TELEDYNE
RD INSTRUMENTS
A Teledyne Technologies Company

WAVES PRIMER:

Wave Measurements and the RDI ADCP Waves Array Technique

Teledyne RD Instruments Inc.
14020 Stowe Drive
Poway CA 92064
Tel: +1 858-842-2600
Fax: +1 858-842-2822
Email: rdisales@teledyne.com for sales related questions
Email: rdifs@teledyne.com for support related questions
Web: www.rdinstruments.com

Introduction:

Acoustic Doppler Current Profilers (ADCPs) gather profiles of water velocity by measuring the Doppler shift of sound reflected from scatterers assumed to be passively following the flow. The measurements are range-gated into a series of bins along three or more beams and then combined to infer the velocity profile encompassed by the beams. This technology is well-proven, and ADCPs are now routinely deployed around the world.

Directional wave measurements, by whatever technique, seek to statistically describe basic wave parameters in terms of the wave amplitude, period and direction. The wave amplitude measurement most commonly employed is known as the Significant Wave Height, H_s; loosely considered to be the average peak-to-peak amplitude of the largest one third of the waves seen during the measurement interval. The peak period, T_p, tells the characteristic frequency of the arriving wave energy (frequency is the inverse of the period). The mean direction of the waves, D_p, tells which way the waves are propagating. To summarize, directional wave measurements tell how much wave energy exists (H_s), at what frequency (T_p) and from what direction (D_p).

Wave Basics

It is perhaps worthwhile to go over some of the fundamental behaviors of ocean waves. To make the process as simple as possible, consider a single wave propagating in the *x* direction. The equations that describe this wave are presented without derivation (see Kundu's 1990 text for an excellent presentation of the derivation):

$$u = a\omega \frac{\cosh k(z+H)}{\sinh kH} \cos(kx - \omega t) \tag{1}$$

$$w = a\omega \frac{\sinh k(z+H)}{\sinh kH} \sin(kx - \omega t) \tag{2}$$

$$\omega = \sqrt{gk \tanh kH} \tag{3}$$

Where: u is the horizontal velocity of a parcel of water
w is the vertical velocity of a parcel of water
a is the amplitude of the wave at the surface
ω is the angular frequency of the wave (which is $2\pi/\tau$ where τ is the period)
k is the wave number of the wave (which is $2\pi/\lambda$ where λ is the wavelength)
H is the water depth
g is acceleration due to gravity

Several important features of ocean waves can be seen from a careful inspection of these equations:

- Equations (1) and (2) show that the amplitude of the velocity fluctuations of a parcel affected by the passage of the wave depends on its depth (z), the total water depth (H), angular frequency (ω), and wave number (k).
- Equations (1) and (2) also show that the amplitude of the horizontal velocity fluctuations (u) is subject to different constraints than the amplitude of the vertical velocity fluctuations (w).

- Equation (3) shows that the angular frequency (ω) is related to both the wave number (k) and the water depth (H). Frequency generally does not change, so a wave of a given frequency that propagates into shallower water will change its wave number, and hence its wavelength.

Some of the interplay between these parameters can be seen with an illustration of a wave propagating into shallower water. Waves are often characterized by their period, and an eight second wave is shown.

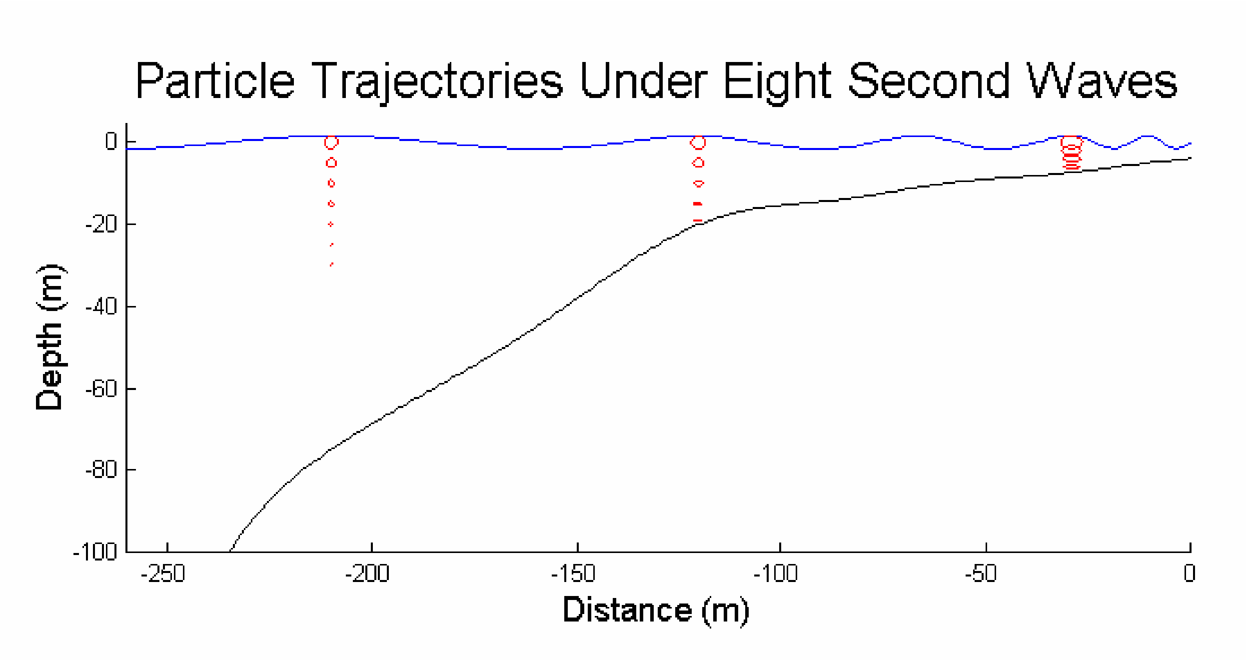

Fig. 1: A two meter wave with an eight second period is shown propagating into shallower water. Each red trace represents a particle's path as the wave passes. Three cases are shown: deep water, intermediate water and shallow water.

In deeper water, the path followed by any parcel of water affected by the wave is a circle whose diameter decreases with depth. This is important because it means the wave energy only propagates to some finite depth, beneath which it can not be seen (or measured). As the wave begins to "feel" the bottom the vertical velocity of the parcel attenuates more rapidly with depth than does the horizontal velocity until, as shown in the shallowest example, the vertical attenuation is so much stronger than the horizontal attenuation that the parcels near the bottom track an entirely horizontal path. It is very important that this depth dependent behavior of waves be understood because the intent is to measure the wave's characteristic properties with a subsurface instrument. Note that the wavelength shortens as the wave propagates into shallower water.

