PUV Wave Directional Spectra
Evaluating Your Results

This section explains some of the limitations of the PUV method so that you can better understand its capabilities--it also explores methods you can use to evaluate and control some of its limitations. We present equations for estimating the uncertainty in direction and directional spreading estimates, and review methods you can use to obtain the best estimate of direction during identifiable wave events.

Wave Scaling

Referring back to the PUV Wave Measurement paper, recall that wave velocities and pressures weaken with depth. This weakening, or attenuation, increases with frequency. At some cutoff frequency, the velocity and pressure signals become so weak, that the waves can no longer be detected. This section shows you how the PUV method corrects for this attenuation, and what happens at the cutoff frequency.

Figure 1 shows example raw spectra from the Torrey Pines data sets. We have included spectra from the Aquadopp Current Profiler deployed nearby because there are important differences that are worth understanding. You can learn more about the Aquadopp data in the NortekUSA Waves2001 paper.

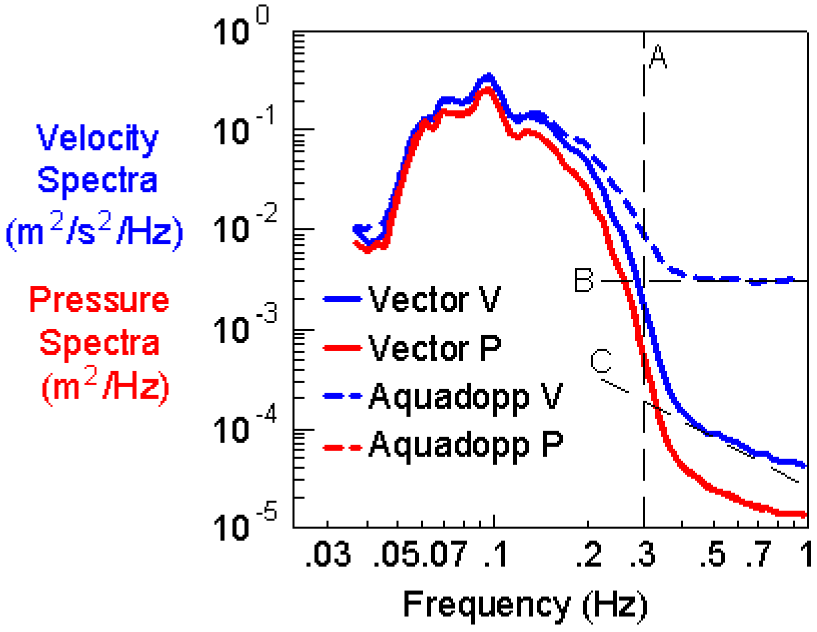

Figure 1. Raw spectra from the Vector Velocimeter and the Aquadopp Current Profiler, both deployed next to each other at Torrey Pines. The dashed lines are:

A) High frequency cutoff
B) Aquadopp velocity noise floor
C) Slope of a turbulent inertial subrange spectrum.

Before we go to the next figure, look closely at the characteristics of the Aquadopp and Vector spectra above the wave cutoff. The flat Aquadopp spectrum at 3 x 10^{-3} $m^2/s^2/Hz$ is the Aquadopp's noise floor. In contrast, the

Vector does not fall to its noise floor, but instead sees the local wave-induced turbulence. These differences are important later when we evaluate directional uncertainties.

The bottom panel of Figure 2 shows how we scale these spectra up to correct for wave attenuation. Wave attenuation is greater at high frequencies than at low frequencies. Around the high frequency cutoff, attenuation increases so quickly that the scale factors literally blow up. You can find the equations for these scale factors in the Waves2001 paper. The top panel shows the wave spectra that result after applying the scale factors to the raw spectra. Real wave spectra fall off at high frequencies, but you can see the opposite happening here--the spectra blow up at high frequencies because of the rapid rise of the scale factors. The only choice is to cut the spectra off above where they run into trouble.

