equiv \mathcal{D}\{z(x_r)\} \equiv \frac{1}{N} \sum_{r=0}^{N-1} z(x_r) e^{-i2\pi\nu_u x_r}.$$

for $u = 0, 1, \dots, N - 1$. Recalling that $\nu_u = u/L = u/(N\Delta x)$ and $x_r = r\Delta x = rL/N$ gives $\nu_u x_r = ru/N$. It is also common to write $z(x_r) = z(r)$ and $\hat{z}(\nu_u) = \hat{z}(u)$, in which case the previous equation becomes

$$\hat{z}(u) \equiv \mathcal{D}\{z(r)\} \equiv \frac{1}{N} \sum_{r=0}^{N-1} z(r) e^{-i2\pi ru/N}. \quad (7)$$

The corresponding inverse discrete Fourier transform is given by

$$z(r) \equiv \mathcal{D}^{-1}\{\hat{z}(u)\} \equiv \sum_{u=0}^{N-1} \hat{z}(u) e^{+i2\pi ru/N}. \quad (8)$$

It was emphasized above that the continuous function $\hat{z}(\nu)$ defined by the continuous Fourier transform (likesection1) is a density function with units of $z(x)$ per spatial frequency interval, e.g., $m/(1/m)$ if z is sea surface height in meters. However, the discrete function $\hat{z}(u)$ defined by the discrete Fourier transform (likesection7) has the same units as $z(r)$. The discrete Fourier transform is a point function that shows how much of $z(r)$ is contained in a *finite* frequency interval $\Delta\nu = 1/L$ centered at frequency $\nu_u = u/L$. Discrete Fourier transforms convert point functions $z(r)$ to point functions $\hat{z}(u)$.

The \mathcal{F} and \mathcal{D} notations will be used to distinguish continuous vs. discrete Fourier transforms. As just seen, the continuous and discrete transforms are different mathematical constructs with different units and interpretations; they must not be confused. Likewise, if necessary, a subscript can be appended to show the frequency variable, e.g., $\mathcal{F}_\nu\{z(x)\}$ as in Eq. (likesection1) or $\mathcal{F}_k\{z(x)\}$ as in Eq. (likesection3). As always, there are competing definitions. Equations (likesection7) and (likesection8) are used in Bracewell and the IDL computer language. *Numerical Recipes* interchanges the i and $-i$. Matlab puts the $1/N$ factor on the inverse transform. Also, Matlab does not support array indices of 0, so the summation indices are shifted from 0 to $N - 1$ to 1 to N , with a corresponding shift from ru to $(r - 1)(u - 1)$ in the exponentials. The devil is in the details, and details like where

to put the $1/N$ factor and differences in array indexing in different computer languages can cause great misery when it comes time to actually write computer programs or to compare results computed by different canned subroutines.

The sums in the last two equations require computing complex exponentials (i.e., sines and cosines), multiplying by the corresponding values of $z(r)$ or $\hat{z}(u)$ and adding up the results. These equations can be evaluated for any value of N . The number of computations required to do this is of order N^2 . The computation time thus increases very rapidly for large N .

However, a classic paper by Cooley and Tukey (1965) showed how these sums can be computed in order $N \log_2 N$ computations, if N is a power of 2. Their technique is now called the *Fast Fourier Transform* or FFT. The difference in computer time becomes enormous for large N . For example, if $N = 2^{12} = 4096$, then $N \log_2 N = 4096 \times 12$, and the difference in computation time is a factor of $N^2/(N \log_2 N) \approx 341$. Thus in the case of $N = 4096$, a roughly 6 minute computer run becomes a 1 second run. The development (or, perhaps, reinvention, since the basic idea goes all the way back to Gauss) of the FFT was a major advance in numerical analysis, which enables the computations on the following pages to be performed extremely efficiently. Subroutines for computing FFTs and inverse FFTs are available in all computer languages commonly used in science (Fortran, C, IDL, Matlab, etc). Fortunately we do not need to concern ourselves here with the details of how the FFT algorithm actually works, any more than we need to worry about how a canned subroutine actually computes the cosine of a number. If you are interested in how the FFT works, a web search will yield many detailed explanations. It is important to remember that the FFT is not another type of transform; the FFT is a numerically efficient way to evaluate the DFT if the number of data values is a power of two.