Now consider a shorter period wave:

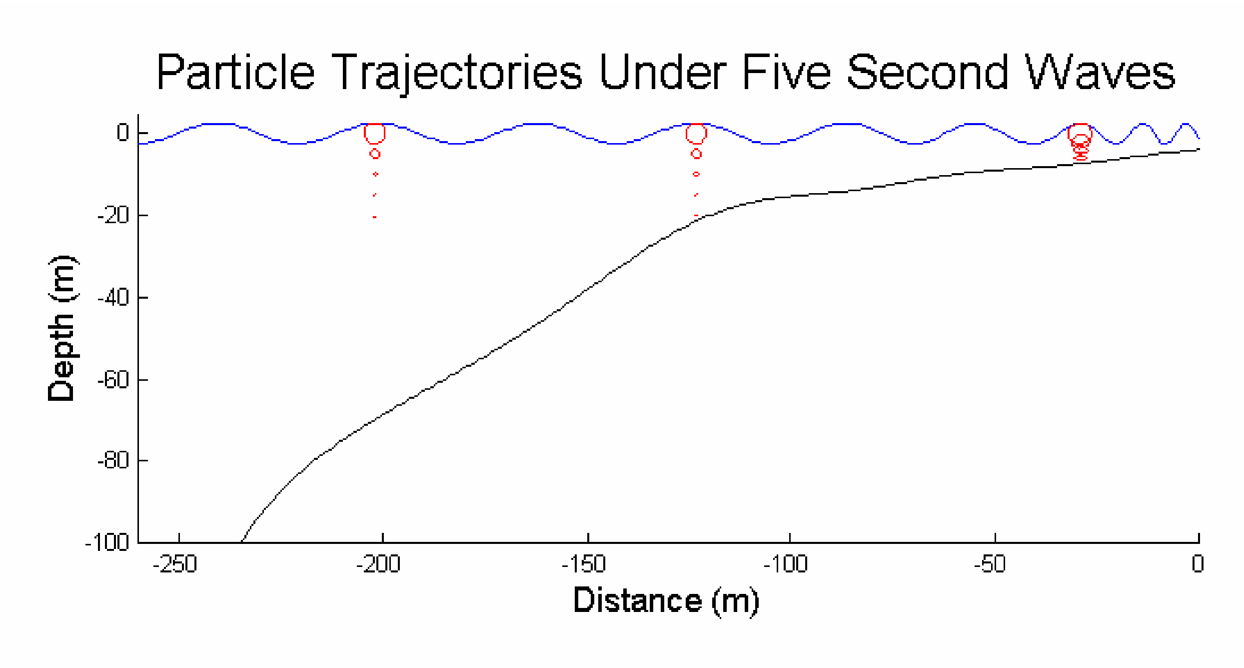

Fig. 2: A two meter wave with five second period is shown identically as in Fig. 1.

Notice that the five second wave behaves substantially differently than the eight second wave. First of all, the amplitude of the parcel excursions decreases with depth more quickly than the eight second wave. This is very important because it shows that higher frequency waves do not penetrate as deeply into the water as lower frequency waves, which in turn means that the frequency of waves that can be measured depends on the depth of the measurement. Because this five second wave is higher frequency than the eight second wave in *Fig. 1* the five second wave is not yet "seeing" the bottom at 20 m depth, and the particle trajectories there are identical to those in deeper water.

The approximate decay of measurable wave periods with depth is shown in *Fig. 3*.

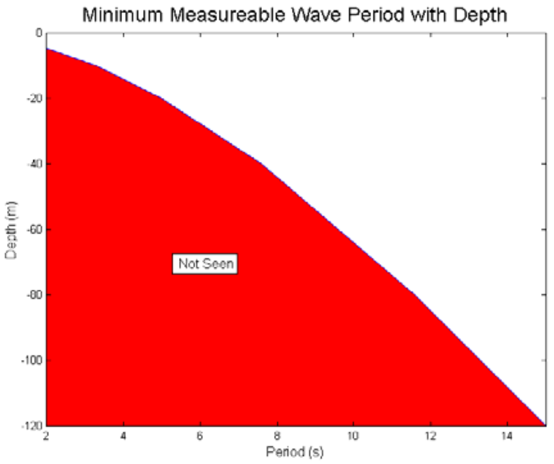

Fig. 3: The approximate decay of wave energy with depth. For example, waves of period shorter than five seconds do not penetrate below 20 m depth.

A Statistical Description

Unfortunately, the real ocean is rarely if ever made up of a single wave. Almost any sea that is carefully watched can be maddeningly irregular. However, there are some very powerful tools readily available if a few basic assumptions are postulated. First is the assumption that the wave field can be described as the summation of sine waves of varying frequency, amplitude and direction. Second is the assumption that the field is statistically stationary – which means that the statistical description of the waves at a given time is essentially the same description that would be obtained at a slightly different time.

Joseph Fourier was a French mathematician who first proposed (in 1807) the idea that any periodic function could be represented as a summation of sine waves. While this idea was initially greeted with outrage, it is safe to say that his ideas have since become one of the primary tools of the physical sciences, and a full exposition is far beyond the scope if this primer. Suffice to say here that Fourier analysis allows a means to reduce a measured time series to a few constituent sine waves, which reveal the dominant frequency components of the process being measured.

Power Spectra

In the plots that follow, the time series of water level is shown in the top panel, and its power spectral density plot is shown in the bottom panel. The power spectral density plot shows how the power is distributed between sine waves of varying frequency. If the entire field consists of a single sine wave, then the power spectral density will be sharply peaked around that frequency. The units for Power Spectral Density are decibels (dB), which is a logarithmic scale – so a decrease of three decibels shows a halving of the power. A 60 dB decrease corresponds to a decrease in power by a factor of one million.

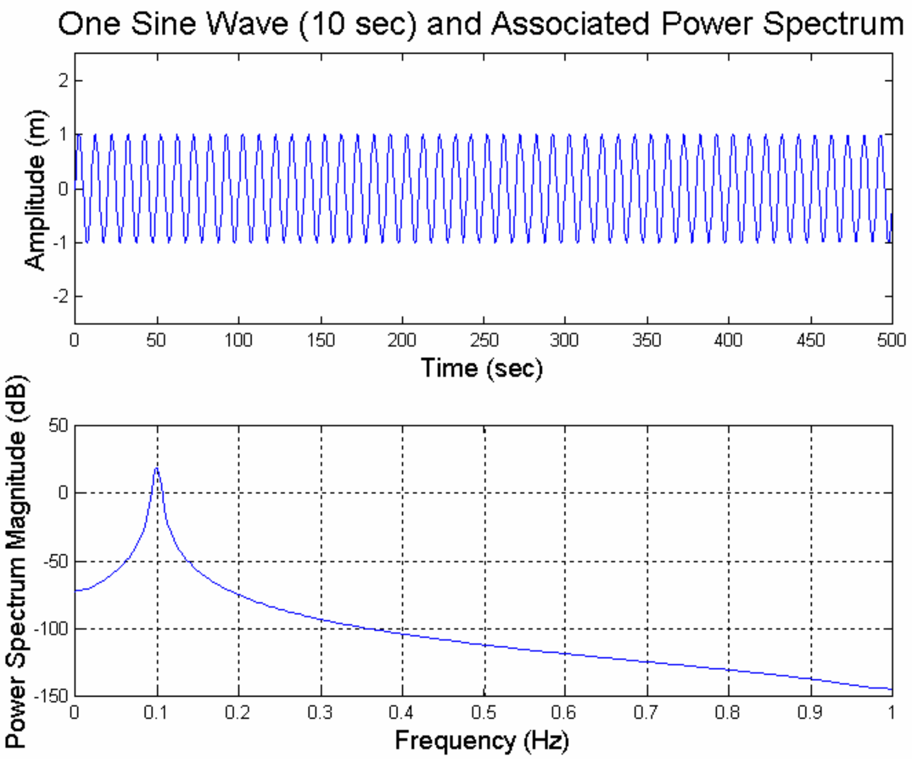

Fig. 4: The power spectral density of a single sine wave with ten second period shows that the energy is concentrated at ten seconds (0.1 Hz).