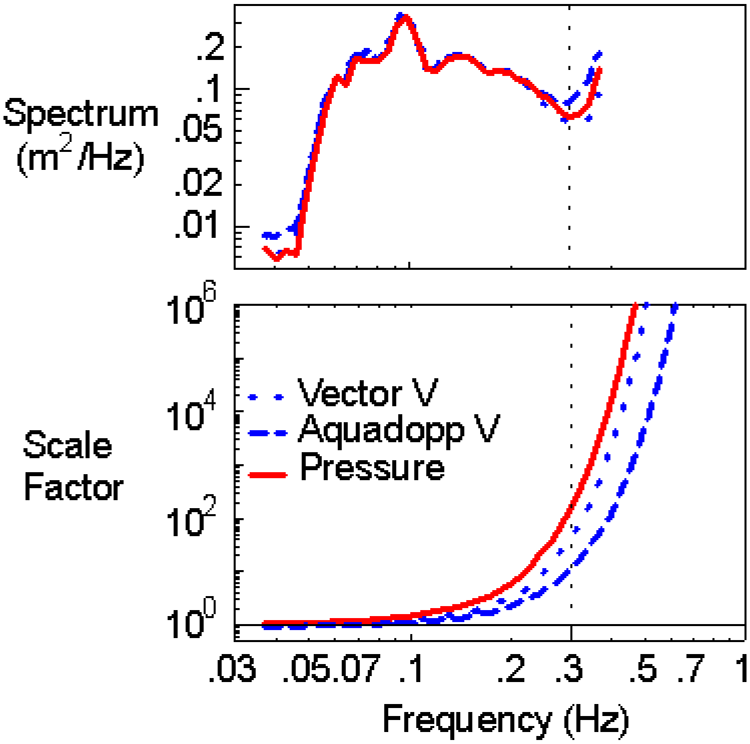

Figure 2. Bottom panel: scale factors that convert velocity and pressure spectra to surface wave spectra. The velocity scale factor has units s^2 *and the pressure scale factor is dimensionless.*

Top panel: wave spectra showing how the spectra blow up past the cutoff.

One sensible approach to cutting off the spectra might be to look for the local minimum at the high frequency end of the spectrum, cutting off everything above the minimum. Another approach would be to cut the spectrum off when the scale factors grow above some value. The scale factors blow up so quickly, however, that it does not really make much difference how you choose to cut off the spectrum.

Figure 3 shows the result of a simple calculation in which the cutoff is computed by limiting the pressure scale factor to the value 200 (assuming the pressure sensor is located 1 m above the bottom). It shows how the cutoff period increases as the water depth increases.

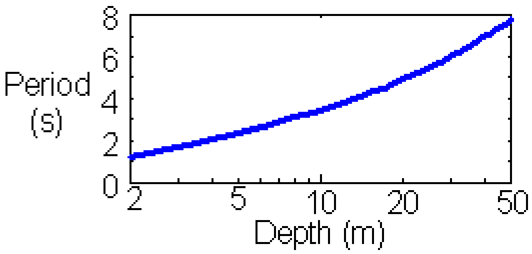

Figure 3. Spectral cutoff period vs. depth.

Scaling Sanity Check

To get the wave spectrum scaled correctly, you must be sure that you have all of the following depths and elevations correct:

1. Mean water depth (keeping track of the tide)
2. Elevation of the pressure sensor above the bottom
3. Elevation of the velocity measurement above the bottom

In most instruments, the positions of the velocity and pressure measurements are fixed relative to each other, so if you know the elevation of the pressure sensor, you will know the elevation of the velocity measurement. On the other hand, one must take care to determine the exact elevation of the instrument above the sea bed, and one must be careful about offsets in the pressure sensor.

Figure 4 illustrates the effect of pressure sensor offsets.

