

This page uses a contrived example to illustrate the preceding theory for a one-dimensional (1D) surface. Detailed comments on this simple example emphasize the mathematical subtleties and physical characteristics of Fourier transforms and wave variance spectra derived from surface elevations. A thorough understanding of this example takes us much of the way to understanding the case of a real ocean surface.

An ad hoc, one-dimensional wave profile is constructed using the formula

$$z(x_r) = z(r) = \sum_{j=0}^{N/2} A(j) \cos[2\pi j x_r \nu_f + \phi(j)] \quad r = 0, \dots, N - 1. \quad (1)$$

The x_r locations are given by $r\Delta x = rL/N$, where L is the length of the sea surface region being sampled and N is the number of samples. $\nu_f = 1/L$ is the fundamental spatial frequency, that is, the spatial frequency or wavenumber of the wave with a wavelength of L . The amplitude of the wave at the j^{th} frequency, $j = 1, 2, \dots, N/2$, is chosen to be

$$A(j) = 0.1 \exp(-3j/N),$$

and $A(0) = 0$. The phase of the j^{th} wave component is randomly distributed over $[0, 2\pi)$ using

$$\phi(j) = 2\pi\mathcal{U}$$

where \mathcal{U} is a uniform $[0, 1)$ random number. A different random number is drawn for each j value.

The upper left panel of Fig. figure1 shows the surface generated in this manner for $L = 10$ m, $N = 16$, and a particular set of random phases. Note that N is a power of 2 as will be needed for the FFT. The thin colored lines show the $N/2 + 1 = 9$ waves for each of the frequencies. The blue line is the wave for the fundamental frequency $\nu_f = 1/L = 0.1 \text{ m}^{-1}$; the thin black line is the two-point wave at the Nyquist frequency $\nu_{\text{Ny}} = 1/(2\Delta x) = 0.8 \text{ m}^{-1}$; the purple line is the constant $j = 0$ wave, which is set to $z = 0$ for the mean sea surface. The black dots connected by the thick black line show the sum of the individual waves. These points represent a discrete sampling of the continuous sea surface elevation.

In this example, the sampled region of the sea surface is $L = 10$ m in length, but the $N = 16$ samples do not include the point at $x = 10$ m. This is because the surface elevation at $x = L$ is always the same as at $x = 0$ when using Fourier techniques. Resolving the surface as a sum of sinusoids that are harmonics of the fundamental frequency $\nu_f = 1/L$ gives sinusoids that always return to their initial value after distance L . Real sea surfaces are of course not periodic, but we do not know the true value at L because it was not measured by the present sampling scheme. (Likewise, we do not know the true surface elevations at points in between the sampled locations.) *When we use Fourier techniques to generate random surface realization, we are always generating a sea surface that is a periodic tiling*; the tile dimension is L . This periodicity is useful if we want to generate a visual rendering of a large region of sea surface from a smaller computed region; the edges of the small tiles will match perfectly and the larger surface will often look reasonable, if you don't look too closely. An example of a tiled two-dimension surface can be seen in Fig. 3.9 of Mobley (2016).

We now take the sequence of the $N = 16$ real wave elevations $z(r)$ seen in Fig. figure1 and feed them into an FFT routine. We soon get back 16 complex numbers, the $\hat{z}(\nu_u) = \hat{z}(u)$ Fourier amplitudes, at a set of 16 corresponding frequencies ν_u . The upper right panel of Fig. figure1 plots the real part of the $\hat{z}(u)$ complex numbers, and the lower-left panel plots the imaginary part.

Note first that the FFT routine returned both negative and positive spatial frequencies:

$$\nu = \left(-\frac{N}{2} + 1\right)\Delta\nu = -0.7 \text{ m}^{-1}, \dots, -0.1, 0, 0.1, \dots, \frac{N}{2}\Delta\nu = 0.8 \text{ m}^{-1},$$

$L = 10 \text{ m}$; $N = 16$; $\nu_f = 0.1 \text{ m}^{-1}$; $\nu_{Ny} = 0.8 \text{ m}^{-1}$; random phases

