

Figure 1: “Sinusoidal” variability sampled at the Nyquist frequency (here 1 cycle/2 seconds), and half the Nyquist frequency (1 cycle/4 seconds).

Spectral Aliasing

Suppose that every day at noon you go outside and measure solar insolation. After a year or so, you could plot a time series that would show the increase in solar insolation in summer and decrease in winter. But since you made your measurements at noon, the mean solar insolation would be quite a bit larger than the true mean. You would have “aliased” the diurnal cycle into the time-averaged mean. In this case, it would be easy to remedy the problem by measuring a few times a day. But if the measurements have already been collected and archived, you might not be able to eliminate aliasing after the fact.

Aliasing would not be a concern if we could measure continuously. In practice although we can measure frequently, we always have discrete measurements rather than continuous measurements. The Nyquist frequency is the highest resolved frequency and is equivalent to one cycle every two data points. Anything that occurs at a frequency faster than that will be aliased.

Let’s look at an example. Imagine that you make measurements at 1 Hz (1 sample per second). The Nyquist frequency will be 1 cycle per 2 seconds (but of course you won’t do a great job sampling at this frequency). Figure 1 shows that although variations at the Nyquist frequency (1 cycle/2 seconds) and half the Nyquist frequency (1 cycle/4 seconds) don’t look very sinusoidal, they are clearly distinguishable. Now what happens when the frequency of the variability exceeds the Nyquist frequency. Dashed lines in Figure 2 show sinusoidal variations corresponding to 1 cycle/3 seconds (blue) and 1 cycle/1.5 seconds (green). Although the curves are distinct, when they are sampled at 1 Hz (solid lines), they are exactly the same. In this case, the variability above the Nyquist frequency is aliased into a frequency below the Nyquist frequency that is able to resolve it.

You can test other cases using Matlab. For example, 1 cycle/1.25 seconds is equivalent to 4 cycles/5 seconds, and it will alias into 1 cycle/5 seconds. Similarly, 1 cycle/1.1 seconds is equivalent to 10 cycles/11 seconds and will alias into 1 cycle per 11 seconds. There’s a simple rule for determining alias frequencies. Take the frequency of the true signal minus the Nyquist frequency, and subtract that from the Nyquist frequency. In other words, $\Delta f = 10 \text{ cycles}/11 \text{ seconds} - 5.5 \text{ cycles}/11 \text{ seconds} = 4.5 \text{ cycles}/11 \text{ seconds}$. Then the aliased frequency is the Nyquist frequency minus $\Delta f = 5.5 \text{ cycles}/11 \text{ seconds} - 4.5 \text{ cycles}/11 \text{ seconds} = 1 \text{ cycle}/11 \text{ seconds}$. Similarly, it’s easy to determine that 11 cycles/20 seconds aliases into 9 cycles/20 seconds.

In other words, if you plot spectra against a linear x-axis in frequency, high frequency variability simply folds back around the Nyquist frequency. This is a good point to remember. Although we often plot spectra on log-log axes, a log-linear axis can be helpful for tracking down the origins of aliased energy.

When spectra are red, you can often assume that spectral energies are substantially higher in the well-sampled frequency range below the Nyquist frequency than they are above the Nyquist frequency. But if spectra are white, then noise that occurs at frequencies above the Nyquist frequency can be expected to




Figure 2: Sinusoidal variability at 1 cycle/3 seconds and 1 cycle/1.5 seconds looks the same when sampled at 1 sample per second.

alias into resolved frequencies. That gives you a good rule of thumb: if spectra are white, then you should probably assume that you're looking at noise or aliased signals.

Finally, let's look at a very real world example. The TOPEX/Poseidon/Jason altimeters pass over the same ground tracks every 9.9156 days. Sea surface height is strongly influenced by the M2 tide, which has a period of 12.4206 hours. What frequency does 12.4206 hours alias into? In this case, the altimeter Nyquist frequency is nowhere near the tidal frequency, and the aliased signal folds back and forth along the x-axis several times. To compute this, we first need to compute frequencies in common units:

```
f_sampling=1/(9.9156); % cycles per day
f_M2=24/12.4206;      % cycles per day
f_Nyquist=f_sampling/2;
M=floor(f_M2/f_Nyquist) % compute the integer ratio of the two frequencies.
                        % In examples above, this was one.
alias = f_M2 - floor(f_M2/f_Nyquist)*f_Nyquist;
1/alias

% Note: if M is odd then reset
if(rem(M,2)~=0) alias=f_Nyquist-alias; end
```

This calculation shows that M is 38 and the tidal energy aliases into a 62.1068 day period. Note that if M is odd, you have to do the calculation slightly differently. As an exercise, you can check what happens if you round off the numbers and use 9.9 for the altimeter sampling frequency and 12.4 for the tidal cycle.