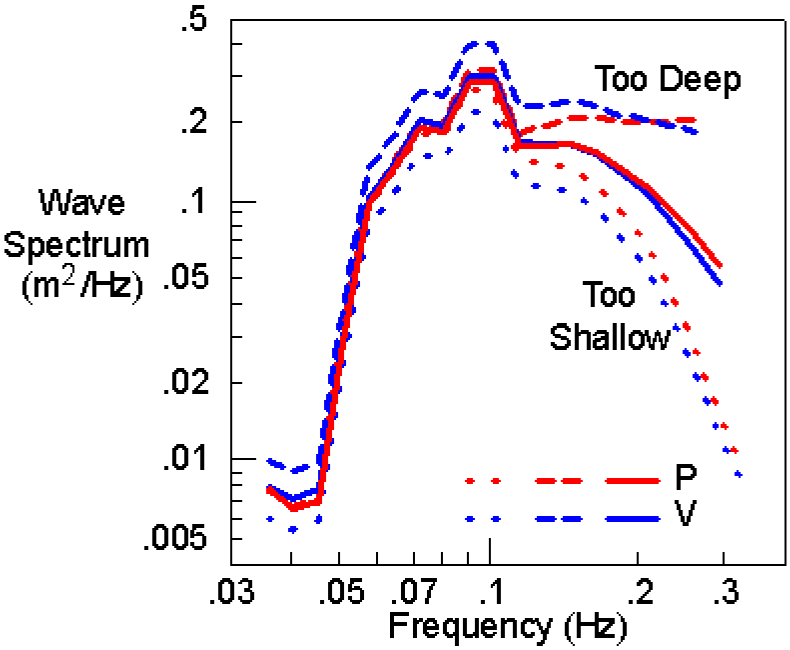

Figure 4. Consequences of pressure offsets. The solid lines show how the wave spectra based on velocity and pressure line up with one another when pressure is measured accurately. The dashed and dotted lines result if the pressure sensor were to read 2 m too deep or 2 m too shallow.

When the pressure and other depths are correct, the wave spectra based on both velocity and pressure line up with each other. Figure 4 shows that pressure offsets not only change the spectra, they make the pressure and velocity spectra different from each other.

Event Overlap

Figure 5 shows a close-up of where two events (Events 2 and 3 from Wave Measurement paper, Figure 6) overlap.

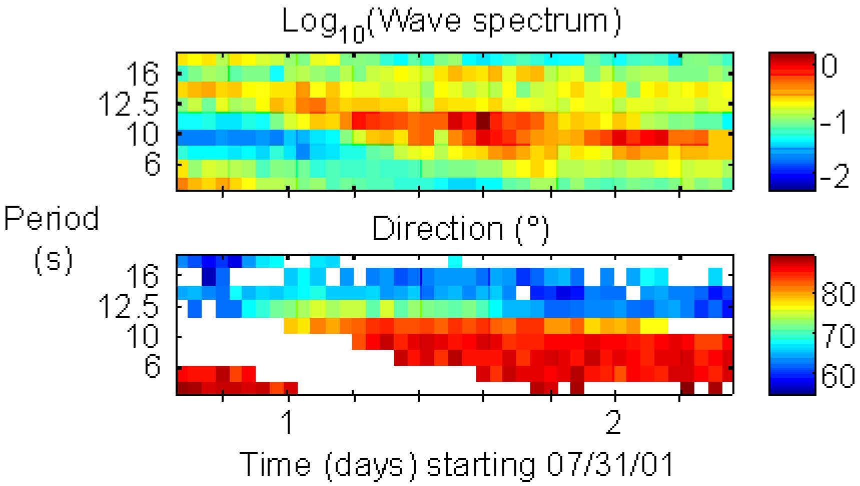

Figure 5. Close-up of where two wave events overlap. The overlap is most visible in the directions.

To understand what happens where the overlap occurs, we will identify four different areas and measure the average wave direction in each. Figure 6 defines these four areas.

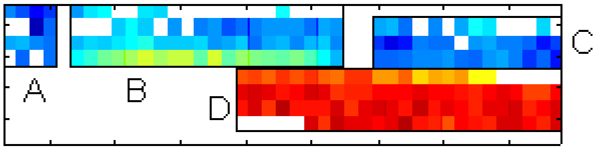

Figure 6. Four areas around where the two events overlap.