If the field is made up of two sine waves then the power spectral density will show the energy distributed between the two peaks. If the two sine waves are of equal amplitude, then the two peaks of the power spectral density will be of the same height. It is interesting to note (*Fig. 5*) that the time series created by the addition of two sine waves already seems fairly complicated, while the power spectral density reveals that it was created quite simply.

Side note: In the example shown in *Fig. 5*, the time series was created by summing two sine waves of similar frequency. It can be shown with basic trigonometry that the sum of two sine waves of similar frequency will result in a function which oscillates at the average frequency, but whose amplitude is bounded by a sine wave whose frequency is the difference of the two frequencies. This phenomenon is known as "beating", and it is a very common feature of real ocean waves. Any beach observer knows that waves tend to occur in fairly regular "sets" of large and small amplitude, which is readily attributable to the effect of adding waves of similar frequency as shown in *Fig. 5*.

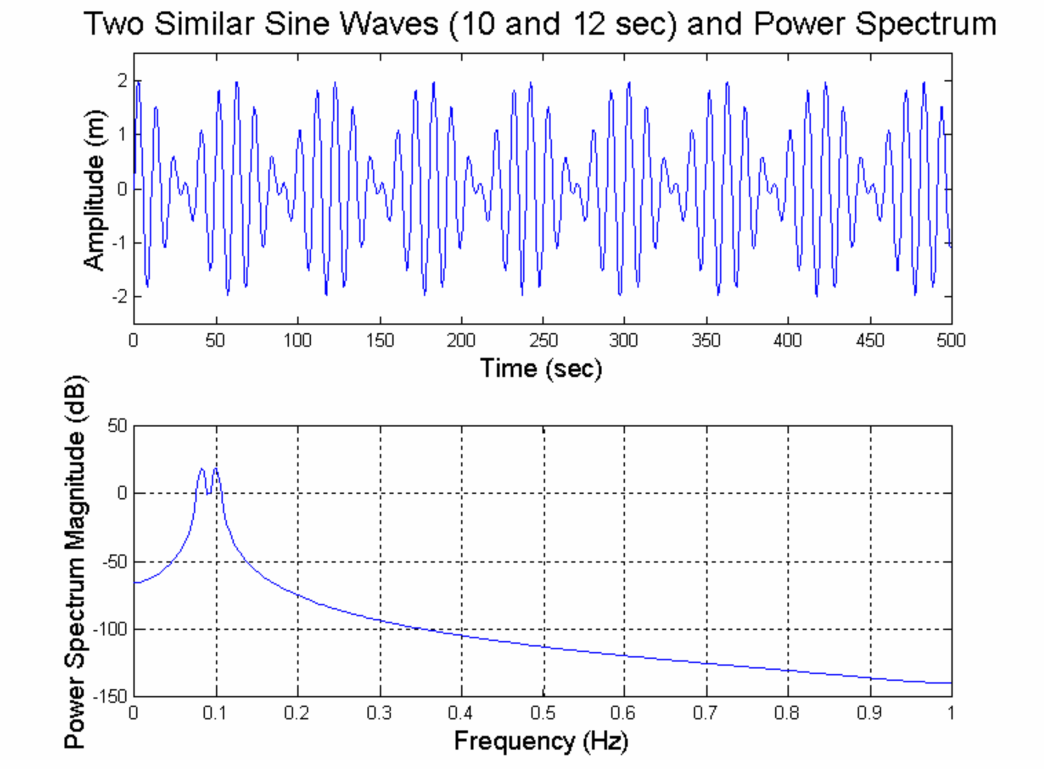

Fig. 5: The power spectral density plot reveals that the complicated time series shown is made up entirely of only two sine waves of similar frequency and equal amplitude.

Now, consider a “real” ocean measurement:

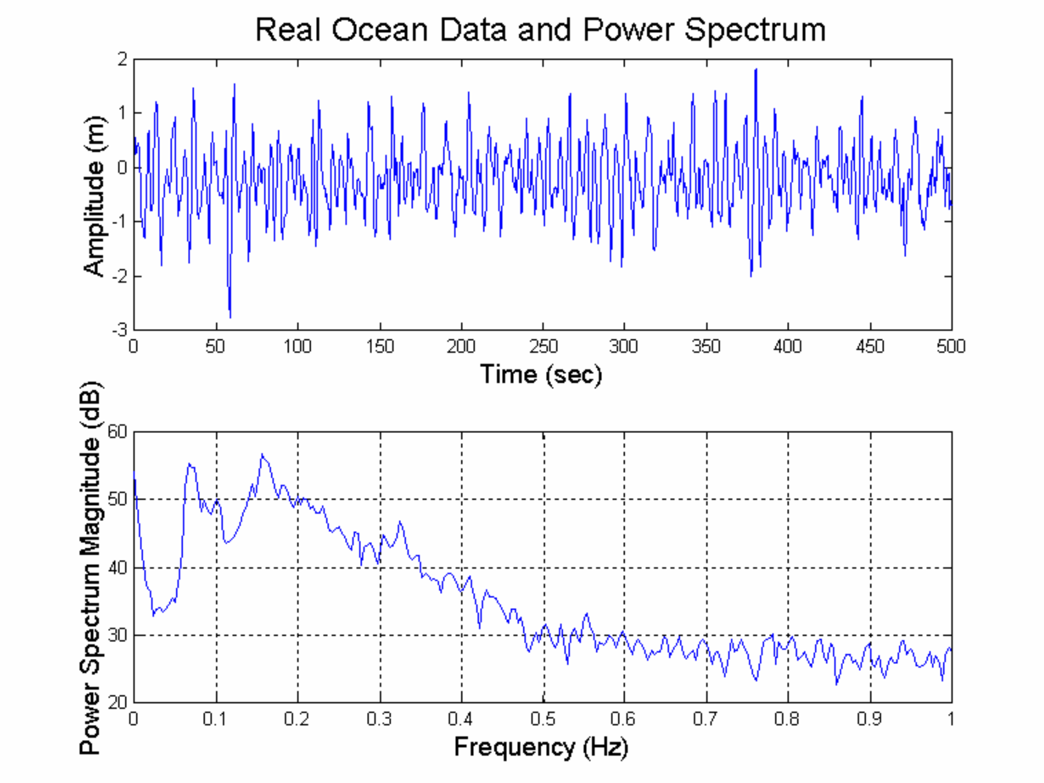

Fig. 6: The power spectral density plot of a real ocean measurement shows that the seemingly hopelessly complicated time series is in fact dominated by only two waves with periods of about 6 and 14 seconds.

