

The previous page showed the major features of the Fourier analysis of a sea-surface elevation record. We started with a sample of a random sea surface and ended with the corresponding discrete variance spectrum (or estimate of the variance spectral density after division by the frequency interval). This page shows how to go in the reverse direction: start with a variance spectrum and generate a random realization of the corresponding sea surface. We first outline the theory, and then show a specific example.

Theory for 1-D Surfaces

To create a one-dimensional (1-D) slice through a sea surface, the essence of the process is as follows:

1. **Choose the domain size.** To create a sea surface over a spatial domain at a given time, we pick the length L of the region $[0, L]$. To generate a time series at a given point, we pick the length of the time series.
2. **Choose the number of points for the DFT.** This number N is the number of frequencies at which we will sample the variance spectrum, and equals the number of samples of the sea surface that will be generated. In normal usage, pick N to be a power of 2 so that an inverse FFT routine can be used to evaluate the inverse DFT.
3. **Choose the frequency variable.** To generate a sea surface over a spatial domain at a given time, we can use either wavenumber ν or angular spatial frequency k . To generate a time series at a given point, we can use either f or ω .
4. **Choose a variance spectrum** The variance spectrum must be expressed in terms of the chosen frequency variable.
5. **Choose the wind speed.** Pick a wind speed, and perhaps other physical parameters such as the age of the waves to be generated if required by the chosen variance spectrum.
6. **Create random Hermitian amplitudes.** This is the tricky part. We must create an array of randomized discrete Hermitian Fourier amplitudes $\hat{z}(u)$, starting with the chosen continuous variance spectrum.
7. **Take the inverse DFT of the random amplitudes.** The inverse DFT converts the Fourier amplitudes to the physical wave heights.
8. **Extract the sea surface heights.** The inverse DFT of the complex amplitudes returns a complex array. The real part of this array is the random realization of the sea surface heights, and the imaginary part is zero.
9. **Check your results.** This is extremely important during code development. For example, take the DFT of the generated surface heights to see if you get back to the Fourier amplitudes and variance spectrum you started with.

We now proceed through these steps and discuss them in detail for a specific example.

Steps likesection1 and likesection2: Let us generate a sea surface over the region from $x = 0$ to $x = L = 100$ m. The longest wavelength that can be resolved is then 100 m. We use $N = 1024$, which gives a spatial grid resolution of $\Delta x = L/N = 0.0977$ m. This means that the shortest wavelength that can be resolved, the two-point wave or Nyquist wavelength, is $2\Delta x = 0.1954$ m.

Step likesection3: We used wavenumber ν in the previous pages because of its easy interpretation. Now let's use angular spatial frequency k , which is more commonly used. The fundamental frequency is then $k_f = 2\pi/L = 0.0628 \text{ rad m}^{-1}$. The highest frequency sampled, the Nyquist frequency, is $k_{\text{Ny}} = (N/2)k_f = 32.15 \text{ rad m}^{-1}$. Note that $(2\pi)/k_{\text{Ny}} = 0.1954$ m which, as noted above, is the wavelength of the two-point wave.

Step likesection4: For this example, we use the Pierson-Moskowitz variance spectrum in terms of angular spatial frequency k . This is given by Eq. (1) of the Wave Variance Spectra: Examples page. *Note that this is a one-sided spectrum*, which is defined for positive k values.

Step likesection5: The wind speed at 10 m elevation is $U_{10} = 5 \text{ m s}^{-1}$. The wind speed is the only input to the Pierson-Moskowitz spectrum.