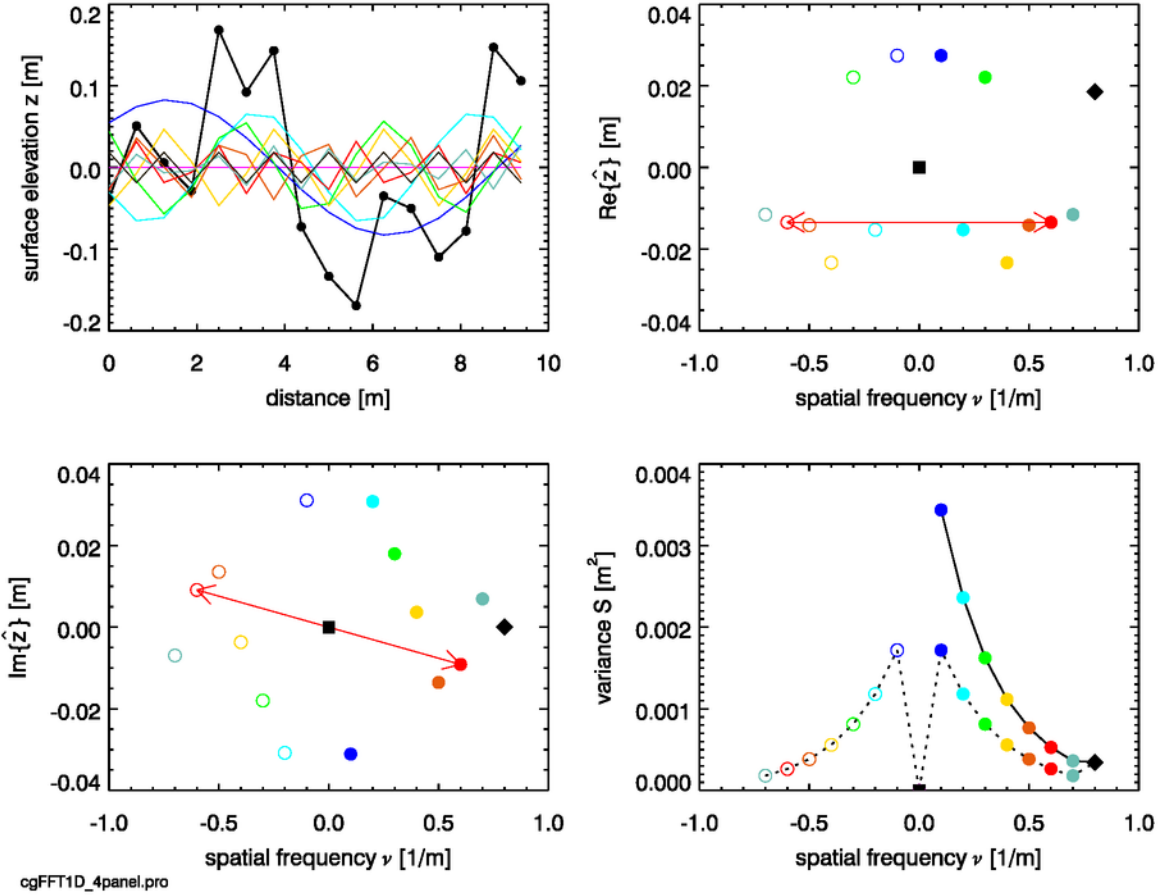


Figure 1: A one-dimensional random sea surface and its Fourier transform and variance spectrum. The black dots in the upper-left panel show the points $z(u)$ of the sampled sea surface. The light colored lines show the sinusoidal components used to create the surface. The upper-right panel shows the real part of the complex amplitude $\hat{z}(\nu_u)$ and the lower-left panel shows the imaginary part. The lower right panel shows the two-sided and one-sided discrete variance spectra. Compare with Figs. figure2 and figure3.

for a total of $N = 16$ discrete spatial frequencies. Note that the frequency spacing $\Delta\nu$ equals the fundamental frequency $\nu_f = 1/L$. The Fourier Transforms page discussed the interpretation of negative frequencies. (The order of the frequencies as returned by the FFT routine was the "FFT order" discussed on the Fourier Transforms page. The frequencies, and the corresponding amplitudes, were reordered to get the "math order" used for plotting.)

Note next that the real parts of the complex amplitudes $\hat{z}(u)$ are even functions of frequency: $Re\{\hat{z}(-\nu_u)\} = Re\{\hat{z}(+\nu_u)\}$. The imaginary parts are odd functions of frequency: $Im\{\hat{z}(-\nu_u)\} = -Im\{\hat{z}(+\nu_u)\}$. In Fig. ?? the positive and negative frequencies are connected by a red arrow for one particular frequency pair, $\nu_u = \pm 0.6 \text{ m}^{-1}$. If a complex function $c(\nu) = a(\nu) + ib(\nu)$ has an even real part $a(\nu)$ and an odd imaginary part $b(\nu)$, then $c^*(-\nu) = c(\nu)$, i.e. the function is Hermitian. Thus the plots verify that the amplitudes are Hermitian, as is always the case for the Fourier transform of a real function. The Hermitian character of the complex amplitudes means

that these $N = 16$ complex numbers contain only 16 independent real and imaginary numbers, not 32 as would be the case for 16 arbitrary complex numbers. In general the FFT of N real numbers (e.g., N spatial samples of a sea surface) gives back N independent numbers, so that the "information content" of the physical and Fourier representations is the same.