The power spectrum of the real ocean contains energy at all frequencies, rather than the artificial constructs shown above, which have all energy concentrated at the set frequencies. This spectrum clearly shows that the bulk of the energy in this very complicated time series is contained in two frequencies of about 0.07 and 0.16 Hz.

Cross Spectra

Cross Spectra provide a relatively straightforward means to compare whether two different measured parameters are varying together. If two different measured parameters are varying at the same frequency, then it is likely that they are related. An obvious example is to compare the water level and velocity of a simple plane wave (*Fig. 7*).

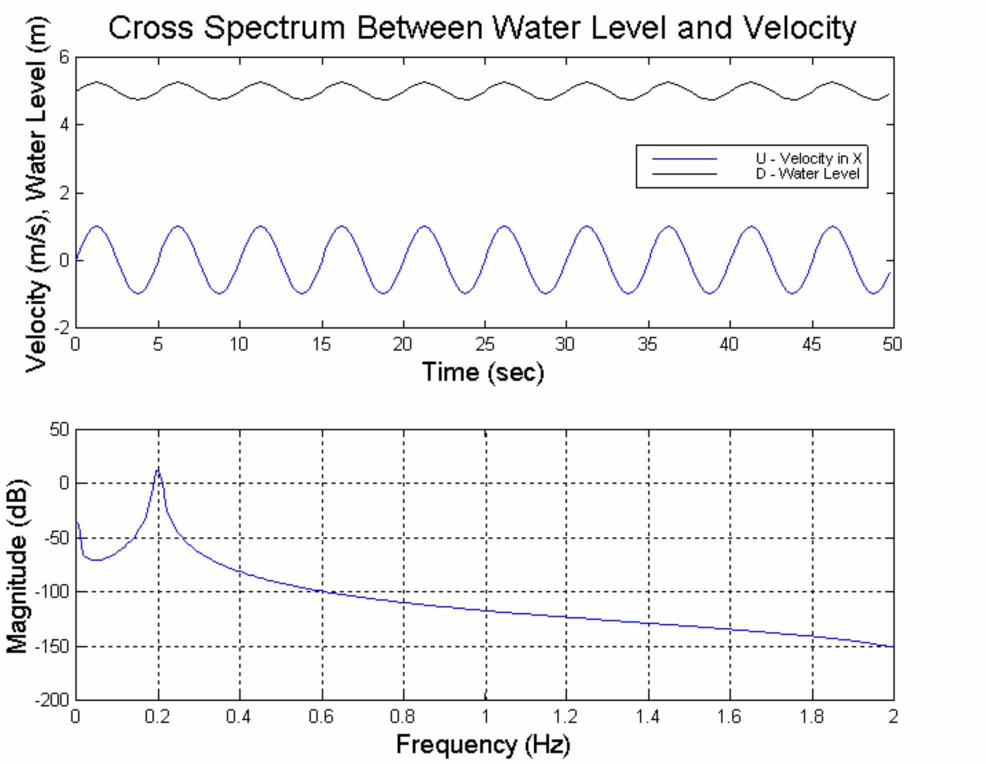

Fig. 7: The time series of water level and horizontal velocity are displayed in the top panel, and their cross spectral density is plotted in the lower panel. It is clear that water level and horizontal velocity are both varying at the same frequency (0.2 Hz)

If we complicate the water level measurement with a second sine wave that is not present in the horizontal velocity (because the second wave is propagating perpendicularly to *U*), then we see (*Fig. 8*) that the cross spectrum of *U* with *D* shows tremendous energy at 0.2 Hz, and only a slight bump at a second frequency.

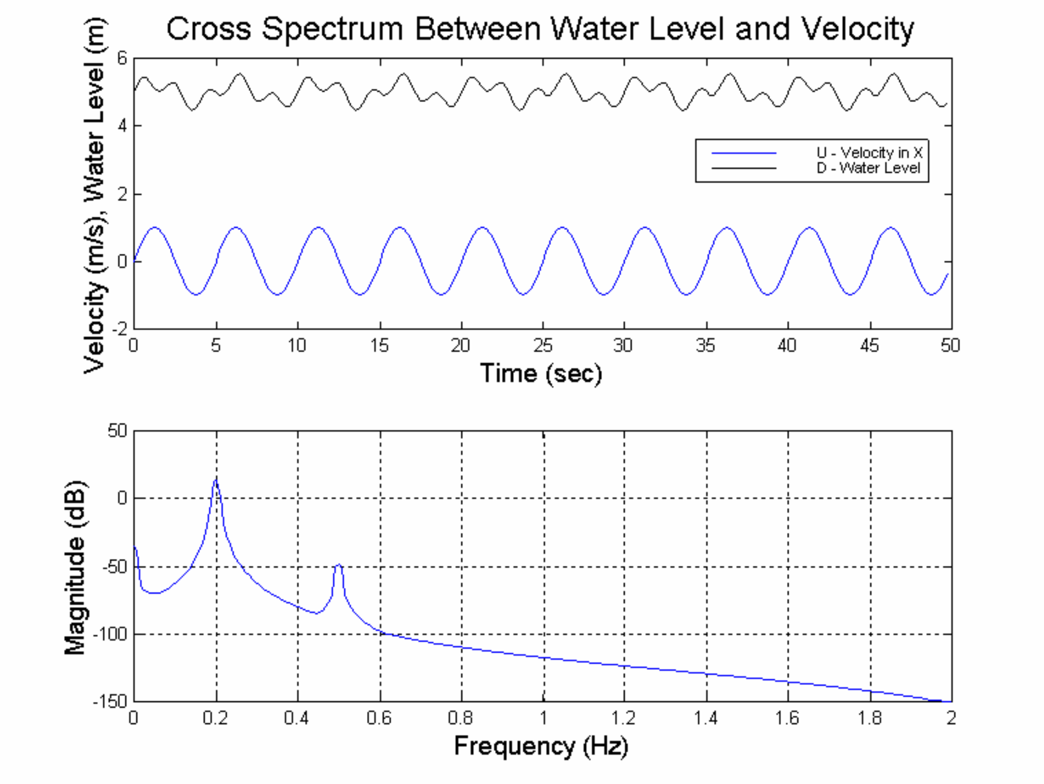

Fig. 8: As in Fig. 7., but a second wave has been added that is propagating perpendicular to the measured velocity.