Step likesection6: We now discuss in detail how to generate a set of random discrete Fourier amplitudes that are physically consistent with the chosen variance spectrum. These amplitudes must be defined for both positive and negative frequencies, and the amplitudes must be Hermitian. We first define

$$\hat{z}_o(k_u) \equiv \frac{1}{\sqrt{2}} [\rho(k_u) + i\sigma(k_u)] \sqrt{\frac{\mathcal{S}_{1s}(k = k_u)}{2} \Delta k}. \quad (1)$$

Here $\mathcal{S}_{1s}(k = k_u)$ denotes the continuous spectral density $\mathcal{S}_{1s}(k)$ evaluated at $k = k_u$. $\Delta k = k_f$ is the spatial frequency sampling interval. $\hat{z}_o(k_u)$ must be defined for both positive and negative discrete frequencies in order to create the Hermitian amplitudes for use in the inverse DFT. As was mentioned on the previous page, we must convert the one-sided continuous spectrum $\mathcal{S}_{1s}(k)$ into a two-sided discrete variance function by

1. dividing its magnitude by 2, assuming that $\mathcal{S}_{2s}(-k) = \mathcal{S}_{2s}(k)$;
2. multiplying the continuous spectral density by the fundamental frequency interval Δk , which gives the variance contained in a finite frequency interval at each frequency k_u .

To emphasize the discrete vs continuous functions, and for brevity of notation, let us write the frequency index u for the frequency k_u . Then Eq. (likesection1) becomes

$$\hat{z}_o(u) = \frac{1}{\sqrt{2}} [\rho(u) + i\sigma(u)] \sqrt{\mathcal{S}_{2s}(u)}, \quad (2)$$

where $\mathcal{S}_{2s}(u)$ denotes the *two-sided* discrete variance spectrum at frequency k_u . $\hat{z}_o(u)$ can now be evaluated for both positive and negative k_u . The 0 and Nyquist frequencies are always special cases: set $\mathcal{S}_{2s}(0) = 0$ and $\mathcal{S}_{2s}(k_{\text{Ny}}) = \mathcal{S}_{1s}(k_{\text{Ny}})$. $\rho(k_u) \equiv \rho(u)$ and $\sigma(k_u) \equiv \sigma(u)$ are independent random numbers drawn from a normal distribution with zero mean and unit variance, denoted $\rho, \sigma \sim \mathcal{N}(0, 1)$. A different pair is drawn for each u value.

$\hat{z}_o(u)$ is a random variable. Let $\langle \dots \rangle$ denote the expectation of the enclosed variable. The expected value of $|\hat{z}_o(u)|^2$, $\langle \hat{z}_o(u)\hat{z}_o^*(u) \rangle$, gives back whatever variance function is used for $\mathcal{S}_{2s}(u)$:

$$\begin{aligned} \langle \hat{z}_o(u)\hat{z}_o^*(u) \rangle &= \&\left\langle \left\{ \frac{1}{\sqrt{2}}[\rho(u) + i\sigma(u)]\sqrt{\mathcal{S}_{2s}(u)} \right\} \left\{ \frac{1}{\sqrt{2}}[\rho(u) - i\sigma(u)]\sqrt{\mathcal{S}_{2s}(u)} \right\} \right\rangle \\ &= \&\frac{\mathcal{S}_{2s}(u)}{2} [\langle \rho^2 \rangle + \langle \sigma^2 \rangle] = \mathcal{S}_{2s}(u) \end{aligned}$$

because $\langle \rho^2 \rangle = \langle \sigma^2 \rangle = 1$ for $\mathcal{N}(0, 1)$ random variables. Thus $\hat{z}_o(u)$ is consistent with the chosen variance spectrum. However, $\hat{z}_o(u)$ is not Hermitian, so the inverse DFT would not give a real sea surface.

Next define $\hat{z}(u)$ as

$$\hat{z}(u) \equiv \frac{1}{\sqrt{2}} [\hat{z}_o(u) + \hat{z}_o^*(-u)] . \quad (3)$$