The positive frequency at $\frac{N}{2}\Delta\nu = 0.8 \text{ m}^{-1}$ is the Nyquist frequency. There is, however, no value for the negative of the Nyquist frequency. Note also in the lower left panel that the imaginary part of the amplitude is identically zero at the Nyquist frequency. We will explain these values below.

The lower-right panel of the figure shows the values of $|\hat{z}(\nu_u)|^2$. The values at the negative to positive frequencies are connected by the black dotted line. These points constitute the *two-sided discrete variance spectrum*,

$$\mathcal{S}_{2s}(\nu_u) = |\hat{z}(\nu_u)|^2 \quad \text{for } u = -\frac{N}{2} + 1, \dots, \frac{N}{2}. \quad (2)$$

"Two-sided", denoted by the subscript 2S, refers to spectra showing both the negative and positive frequencies. The variance at zero frequency is the variance contained in the constant mean sea level. This value is zero because we have set the mean sea level to zero.

Oceanographers are often concerned only with the variance at a given magnitude of the spatial frequency, and not with whether the frequency is negative or positive. Nor is there any reason to plot the point at zero frequency, which is usually zero by choice of zero for the mean sea level. It is therefore customary to define the *one-sided* variance spectrum

$$\mathcal{S}_{1s}(\nu_u) = \mathcal{S}_{2s}(-\nu_u) + \mathcal{S}_{2s}(\nu_u), \quad (3)$$

for $u = 1, 2, \dots, \frac{N}{2} - 1$, and $\mathcal{S}_{1s}(\nu_{Ny}) = \mathcal{S}_{2s}(\nu_{Ny})$. Then only the positive frequencies are plotted. The points connected by the solid black line in the lower-right panel of Fig. figure1 comprise the one-sided variance spectrum. In the present simulation, the two-sided spectrum is symmetric for positive and negative frequencies (except for the Nyquist frequency, which does not have a negative counterpart and is always a special case), and the one-sided function is simply twice the value of the two-sided function for the positive frequencies, except for the Nyquist frequency. When you read a paper and it refers to or plots "the variance (or energy or power) spectrum" without further comment, it is always the one-sided spectrum. However, on the next pages we will have to use two-sided spectra, in which case we will have to account for the magnitude difference in one- and two-sided spectra.

There is an important detail to note in the computation and plotting of $\mathcal{S}(u)$, as in the lower-right panel of Fig. figure1. The values of $\mathcal{S}(u)$ were obtained by the discrete Fourier transform of Eq. (??), and $\mathcal{S}(u)$ gives the *variance contained in a finite frequency interval $\Delta\nu = 1/L$ at the discrete frequency ν_u* . $\Delta\nu$ equals the fundamental frequency and is the frequency interval used in the calculations and the plot. As noted in the discussion of the discrete transform, $\mathcal{S}(u)$ is a point function. As was seen in Eq. (13) of the Fourier Transforms page, if we wish to convert the discrete $\mathcal{S}(u)$ to an estimate of the continuous variance spectral *density* $\mathcal{S}(\nu)$, we must divide by the discrete function by the frequency interval: $\mathcal{S}(\nu) = \mathcal{S}(u)/\Delta\nu$. The units of $\mathcal{S}(\nu)$ are then $m^2/(1/m)$, as expected for a spectral density function of spatial frequency. It is important to distinguish between a discrete variance point function and a continuous variance spectral density.

Now return to Eq. (1) and set all of the phases $\phi(j)$ to zero. We are then adding together cosines to create the surface wave profile, which is seen in the upper left panel of Fig. figure2. The FFT of this profile gives the real part of $\hat{z}(\nu_u)$ as positive numbers except for the 0 frequency, and the imaginary part is identically zero for all frequencies.

If we set all of the phases $\phi(j)$ to $\pi/2$, we are then adding together sines to create the surface wave profile, which is seen in the upper left panel of Fig. figure3. The FFT of this profile gives the

real part of $\hat{z}(\nu_u)$ identically zero and the imaginary part is nonzero except for the 0 and Nyquist frequencies.