Of the four areas defined in Figure 6, areas A and C are entirely from Event 3, area D is entirely from Event 2, and area B is where the two events overlap. Table 1 shows that the directions in areas A and C are within 1° of each other. Given the statistical uncertainty of the mean direction in these areas (see discussion below), 1° is not different from zero. The direction in area C is 84°. The direction in area B where the two events overlap, 67°, is between 62° and 84°, and different from both 62° and 84° with statistical confidence.

Area	Mean Direction
A	62
B	67
C	63
D	84

Table 1.

Measurement Errors and Uncertainties

Measurements are estimates of the the value of something real. Given a real wave direction, each measurement is an estimate of this direction. If you happened to be in waves with infintely-long, parallel crests (i.e. the kind you find most often in Waves 101 classes!), wave direction is easy to define--it is perpendicular to the wave fronts. However, real waves are rarely so simple. At any given time, wave spreading blurs the wave direction, making the real wave direction meaningful only as an average.

Wave direction measurements are similar. Each direction measurement is an estimate of the mean wave direction, but if you make many independent estimates, they differ from one another. Averaging many independent wave direction measurements enables you to get a better estimate of the actual mean direction of the waves you are observing.

Unbiased vs. biased estimators. Averaging estimates together usually improves your measurement. If your estimator is unbiased, then the more you average, the better your estimate becomes. Some estimators are biased, however. No matter how much you average a biased estimator, you will always have a residual error, that is, a residual difference between your mean estimate and what you are measuring. Even so, there is value in averaging biased estimates, for two reasons. 1) The bias is often smaller, or even *much* smaller, than the random errors you can remove with averaging. 2) If you understand the characteristics of the bias, you can correct a biased estimator to obtain a better estimate.

Uncertainties: Direction and Directional Spreading

Equations for the uncertainty in direction and the directional spreading are in the Waves2001 paper so we will not repeat them here. There are three primary factors in the uncertainty:

1. The actual directional spread of the waves themselves. Uncertainty in the mean wave direction is proportional to the spreading of the waves.
2. SNR or signal/noise ratio. A noisy measurement increases the apparent spreading and the uncertainty of the measurement.
3. Averaging. Like most estimators, averaging produces more accurate estimates.

The directional estimator is unbiased, so averaging should always reduce the uncertainty. In contrast, the spreading estimator is biased. Averaging still helps, but you will always have a residual bias, the magnitude of which depends on the amount of spreading. The uncertainty model in the Waves2001 paper should enable you to remove this bias, at least to first order.

How to Measure Direction Accurately

Keep in mind that *accuracy* means different things to different people, because different people have different requirements. In this section, we will consider a specific wave generation event somewhere in the ocean, and look at how to obtain the best estimate of the mean direction of the waves coming from this event. Hopefully, this discussion will assist your evaluation when you have other measurement objectives.

When we consider a wave event, we will assume the following:

1. Waves come from one direction originating from a single source; we are not mixing together waves from different sources.
2. Wave direction is stationary; it does not change with time.
3. Wave spreading is also stationary.

If waves arrive from two events at the same time, we can consider them separate if they do not overlap in frequency. Events 1 and 3 are simultaneous, but well separated in frequency, so they are separate events. Events 2 and 3 are mostly separate, but we have to be more careful where they overlap.

The best estimate of wave direction in this situation comes by averaging direction over the full extent of the event. Wave direction algorithms typically provide direction estimates over a range of frequencies, where each estimate is independent of the others. Furthermore, when you make repeat measurements periodically, you normally see waves from the same event repeatedly, which allows you to average over time in addition to frequency. The best estimate of the wave direction comes from averaging all of your direction estimates across the frequencies and the times occupied by the event.

The key to getting the most accurate direction measurement is to maximize the number of *degrees of freedom,* or DOF. One DOF represents one independent estimate. You get twice as many DOF when you make 17-minute measurements every hour instead of every 2 hours. A 34-minute measurement every hour also gives you twice as many DOF as a 17-minute average every two hours.