This function is clearly Hermitian, so the inverse DFT applied to $\hat{z}(u)$ will give a real-valued $z(x_r) \equiv z(r)$. Moreover, this $\hat{z}(u)$ is consistent with the variance spectrum:

$$\begin{aligned} \langle |\hat{z}(u)|^2 \rangle &= \&\langle \hat{z}(u)\hat{z}^*(u) \rangle \\ &= \&\left\langle \frac{1}{\sqrt{2}} \left[\frac{1}{\sqrt{2}}[\rho(u) + i\sigma(u)]\sqrt{\mathcal{S}_{2s}(u)} + \frac{1}{\sqrt{2}}[\rho(-u) - i\sigma(-u)]\sqrt{\mathcal{S}_{2s}(-u)} \right] \times \right. \\ &\quad \&\left. \frac{1}{\sqrt{2}} \left[\frac{1}{\sqrt{2}}[\rho(u) - i\sigma(u)]\sqrt{\mathcal{S}_{2s}(u)} + \frac{1}{\sqrt{2}}[\rho(-u) + i\sigma(-u)]\sqrt{\mathcal{S}_{2s}(-u)} \right] \right\rangle \\ &= \&\frac{1}{4} \left\langle [\rho^2(u) - i\rho(u)\sigma(u) + i\sigma(u)\rho(u) + \sigma^2(u)] \mathcal{S}_{2s}(u) + \right. \\ &\quad \&[\rho(u)\rho(-u) + i\rho(u)\sigma(-u) + i\sigma(u)\rho(-u) - \sigma(u)\sigma(-u)] \sqrt{\mathcal{S}_{2s}(u)}\sqrt{\mathcal{S}_{2s}(-u)} + \\ &\quad \&[\rho(-u)\rho(u) - i\rho(-u)\sigma(u) - i\sigma(-u)\rho(u) - \sigma(-u)\sigma(u)] \sqrt{\mathcal{S}_{2s}(-u)}\sqrt{\mathcal{S}_{2s}(u)} + \\ &\quad \&[\rho^2(-u) + i\rho(-u)\sigma(-u) - i\sigma(-u)\rho(-u) + \sigma^2(-u)] \mathcal{S}_{2s}(-u) \left. \right\rangle \\ &= \&\frac{1}{2} [\mathcal{S}_{2s}(u) + \mathcal{S}_{2s}(-u)] = \mathcal{S}_{2s}(u) . \end{aligned}$$

Here we have noted that $\langle \rho(u)\rho(-u) \rangle = 0$, etc., because the random variables are uncorrelated for different u values.

Equations (likesection1) and (likesection3) are the key to generating random sea surfaces from variance spectra. $\hat{z}(u)$ defined by these equations contains random noise, which leads to a sea surface with random amplitudes and phases for the component waves of different frequencies. Any one of these surfaces has a variance spectrum that looks like the chosen spectrum plus random noise. However, on average over many realizations, the the noise in these spectra will average out, leaving the variance spectrum. Figure figure1 below illustrates these important ideas, but first we must complete the surface generation.

Step likesection7: Compute the inverse DFT of the $\hat{z}(u)$ of Eq. (likesection3). The result is a complex function $Z(x)$:

$$Z(x_r) \equiv Z(r) = \mathcal{D}^{-1}\{\hat{z}(u)\} .$$

A crucial warning to this step is that the $u = 0, \dots, N - 1$ elements of the $\hat{z}(u)$ array *must* be in the FFT frequency order given by Eq. (13) of the Fourier Transforms page when using an FFT routine to evaluate the DFT. $Z(r)$ is returned with x_r values in the order from $x_0 = 0$ to $x_{N-1} = (N - 1)\Delta x$.

Step likesection8: Extract the surface. The inverse DFT returns a complex array $Z(x_r)$ whose real part is the surface elevations $z(x_r)$ and whose imaginary part is 0. The surface elevations are extracted as the real part of $Z(x_r)$:

$$z(x_r) = \text{Re}\{Z(x_r)\}.$$

Step likesection9: Check the results! There are many places along the way to lose a $\sqrt{2}$ or to mess up array indexing. At the minimum, it is worthwhile to check that the mean of the generated surface is zero, and that the imaginary part of $Z(x_r) = 0$ (to within a small amount of numerical roundoff error).