The slight bump at the other frequency can be neglected because it is so much lower in power (60 dB corresponds to one million times less power). It is present because the modeled water level is made up of several thousand measurements of two simple sine waves with different periods (2 and 5 seconds), so there will be a subset of measurements (every ten seconds) where the water level seems to vary with *U* at the other frequency. That the power in this frequency is so much lower is the tip-off that it is not real – or at least not important.

Wave Parameter Measurements

It is useful to separate a discussion of the wave field into two parts: the parameters that do not depend on the direction of the wave propagation, H_s and T_p; and the parameter that does, D_p.

Non-Directional Wave Parameter Measurement

Peak Period (T_p)

Wave frequency does not change with depth or direction, making peak period the easiest to measure of the three commonly reported wave parameters. A time series of any of a number of parameters that vary with the wave frequency is all that is required. Such parameters include water level, pressure and the orbital velocity of the passing waves. Methods for directly measuring water level include capacitance wave gauges, laser altimeters (looking down in air) and inverted echo sounders (looking up in water). Subsurface instruments can be used to measure pressure and/or orbital velocities, but it is important to remember that the higher frequency waves may not penetrate to the measurement depth. The practical effect of this is that a tradeoff exists between the depth of the measurement location and the ability to measure (or even see) the higher frequency waves. Once the time series of the varying parameter is obtained, its power spectral density function will reveal the frequency with the most energy, and hence the peak period.

Significant Wave Height (H_s)

The time series of the water level can be obtained directly by a number of means, including capacitance wave gauges, laser altimeters (looking down in air) and inverted echo sounders (looking up in water). It is also quite common to measure the pressure and/or orbital velocities from subsurface instruments. However, these parameters decay with depth as described above and therefore must be transferred to their equivalent surface values using equations (1)-(3) for the velocities, and a similar equation for the pressure.

Significant wave height is the average peak-to-peak amplitude of the largest one third of the waves in a given field. This value is roughly equivalent to four times the square root of the value obtained by integrating the non-directional spectrum from the time series of the surface level, and in practice it is usually this integration that is actually reported as the significant wave height. There are a couple of practical issues that arise from calculating the significant wave height in this manner:

1) Spectra from the real ocean are quite noisy, so in practice it is common to band average the frequency spectrum – that is, reduce the spectrum to frequency bands whose reported energy is the mean energy within the band.
2) In any real measurement there is a noise floor, that is, a lower bound that the spectrum asymptotically approaches. This is readily seen in *Fig. 5* where the spectrum is seen to "flatten out" to 25 dB or so at about 0.6 Hz. This noise has several sources, including: that the measurement is not perfect, that the instrument itself will have noise, turbulent fluctuations of the water, etc. This noise floor is usually removed before calculating significant wave height because including it will bias the calculation high.
3) The high and low frequency cutoffs of the spectrum can vary due to the sampling and the noise levels. It is important that the main peak of the characteristic wave energy is located well within the high and low frequency cutoffs or some of the energy is "lost" which will result in an underestimation of the significant wave height. For example, in *Fig. 9* below we show the spectrum from *Fig. 4* (a single sine wave of ten second period). If we set the low frequency end of our spectrum too high, say eleven seconds here, a significant portion of the wave energy is actually at frequencies lower than we include in the integration, and our significant wave height will be underestimated.

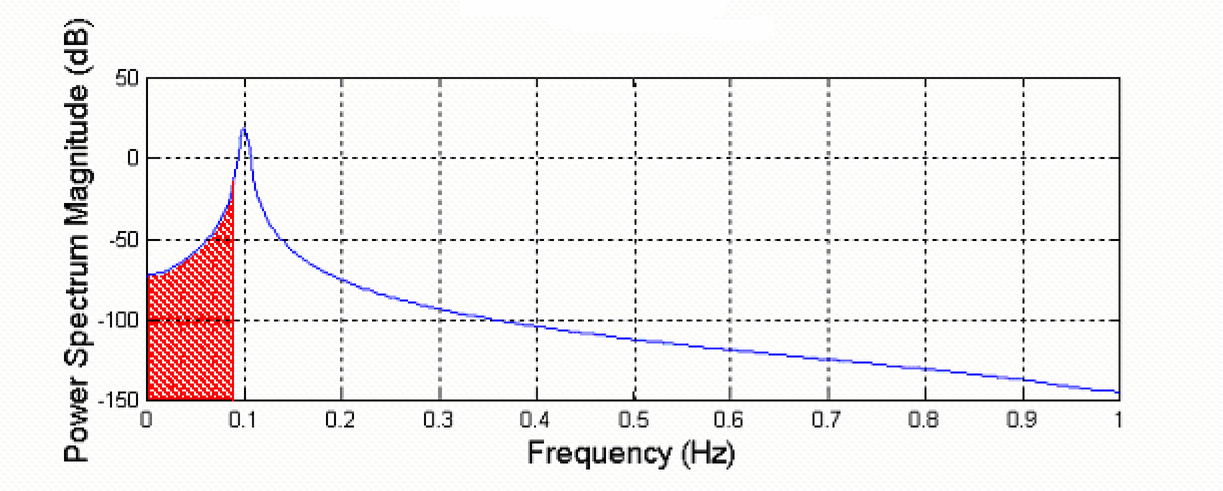

Fig. 9: The power spectral density function as generated in Fig. 4. If we set the low frequency cutoff to eleven seconds, then all of the areas marked with red cross-hatches would be neglected in our significant wave height calculation. Since this spectrum was generated by a wave of ten seconds, quite near our cutoff, excluding the marked area will result in an underestimate of the significant wave height.

The above three points apply to whatever method is used to measure the time series of water level. There is an additional consideration that must be made when the water level measurement is inferred from a subsurface velocity or pressure measurement, and that is the effect of background currents on those measurements. *Fig. 10* shows the same eight second wave as was shown in *Fig. 1*, but an onshore background current of 0.75 cm/s has been added.