The one-dimensional (1-D) equations seen above are easily extended to two or more dimensions. For two dimensions (x, y) , we can sample a region of size L_x by L_y meters over N_x points in the x direction and N_y points in the y direction, with N_x and N_y both powers of 2 so we can use FFTs. Equations (likesection5) and (likesection6) then become

$$\begin{aligned} z(x_r, y_s) &= \& \frac{a_0}{2} + \sum_{u=1}^{N_x/2} \sum_{v=1}^{N_y/2} [a_{u,v} \cos(k_u x_r + k_v y_s) + b_{u,v} \sin(k_u x_r + k_v y_s)] \\ &= \& \sum_{u=-N_x/2+1}^{N_x/2} \sum_{v=-N_y/2+1}^{N_y/2} c_{u,v} e^{i(k_u x_r + k_v y_s)}. \end{aligned} \quad (9)$$

The corresponding 2-D DFT pair is

$$\hat{z}(u, v) \equiv \mathcal{D}\{z(r, s)\} \equiv \frac{1}{N_x N_y} \sum_{r=0}^{N_x-1} \sum_{s=0}^{N_y-1} z(r, s) e^{-i2\pi(ru/N_x + sv/N_y)} \quad (10)$$

and

$$z(r, s) \equiv \mathcal{D}^{-1}\{\hat{z}(u, v)\} \equiv \sum_{u=0}^{N_x-1} \sum_{v=0}^{N_y-1} \hat{z}(u, v) e^{+i2\pi(ru/N_x + sv/N_y)}. \quad (11)$$

It will be important below to keep notational track of continuous vs. discrete versions of various functions. For any variable \mathcal{S} , $\mathcal{S}(k)$ will denote a continuous function of k , $\mathcal{S}(k = k_u)$

will denote the continuous function evaluated at the discrete value k_u , and $\mathcal{S}(k_u) = \mathcal{S}(u)$ will denote a discrete point function. Keep in mind that the density function $\mathcal{S}(k = k_u)$ and the point function $\mathcal{S}(k_u)$ have different units. In the following pages $\mathcal{S}(k)$ and $\mathcal{S}(k_u)$ will denote the continuous and discrete versions, respectively, of wave elevation variance spectra.

The differences in units between continuous and discrete Fourier amplitudes sometimes makes it tricky to make the transition between discrete and continuous versions of the same quantity. In particular, it will be necessary to explicitly include the frequency intervals in some of the later calculations that involve both continuous and discrete variables. For example, if we have a continuous density function and we need to convert to a corresponding discrete function, we must multiple the continuous function by the finite frequency interval, e.g.

$$\hat{z}(u) = \hat{z}(\nu = \nu_u)\Delta\nu. \quad (12)$$

Conversely, if we have discrete amplitudes $\hat{z}(u)$ and we wish to estimate the continuous spectral density $\hat{z}(\nu)$, then we must divide by the frequency interval:

$$\hat{z}(\nu) = \hat{z}(u)/\Delta\nu. \quad (13)$$

(If you are an optical oceanographer familiar with the scattering phase function, you can find an analogous situation in the estimation of the scattering phase function from measurements of scattered light. The scattering phase function is a measure of how much light energy is scattered from an incident direction into a particular direction, per unit solid angle; it therefore has units of 1/steradian. If you measure the scattered light using an instrument with a finite solid angle $\Delta\Omega$, then you get the total amount of energy scattered into the solid angle $\Delta\Omega$. To estimate the phase function from this measurement, you must divide the measured value by the solid angle of the instrument; this gets you back to units of 1/steradian.)

Frequency Ordering

There is a peculiarity to most (perhaps all) FFT subroutines. The discrete FFT of Eq. (like section 7) returns an array of N complex numbers $\hat{z}(u)$, which are associated with N discrete spatial frequencies. What is peculiar is the order in which the elements of the $\hat{z}(u)$ array are returned by an FFT subroutine.

Let Δ_f represent the fundamental frequency. If wavenumber ν is the frequency variable, then $\Delta_f = 1/L$. If angular spatial frequency k is the frequency variable, then $\Delta_f = 2\pi/L$; for temporal angular frequency ω , $\Delta_f = 2\pi/T$. In any case, the discrete frequencies associated with the discrete Fourier amplitudes are equally spaced at intervals of Δ_f and are in the negative-to-positive order

$$\left\{ -\frac{N}{2} + 1, -\frac{N}{2} + 2, \dots, -1, 0, 1, \dots, \frac{N}{2} - 1, \frac{N}{2} \right\} \Delta_f. \quad (14)$$

I'll call this the “math frequency order” because this is the natural order of arranging values in mathematics. This frequency order is convenient for plotting all of the amplitudes, as

will be seen in the following pages. Plots showing both negative and positive frequencies are called “two-sided,” and examples will be seen below.

However, FFT routines generally return their amplitudes corresponding to frequencies in the order of 0 first, then the positive frequencies from the smallest to the Nyquist frequency, then the negative frequencies in reverse order:

$$\left\{ 0, 1, \dots, \frac{N}{2} - 1, \frac{N}{2}, -\frac{N}{2} + 1, -\frac{N}{2} + 2, \dots, -1, \right\} \Delta_f. \quad (15)$$

I’ll call this the “FFT frequency order.” Given the Hermitian symmetry of the amplitudes about the 0 frequency, the FFT order is convenient for plotting amplitudes only for the positive frequencies, with the negative-positive symmetry of \hat{z} understood. Such plots are called “one-sided” or “folded.”