When developing computer code, or when first learning this material, it is also a good idea to take the forward DFT of $Z(x_r)$ to make sure that the input Fourier amplitudes $\hat{z}(u)$ are recovered, and that the variance spectrum corresponding to $z(x_r)$ is consistent with the one chosen in Step likesection4. Indeed, it was the failure of this check in surfaces I was generating using equations from the literature that led me to develop the Mobley (2016) tutorial and these web pages.

Equations (likesection1) and (likesection3) are, with minor changes in notation, Eqs. (42) and (43), respectively, of Tessendorf (2004). However, Tessendorf’s version of Eq. (likesection1) appears to use a one-sided variance spectrum (his example used the one-sided Phillips spectrum of his Eq. (40)) without the division by 2 seen in Eq. (likesection1), which is needed to convert the one-sided spectrum to a two-sided spectrum. Nor does he show the Δk factor needed to convert a continuous spectral density to a discrete function. His version of Eq. (likesection3) does not contain the overall factor of $1/\sqrt{2}$ seen above. These missing factors mean that in a round-trip calculation

$$\text{variance spectrum} \rightarrow \text{DFT}^{-1} \rightarrow \text{sea surface} \rightarrow \text{DFT} \rightarrow \text{variance spectrum},$$

you do not get back to the original variance spectrum. In other words, *the Tessendorf equations do not conserve wave variance (i.e., wave energy)*. Even if he included the Δk factor in his actual computations, the missing factors of $1/\sqrt{2}$ in his versions of our Eqs. (likesection1) and (likesection3) give an overall factor one-half on the amplitudes, which corresponds to a factor of four error in the variance. That is, waves generated using the Tessendorf equations have amplitudes that are too large.

Tessendorf (2004) discusses much more than just Fourier transform techniques, and his notes have been very influential in the computer graphics industry. In 2008 he deservedly received an Academy Award for Technical Achievement for showing the movie industry how to generate and render sea surfaces, as well as for his many other pioneering accomplishments in efficiently computing and rendering fluid motions into visually appealing images. (The first movie to use his techniques was *Waterworld*, followed by dozens of others including *Titanic*.) When I checked with him about the missing numerical factors, he readily acknowledged that Eqs. (likesection1) and (likesection3) are the correct ones, but pointed out that “Hollywood doesn’t care about conservation of energy.” I suppose that should be

no surprise, since movies seem to have no problem with rockets going faster than the speed of light, sound propagating through the vacuum of outer space, or time travel that violates causality. Tessendorf’s equations are widely cited (especially in the computer graphics literature), always without comment about the missing scale factors. Even if Tessendorf had included the needed numerical factors in his equations, graphics artists would distort the resulting images to make them look “better,” e.g., to make the ocean waves look bigger than nature would allow. That may be acceptable in a fantasy world, but such laxness is not permissible if we wish to use numerically generated waves to compute sea surface optical properties.

Example: A Roundtrip Calculation

Figure figure1 shows an example of 1-D surface waves generated using the Pierson-Moskowitz spectrum for a wind speed of 5 m/s, and the recovery of the variance spectrum from the generated surface. The blue curve in the upper-left panel shows the Pierson-Moskowitz spectrum as defined by Eq. (1) on page Wave Variance Spectra: Examples. The red dots show the frequencies at which the continuous spectrum is sampled. Those dots blur together at the higher frequencies because of the log scale, but the k_u points are equally spaced at intervals of the fundamental frequency $\Delta k = k_f = 2\pi/L = 0.0628$ rad/m. The last sampled frequency is $k_{Ny} = 32.17$ rad/m. The bottom panel shows the sea surface elevations $z(x_r)$ generated for a particular sequence of random numbers $\rho(u), \sigma(u)$.