These two figures show that the real part of the complex amplitude $\hat{z}(\nu_u)$ tells us how much of $z(x_r)$ is composed of cosine waves, and the imaginary part shows how much of $z(x_r)$ is composed of sine waves. This explains why the imaginary part of the amplitude is zero at the Nyquist frequency. The two-point wave at the Nyquist frequency is inherently a cosine wave because, as noted previously, a two-point sine wave is sampled only at its zero values. The general case of a wave component with a phase that is neither 0 (nor a multiple of 2π) nor $\pi/2$ (nor an odd integer multiple of $\pi/2$) can be written as a sum of cosine and sine waves, as in Eq. (??). In that case, both the real and imaginary parts of the amplitudes are nonzero (except for the special cases of the 0 and Nyquist frequencies).

Note that in each of these three simulations, which differ only by the phases of the component sinusoidal waves, the variance spectrum is exactly the same (except at the Nyquist frequency), as seen in the lower right panels of Figs. figure1-figure3. That is to say, the variance contained in a wave does not depend on the reference coordinate system used to describe it, even though the Fourier amplitudes $\hat{z}(u)$ do depend on the coordinate system. The variance depends only on the amplitude of the wave. The variance at the Nyquist frequency is largest when cosine waves are added and is zero when sine waves are added. In the first case, we have the maximum possible amplitude of the two-point cosine wave, and in the latter case there is no two-point wave.

$L = 10 \text{ m}; N = 16; \nu_t = 0.1 \text{ m}^{-1}; \nu_{Ny} = 0.8 \text{ m}^{-1}; \text{phases} = 0$