Either frequency order can be obtained from the other by a repeated *circular shift*, which moves an array element off of one end of an array and copies it to the other end of the array, shifting all elements right or left by one position in the process. The detail to watch is that different computer languages define a circular shift in different ways. For example, the IDL routine SHIFT (and the Matlab routine CIRCSHIFT) moves the array elements to the *right* for a “positive” shift (a negative shift moves elements to the left), whereas the Fortran 95 CSHIFT routine moves the array elements to the *left* for a positive shift (a negative shift moves elements to the right). Thus

In IDL:

```
math order = SHIFT(FFT order, N/2-1)
FFT order = SHIFT(math order, -(N/2-1))
```

In Fortran95:

```
math order = CSHIFT(FFT order, -(N/2-1))
FFT order = CSHIFT(math order, N/2-1)
```

Some FFT routines allow the user to set a flag specifying which frequency order is to be returned. In any case, sorting out the frequency order of the amplitudes returned by a particular FFT routine, and figuring out how to shift between math and FFT frequency orders in a particular computer language, can drive you to tears.

Note finally that in Eq. (likesection7) you are simply providing an array of $z(r)$ values and getting back the same number of $\hat{z}(u)$ values. What $x(r)$ values correspond to the $z(r)$ values is irrelevant. That is, $x(r), r = 0, \dots, N - 1$ might correspond to the spatial range from $x = 0$ to L , or from $x = -L/2$ to $L/2$, or to any other x range. It is only the number of samples N and the corresponding $z(u)$ values that matters. In other words, the amplitudes coming out of the Fourier transform are independent of the origin of the spatial coordinate system. The frequencies depend on both the number of points N and the physical size of the sampled region via the presence of L (or time interval T) in the fundamental frequency Δ_f . Thus the size of the region sampled and the number of samples, along with the sample values themselves, fully define the associated Fourier transforms.

Interpretation of Negative Frequencies

The appearance of negative frequencies in Fourier transforms may seem somewhat mysterious. Frequency, after all, is a physical quantity that is inherently a positive number, e.g., the number of wavelengths in a given distance. However, one way to interpret mathematically negative frequencies is that they are simply the mathematical price we pay for the convenience of using complex numbers. Consider, for example, the representation of the cosine as a sum of complex exponentials for the u^{th} frequency:

$$\begin{aligned}\cos(2\pi\nu_u x_r) &= \& \frac{1}{2} (e^{i2\pi\nu_u x_r} + e^{-i2\pi\nu_u x_r}) \\ &= \& \frac{1}{2} (e^{i2\pi(+\nu_u)x_r} + e^{i2\pi(-\nu_u)x_r}) .\end{aligned}$$

We can interpret the complex representation of the real cosine as having one term with a positive frequency $+\nu_u$ and one term with a negative frequency $-\nu_u$. A similar equation holds for the complex representation of $\sin(2\pi\nu_u x_r)$. The Fourier transform of a real function always contains both negative and positive frequencies, which arise from the complex exponentials in the definition of the transform.

Additional comments about negative frequencies will be made on later pages, where it will be seen that positive frequencies can be associated with waves propagating in the $+x$ direction, and negative frequencies correspond to waves propagating in the opposite, $-x$, direction.

Parseval's Relations

The physical and spectral variables of a continuous Fourier transform pair satisfy

$$\int_{-\infty}^{\infty} |z(x)|^2 dx = \int_{-\infty}^{\infty} |\hat{z}(\nu)|^2 d\nu , \quad (16)$$

which is known as Parseval's relation. For complex amplitudes, $|\hat{z}|^2 = \hat{z}\hat{z}^*$. The corresponding equation for the discrete Fourier transform pair defined by Eqs (likesection7) and (likesection8) is

$$\sum_{r=0}^{N-1} |z(r)|^2 = N \sum_{u=0}^{N-1} |\hat{z}(u)|^2 . \quad (17)$$

The extension to the two-dimensional case is straightforward:

$$\sum_{r=0}^{N_x-1} \sum_{s=0}^{N_y-1} |z(r, s)|^2 = N_x N_y \sum_{u=0}^{N_x-1} \sum_{v=0}^{N_y-1} |\hat{z}(u, v)|^2 . \quad (18)$$

The discrete forms of Parseval's relations provide important checks on numerical calculations. For example, it is easy to misplace factors of N , which appear in different places depending on the exact form used for the definition of the discrete transforms.

You can take a class in Fourier transforms and prove many more beautiful theorems about their properties. However, we have now assembled the mathematical tools needed to describe wind-blown sea surfaces, and so we can now get back to oceanography.