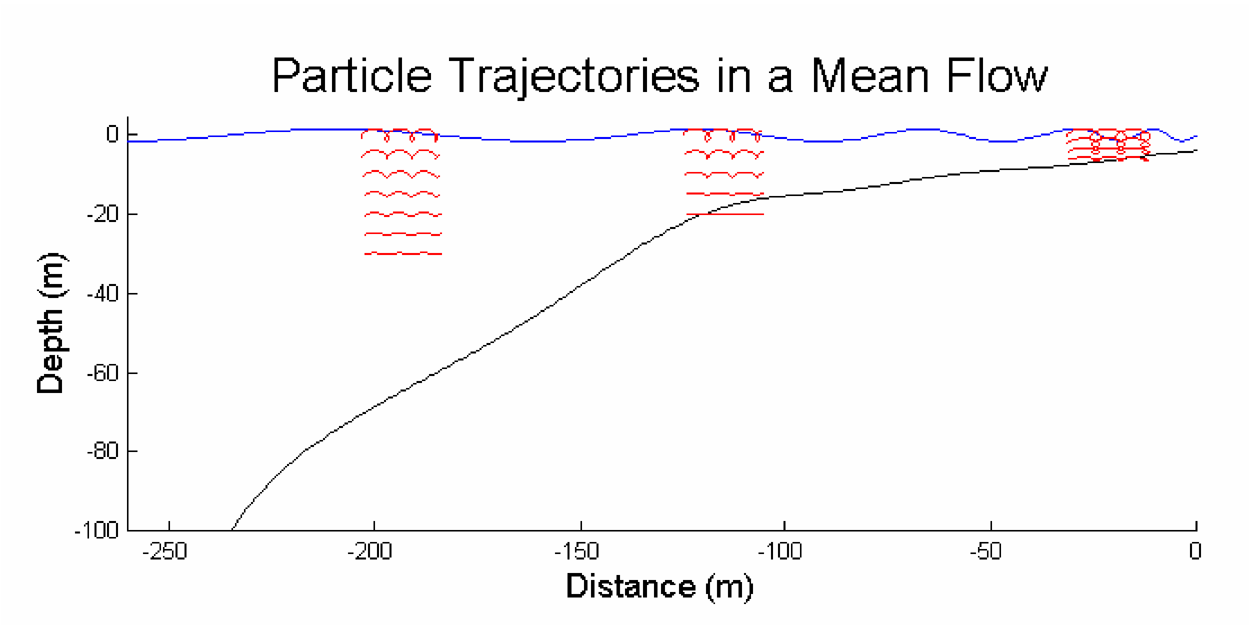

Fig. 10: An eight second wave as in Fig. 1, but with an onshore background current of 0.75 cm/s. Note the substantial elongation of the particle paths in comparison to Fig. 1.

The background current results in a substantial elongation of the particle paths. Any subsurface measurement of velocity (or pressure – any change in velocity is accompanied by a proportional change in pressure) will be larger than if the background current were not there. The subsurface measurement simply reports the total velocity or pressure; it is incapable of distinguishing the changes due to waves from other forcing. Since the subsurface measurements must be transferred to a surface level for wave height estimation, these artificially large measurements will result in an artificially large measurement of the water level, and therefore a reported significant wave height that is too high. The transfer to the surface is specifically correcting for an assumed decay in the velocity, so deeper measurements will be amplified more than shallower measurements. Hence, deeper measurements that do not correct for the presence of background currents will report higher significant wave heights than shallower measurements in the same background current.

This effect can be removed by including the effects of background currents in the dispersion relation (equation (3)) as follows:

$$\omega - kU\cos\alpha = \sqrt{gk\tanh kH} \tag{4}$$

where U is the vertically weighted magnitude of the background current
α is the angle between the background current and the wave direction
all other variables are as in equation (3)

Note that equation (3) is recovered identically if the background current is zero. It is also worth noting that the appropriate velocity is the vertically weighted velocity, which requires a profile of the velocity through the water column.

Directional Wave Parameter Measurement

It is common practice to separate the directional wave spectrum into a frequency spectrum and a directional spreading function, as follows:

$$DW(\omega,\theta) = S(\omega)D(\omega,\theta) \tag{5}$$

The frequency spectrum $S(\omega)$ is the same non-directional spectrum used to determine parameters like the significant wave height and peak period as described in the previous section. There are a number of ways to estimate the directional spreading function, $D(\omega,\theta)$, most of which involve comparing the cross spectra between measurements of some number of independent parameters.

It is very common to compare measurements of pressure and horizontal velocity to determine the wave direction. This is known as the PUV method, which really boils down to answering a simple question: "Does the variation in water level (P) vary more with the velocity in the *x* direction (*U*) or the velocity in the *y* direction (*V*)?" Consider the case where the wave is propagating entirely in the *x* direction:

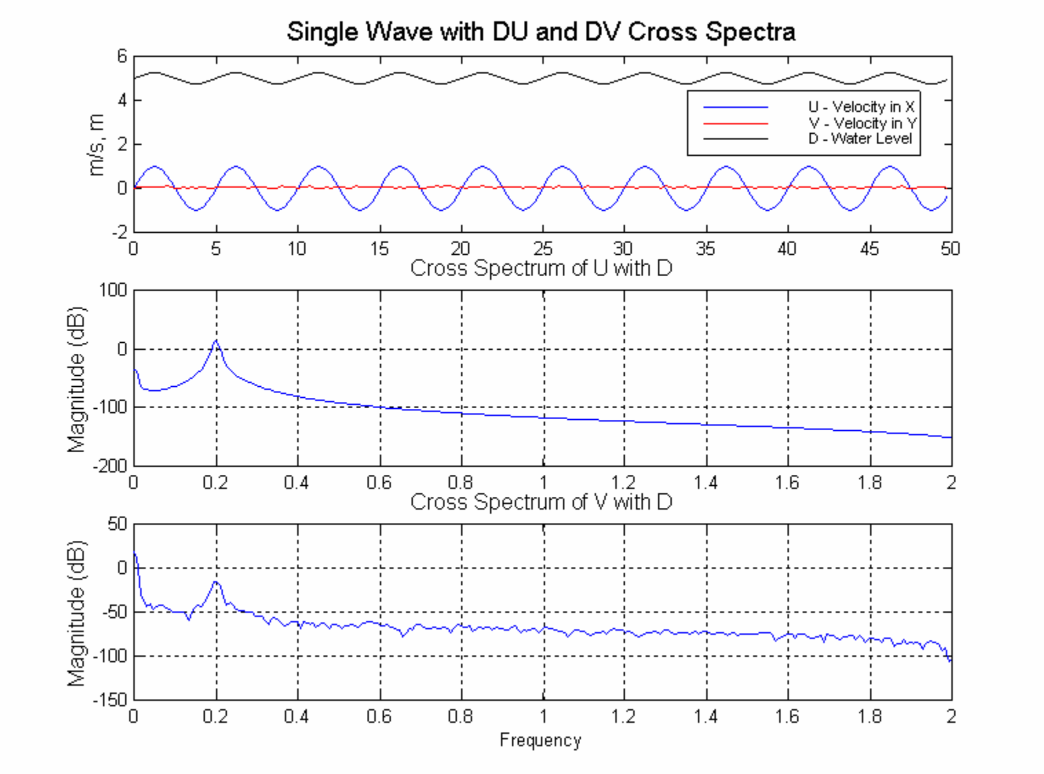

Fig. 11: A simple sine wave propagating in the x direction and the cross spectra between water level and each component of the horizontal velocity.

Both cross spectra in *Fig. 11* show some energy at 0.2 Hz, but that there is far more energy in the *x* direction than in the *y* direction. In the PUV method the direction is obtained by comparing the magnitude of the cross spectra at each frequency (or at each averaged frequency band):

$$D_P(f) = \tan^{-1}(C_{DU}(f)/C_{DV}(f)) \tag{6}$$

Where C_{DV} is the cross spectrum of *D* with *V*
C_{DU} is the cross spectrum of *D* with *U*
f is the frequency (or frequency band) of interest

This technique will also work if waves of different frequency are propagating in different directions. Consider the case where one wave is propagating along the *x* direction and the other is propagating along the *y* direction.