The red line in the upper-right panel shows the function

$$\mathcal{P}(u) \equiv \frac{|\hat{z}(u)|^2}{\Delta k} = \frac{1}{\Delta k} |\mathcal{D}\{z(r)\}|^2. \quad (4)$$

$\mathcal{P}(u)$ is the discrete variance function for this particular $z(x_r)$ surface. Schuster (1898) called $\mathcal{P}(u)$ a *periodogram*. The periodogram $\mathcal{P}(u)$ contains random noise because $z(x_r)$ is a random realization of the sea surface, which was generated by applying random noise to the the variance spectrum. This particular $z(x_r)$ is analogous to a particular measurement of the sea surface. Had we drawn a difference sequence of random numbers for use in Eq. (like section2), we would have generated a different sea surface, and a different $\mathcal{P}(u)$. However, we can expect that if we average together many different $\mathcal{P}(u)$, corresponding to many different sets of $z(x_r)$, the noise would average out and we would be left with a curve close to the variance spectrum we started with, which is shown in blue. Numerical experimentation shows that averaging 100 $\mathcal{P}(u)$ generated from 100 independent sea surface realizations gives an average $\mathcal{P}(u)$ that is almost indistinguishable from the blue curve at the scale of this plot. Thus $\mathcal{P}(u) \equiv \mathcal{P}(k_u)$ is an approximation of the variance density spectrum $\mathcal{S}(k)$, denoted $\mathcal{P}(k_u) \doteq \mathcal{S}(k)$. This averaging processes leads to the topic of *spectrum estimation*, which considers such problems as how many sets of measurements of $z(x_r)$ are needed to estimate the variance spectrum to within certain error bounds. Fortunately, we need not pursue that here. (The noise in the upper right panel is Gaussian distributed about the theoretical spectrum. However, the log axis makes it look asymmetric about the blue curve.)

At the minimum, you should always check to see that Parseval’s relation, Eq. (17) of the