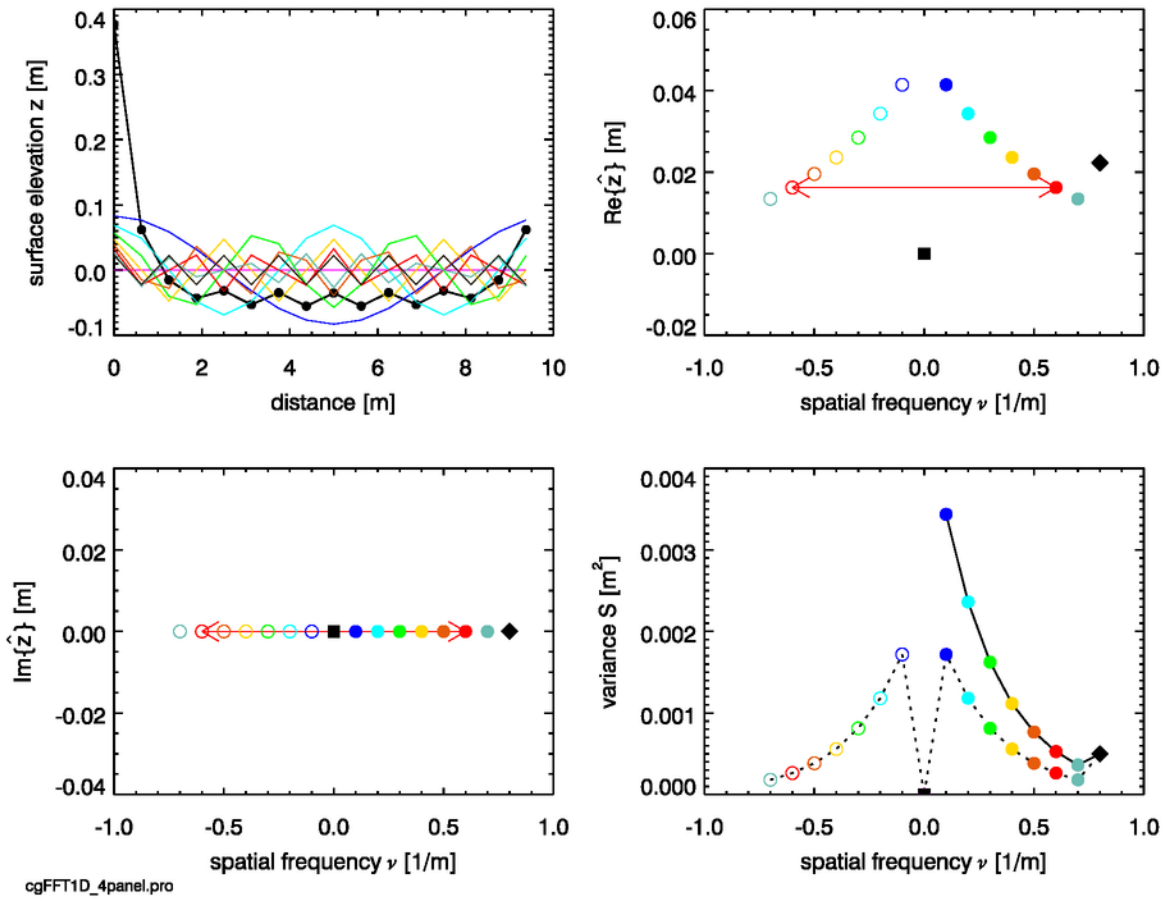


Figure 2: A one-dimensional surface composed of cosine waves. Compare with Figs. figure1 and figure3.

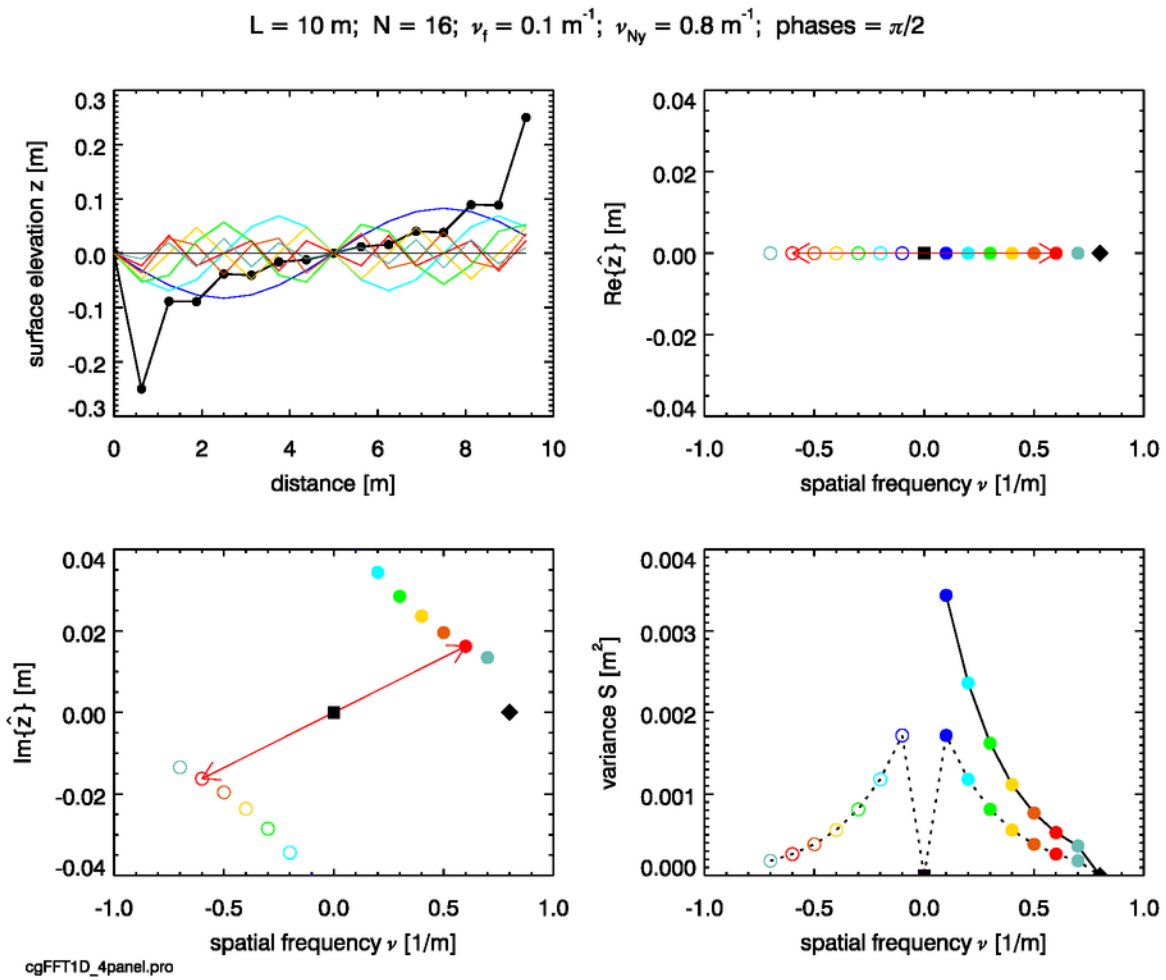


Figure 3: A one-dimensional surface composed of sine waves. Compare with Figs. figure1 and figure2.