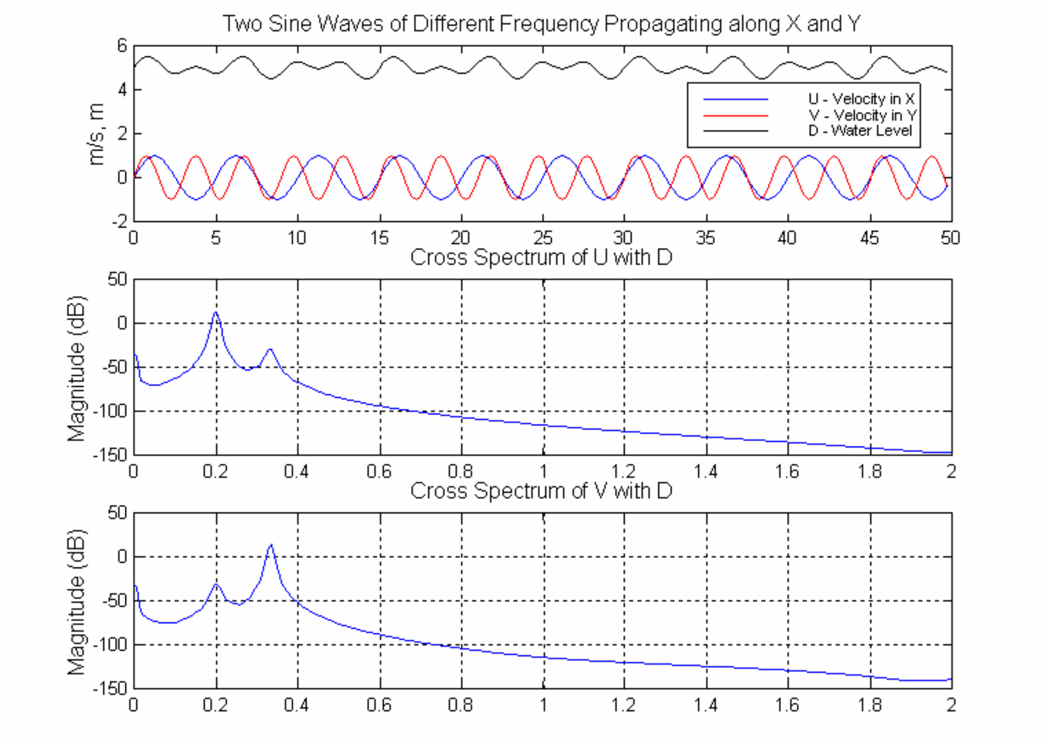

Fig. 12: The top panel shows the time series of horizontal velocity and water level for two waves of different frequency propagating in different directions. The middle panel is the cross spectral density of U with D and the bottom panel shows the cross spectral density of V with D.

The cross spectral density plots of *Fig. 12* clearly recover the fact that our model consists of a five second wave propagating in the *x* direction (velocity *U*) and a three second wave propagating in the *y* direction (velocity *V*).

The PUV method has the primary advantage that a single, bottom-mounted instrument can provide all of the measurements necessary to make the calculations – making it very simple to deploy. There are a couple of significant limitations to the PUV technique however:

1) The PUV technique measures velocity at a single level. In the presence of a background current a weighted average of the background current over the full water column is needed to properly transfer the measurements to the surface using equation (4).
2) At any given frequency or frequency band, the PUV technique will report a single direction. This is a fundamental limitation to the technique that is attributable to the limited number of measurements that have been made. If there are two or more wave trains of similar frequency propagating in different directions across a PUV gauge, the PUV gauge will choose some intermediate direction which might not actually have any wave energy at all. This is important because crossing wave trains will result in different forcing, and different coastal processes, than a single wave train from an intermediate direction.

Arrays

In order to resolve a multidirectional wave distribution at a given frequency more measurements must be made, which generally means deploying an array of instruments. Making additional measurements from instruments deployed with specific spatial separation provides the additional information needed to resolve the full multidirectional wave field for each frequency. A typical setup might involve careful installation of several water level sensors (which could be pressure sensors, capacitance wave gauges, inverted echo sounders, laser altimeters, etc.) at precisely surveyed locations. The cross spectra of these measurements are then compared in any number of ways to yield the directional spreading function.

The cross spectrum of the elements in the array can be defined as:

$$C(\omega) = \int H(\omega,\theta) D(\omega,\theta) H^{\dagger}(\omega,\theta) d\theta \qquad (7)$$

Where $C(\omega)$ is the cross spectra of the measurements

$H(\omega,\theta)$ is the array response to a uniform plane wave of unit amplitude propagating in direction θ

$D(\omega,\theta)$ is the directional spreading function to solve for

$H^{\dagger}(\omega,\theta)$ is the Hermitean transpose of H

The fundamental limitation of this equation is that $C(\omega)$ is finite (limited by the number of elements in the array), while there are a multitude of solutions $D(\omega,\theta)$ that will satisfy equation (7).

There are a number of proven approaches to solving this equation for the directional spreading function, only one is presented here. The Maximum Likelihood Method (MLM) assumes a solution at each frequency is of the form:

$$D(\theta) \propto \frac{1}{H(\theta)^{\dagger} C^{-1} H(\theta)} \qquad (8)$$

Where C^{-1} is the inverse of C

In essence, equation (8) is solved for each frequency at each direction to determine if any wave energy exists at that frequency and in that direction. As long as C contains cross spectra from a sufficient number of independent measurements and H is well defined, then it is possible to resolve multiple wave trains of differing direction but similar frequency because the method specifically looks for waves in every direction at each frequency.

The RDI ADCP Waves Array Technique

The basic principle behind wave measurement is that the wave orbital velocities below the surface can be measured by the highly accurate ADCP. The ADCP is bottom mounted, upward facing and has a pressure sensor for measuring tide and mean water depth. Time series of velocities are accumulated, and from these time series, velocity power spectra are calculated. To get a surface height spectrum, the velocity spectrum is translated to surface displacement using linear wave kinematics. The depth of each bin measured and the total water depth are used to calculate this translation.

To calculate directional spectra, phase information must be preserved. Each bin in each beam is considered an independent sensor in an array. The cross-spectrum is then calculated between each sensor and every other sensor in the array. The result is a cross-spectral matrix that contains phase information in the path between each sensor and every other sensor at each frequency band. The cross-spectrum at a particular frequency is linearly related to the directional spectrum at a particular frequency. By inverting this forward relation, we solve for the directional spectrum.