Sampling faster does *not* increase your DOF. Sampling theory tells us that you cannot resolve what you are measuring unless you sample at least twice as fast as the highest frequency in your signal. Common practical advice tells us that sampling four times as fast is even better because it minimizes aliasing.

However, both the Aquadopp and Vector sample continuously during each measurement interval, reporting the average. This kind of sampling effectively suppresses aliasing, so there is no advantage in faster sampling.

Estimating Expected Uncertainty

Figure 7, taken from the Waves2001 paper, shows how uncertainties in direction and spreading vary with the actual spread of the waves and with the SNR. The figure is based on 2 DOF, standard deviations can be reduced proportional to DOF$^{-1/2}$.

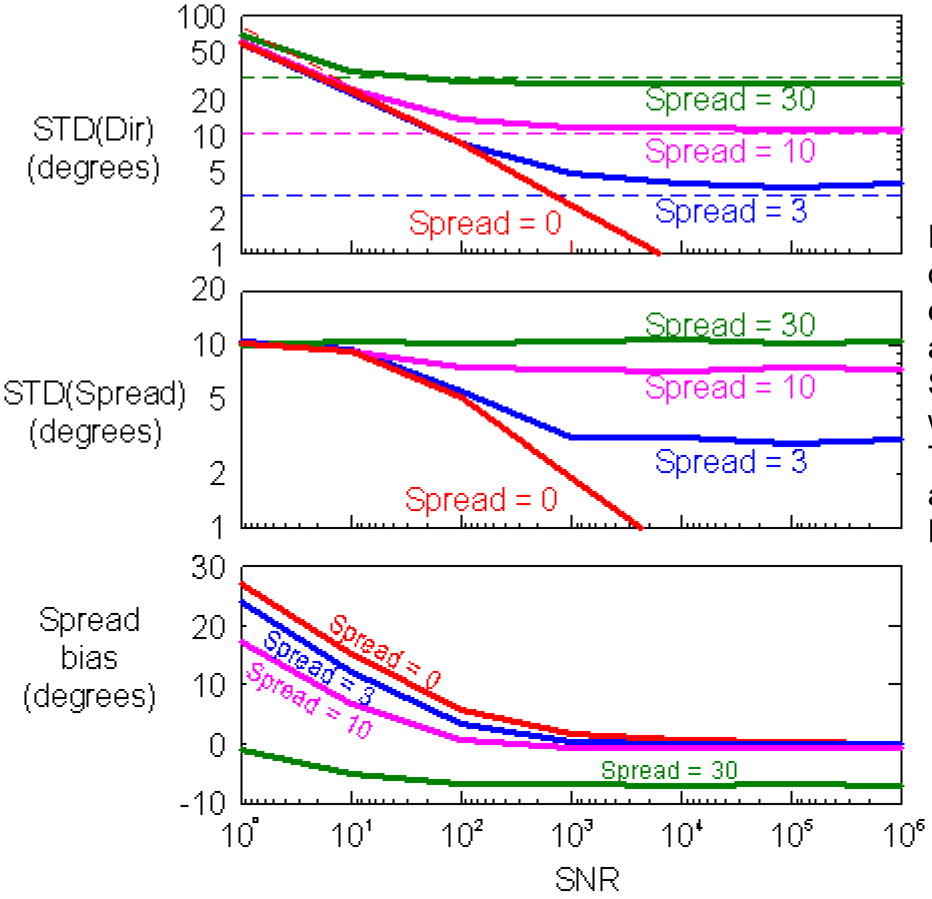

Figure 7. Standard deviations of estimated direction and spreading, vs. SNR and the real wave spreading. These curves are all based on 2 DOF.

You can estimate expected standard deviations by reading Figure 7, or with equations given in the Waves2001 paper. First, you must estimate the real spreading. If the spreading is less than, say, 30°, and if your SNR is larger than 100, then it is reasonably easy to use Figure 7 to estimate the real spreading based on the estimated spreading (see also Figure 8, below). In addition to the real spreading and the SNR, you also need to know the number of DOF. The following addresses estimating SNR and DOF.