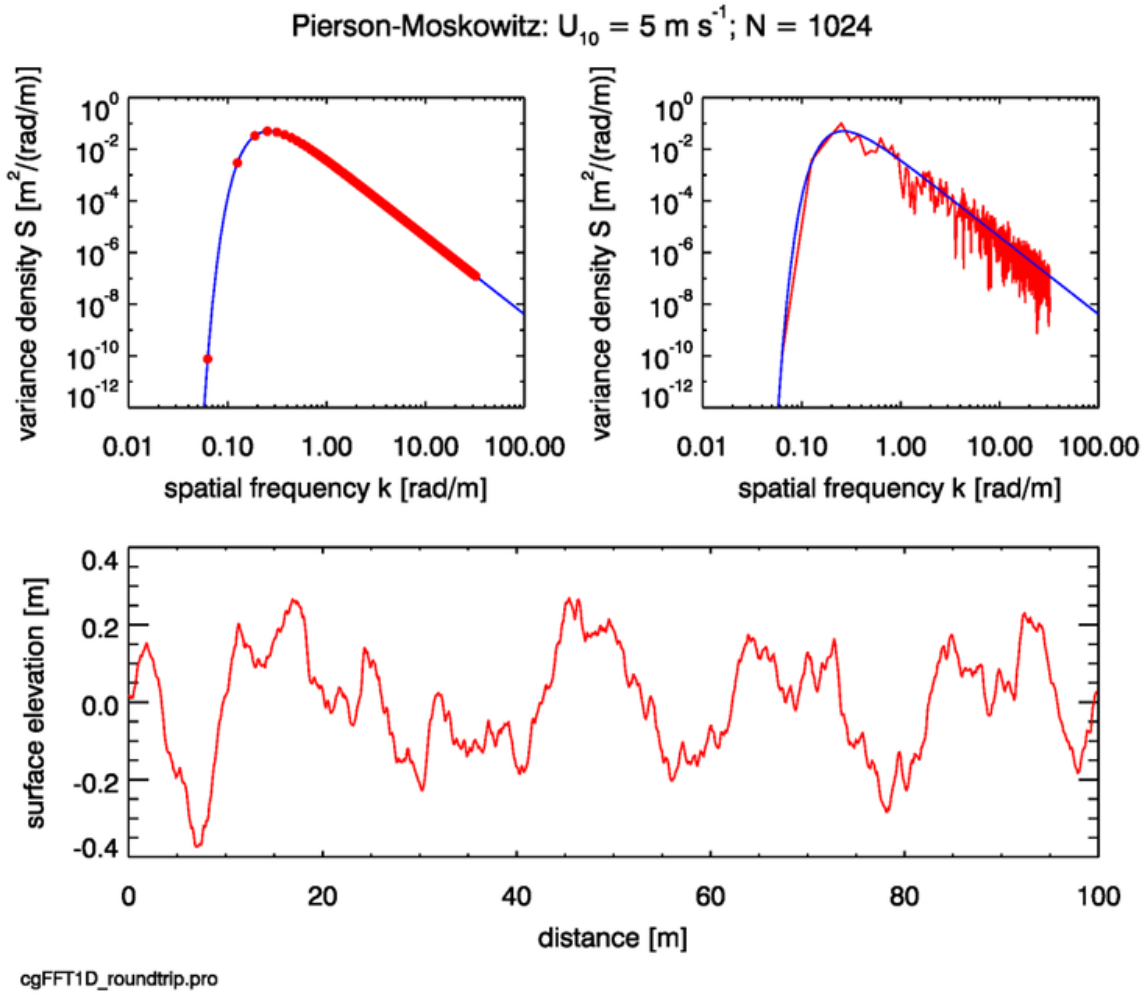


Figure 1: Example of a 1-D random sea surface generated from the Pierson-Moskowitz spectrum for $U_{10} = 5 \text{ m s}^{-1}$ (the wind speed at 10 m above mean sea level). The longest resolvable wave has wavelength $L = 100 \text{ m}$. For $N = 1024$, the two-point wave has wavelength $2L/N = 0.195 \text{ m}$. This is already less than the smallest wavelength (highest frequency) for which the Pierson-Moskowitz spectrum should be used. Generated by IDL routine `cgFFT1D_roundtrip.pro`.

Fourier Transforms page, is satisfied. For the simulation of Fig. figure1, we have

$$\sum_{r=0}^{N-1} |z(r)|^2 = N \sum_{u=0}^{N-1} |\hat{z}(u)|^2 = 19.395 \text{ m}^2 .$$

There are sometimes other checks that can be made. For example, the Pierson-Moskowitz spectrum is simple enough that it can be analytically integrated over all frequencies. This gives

$$\langle z^2 \rangle = \int_0^\infty \mathcal{S}_{\text{PM}}(k) dk = 3.04 \cdot 10^{-3} \frac{U_{10}^4}{g^2} , \quad (5)$$

where we have recalled that the variance spectral density is related to the variance of the sea surface. The variance of the generated zero-mean sea surface can be computed from

$$\text{var}(z) = \frac{1}{N} \sum_{r=0}^N z^2(x_r) \quad (6)$$

and compared with the analytical expectation. For the surface seen in Fig. figure1, Eq. (equation6) gives $\text{var}\{z\} = 0.0189 \text{ m}^2$ vs. the theoretical value of $\langle z^2 \rangle = 0.0197 \text{ m}^2$ from Eq. (equation5). This agreement to within a small amount of random noise indicates that all is probably well with the calculations. Indeed, the average $\text{var}(z) \pm$ one standard deviation for 100 independent simulations is 0.020 ± 0.007 , which agrees well with the theoretical value.

The significant wave height $H_{1/3}$ is by definition the height (trough-to-crest distance) of the highest one-third of the waves. To a good approximation, this is related to the expectation of the variance by

$$H_{1/3} = 4\sqrt{\langle z^2 \rangle} .$$

In the present example, this formula gives $H_{1/3} = 0.55 \text{ m}$. The average significant wave height for 100 simulations is 0.56 ± 0.09 (average \pm one standard deviation). If you spend enough time in a sea kayak to develop intuition about wave heights as a function of wind speed, a half-meter height for the largest waves is about right for a 5 m s^{-1} or 10 knot wind speed.

To summarize this page: we have learned how to start with a wave variance spectral density function and generate random discrete Fourier amplitudes. The inverse DFT of those amplitudes gives a random realization of a sea surface that is physically consistent with the chosen variance spectrum. That is all that we can ask of the Fourier transform technique.