Estimating SNR

The most direct way to estimate SNR is to create spectra like the ones in Figure 1. You must average the spectra sufficiently to allow the noise floor to emerge from the randomness. Once you know the noise level, then you can compute the average spectrum in the wave band you are interested in, and compute the SNR. For example, in Figure 1, the noise level of the Aquadopp is 3×10^{-3} $m^2/s^2/Hz$. An eyeball average at the peak of the spectrum is around 2×10^{-1} $m^2/s^2/Hz$. The ratio is 60. Once you define the limits of an event in time and frequency, you can also compute the average spectrum within the event. Then compute SNR as the ratio of the average spectrum to the noise level.

Estimating DOF

Fourier transforms create spectra in a series of evenly spaced fundamental frequency bands each of which has 2 DOF. Estimates of direction are typically computed from averages across a number of these fundamental frequency bands and the total number of DOF in a single estimate is twice the number of fundamental frequencies.

The events identified in Figure 6 of the Wave Measurement paper include a large number of similar direction estimates. Once you identify the range, both in time and frequency, of a stationary event, you can obtain an even better estimate of the mean direction by averaging all of the estimates together. The total number of DOF of the large average is the sum of all of the DOF in each of the smaller averages.

Direction Bias

There are several sources for directional bias, but the direction estimator itself is not one of them. Common sources of bias include compass error, obstructions or interference around the velocity measurement volume and drifts or miscalibrated velocity sensors. Aquadopp and Vector compasses are good to around 2°, so the compass's direction uncertainty is larger than these instruments' ability to resolve waves. Relative to other sensors, the Aquadopp and Vector are much less sensitive to the other sources of bias listed above. Doppler sensors hold their calibration, so calibration drifts do not bias direction.

A significant advantage of both the Aquadopp and the Vector is that they measure remotely. Remote measurement makes it is easy to mount the instruments to avoid or eliminate flow disturbance. Oscillatory surface waves often cause an instrument's wake to return repeatedly to the instrument. Instruments that measure velocity in a volume within the instrument, or immediately next to the instrument, suffer from self interference--they disturb the flow as they try to measure it. This self interference is a source of direction bias that can be avoided with remote measurements.

Spread Bias

The spreading estimator is biased, and the bias increases with the spreading. Figure 8 shows how the bias varies with real spread. At high SNR and with real spreading less than 10°, the bias reduces the estimated spreading by less than 1° relative to the real spreading, but at 30° the bias is -7°. Noise (low SNR) causes a positive bias, particularly for small spreading.

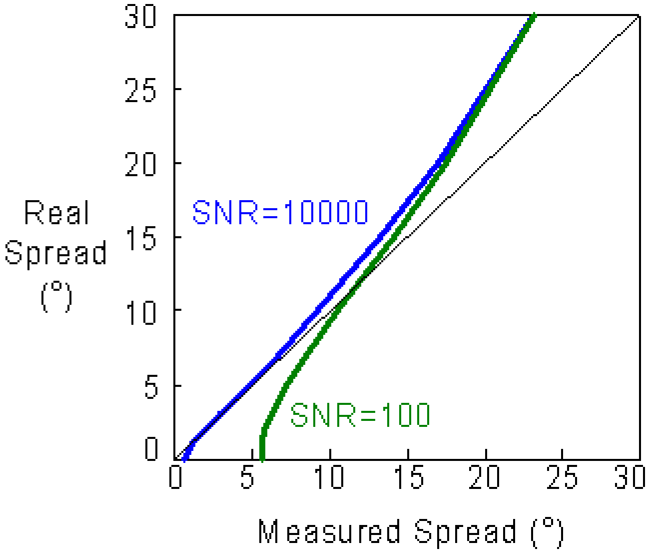

Figure 8. Bias in the spreading estimator vs. the true spreading, for varying SNR.