Background:

The use of Doppler sonar to measure ocean currents is by now well established, and is documented in the RDI publication *Acoustic Doppler Current Profilers, Principles of Operation* (RD Instruments, 1989). Conventional acoustic Doppler current profilers (ADCPs) typically use a Janus configuration consisting of four acoustic beams, paired in orthogonal planes, where each beam is inclined at a fixed angle to the vertical (usually 20 or 30 degrees). The sonar measures the component of velocity projected along the beam axis, averaged over a range cell whose along-beam length is roughly half that of the acoustic pulse. Since the mean current is assumed horizontally uniform over the beams, its components can be recovered by subtracting the measured velocity from opposing beams. This procedure is relatively insensitive to contamination by vertical currents and/or unknown instrument tilts (RD Instruments, 1989).

The situation regarding waves is more complicated. At any instant of time, the wave velocity varies across the array. As a result, except for waves that are highly coherent during their passage from one beam to another, it is not possible to separate the measured along-beam velocities into their horizontal and vertical components. However, the wave field is statistically steady in time and homogeneous in space, so that the cross-spectra of velocities measured at various range cells (either between different beams or along each beam) depend on wave direction. This fact allows us to apply array-processing techniques to estimate the frequency-direction spectrum of the waves. In other words, each depth cell of the ADCP can be considered an independent sensor that makes a measurement of one component of the wave field velocity. The ensemble of depth cells along the four beams constitutes an array of sensors from which magnitude and directional information about the wave field can be determined.

Using an ADCP as a Wave Gauge

The ADCP can use its profiling ability (bins and beams) as an array of sensors. Because the ADCP can profile the water volume all the way to the surface, it can be mounted in much deeper water than a traditional pressure (PUV triplet) based device. Higher frequency waves attenuate more quickly with depth below the surface. The ADCP can measure much higher frequency waves than a PUV and do so in deeper water, because it can make measurements higher up in the water column. Additionally, the ADCP has many independent sensors (bins-beams) so even when sampling at a 2Hz sample rate the data is as quiet as if it had been sampled at 200Hz by a single point meter (example uses 25 bins and 4 beams).

To achieve the best possible solution for wave height spectrum, the height spectrum and the noise spectrum are fit to the bin-beam data using a least squares fit. In addition to the orbital velocity technique for measuring wave spectra, the ADCP can measure wave height spectra from its pressure sensor (with frequency/depth limitations) and from echo ranging the surface. Within the frequency range of the pressure sensor the pressure height spectrum is an old reliable reference for data comparison. The surface track measurement of wave height is reliable most but not all of the time. The advantage of the surface track derived height spectrum is that it is a direct measurement of the surface and can measure wave energy at very high frequencies (higher than 0.9 Hz in some installations). Having three completely independent measures of wave height spectrum that all agree very closely is a solid argument for data quality.

The directional spectrum is much truer and of higher quality than any sort of triplet (PUV, UVW, PRH) and is almost as good as large home-built arrays. The Maximum Likelihood Method used for inversion allows one to independently resolve the wave field in each direction. The full circle (360 degrees) is arbitrarily divided into as many slices as one chooses (up to 360 slices of 1 degree width). Because of this, the RDI directional spectra algorithm can resolve two separate swells arriving from different directions at similar frequency. This feat is impossible using traditional triplet algorithms.

The ADCP measures a sparse array and as such, it cannot achieve the aperture of expensive home built arrays, (e.g. Duck Is.). However, the aperture of the beams gives the ADCP a significant improvement in directional accuracy over single point measurements. A traditional triplet algorithm uses only the first three terms in a Fourier series so it can identify a single directional peak particularly at longer wavelengths. However, buoys, PUVs, and other triplets cannot accurately represent the multiple directional peaks or even the true directional distribution.

In the ADCP wave algorithm there are many sensors giving an array with many degrees of freedom and some aperture. The Maximum Likelihood Method used to calculate the directional spectrum has a smearing kernel associated with the inversion. By using the Iterative Maximum Likelihood Method, the spreading of the directional spectrum can be corrected. The process is repeated until the directional spectrum converges to what the data actually supports. The spectrum will get narrower and sweep up directionally spread power into the peak as long as the measured data supports it. The result is a directional spectrum that more accurately represents the true directional distribution.

Data Screening and Processing

Current Profile Averaging

Ensembles are collected and averaged over the averaging interval independent of ping rate. Velocities that are marked bad (-32768) are not included. The profile averaging interval is independent of the burst sampling interval. For example, averaged ensemble data can be displayed every 6 minutes and bursts processed every 17 minutes (2048 ensembles).

A burst is a collection of continuous ensembles. *WavesMon* collects ensemble data until the requested ensembles per burst have been accumulated. This burst file is then passed to the waves processing module. Because sampling at 2Hz can use up memory and batteries very quickly, the ADCP can be configured for burst sampling as well. The ADCP can be setup to ping continuously for 20 minutes at 2 Hz then sleep for 40.

Discontinuities

If for some reason ensemble data is lost, it is important that waves processing does not try to perform an FFT on data that was collected at 2 different times. Because of this constraint, the software requires that a burst be continuous. It cannot straddle two different burst-sampling intervals. If a discontinuity is detected in the ADCP time stamps for a particular burst, the burst accumulation will be reset to the beginning. This attempts to resynchronize with the burst interval of the data with as little data loss as possible. When a discontinuity is detected, the program status pane on the status bar displays the discontinuity message.

Pressure Time Series

The pressure time series is in fact the pressure sensor derived depth in mm. There is little or no screening required for this data.

Velocity Time Series

The velocity time series are edited for wild points. First, marked bad velocity data (-32768) points are thrown out. Where bad values are found they are filled with an interpolation of adjacent data points. Next, the data is screened for values that are 5 standard deviations or more outside of the mean. This process is done 3 times so that the calculation of mean and STD are not corrupted by the bad data being removed. Once again, each bad value is filled by interpolation. If more than 20 percent of the data has to be thrown out in this manner, the entire time series is marked bad.

Surface Track Time Series

The surface track can and will have wild data points in the time series. There are conditions under which it will not work at all. Most of the time there is enough roughness to the sea surface to make surface track viable.

To find the surface the RSSI (echo intensity) for each beam is first corrected for geometric spreading of the beam. This flattens the RSSI response with range. Next, the largest data point in the vicinity of the surface is found. The general location of the surface is calculated with the pressure sensor derived depth. Once the biggest bump is identified, the X and Y coordinates for the biggest data point and 2 adjacent data points on either side are calculated. With this information, the slopes on either side of the peak can be used to interpolate to the actual peak. This process reduces quantization error introduce by the size of the bin. Once the location is determined, it is stored in the time series for that beam.

Surface track time series are edited for wild points in the same way the velocity data is. Three passes are made through the data in which values greater than 5 standard deviations from the mean are thrown out and interpolated.

Waves Processing

The waves processing algorithms are computationally intensive and may require some time to perform. Depending on what data you have selected you can process wave height spectra from the pressure sensor, from the surface track, and from the velocity data. The velocity data is used to calculate directional spectra.

Height Spectrum

The ADCP has three different independent techniques for measuring non-directional wave height spectrum.

1. The orbital velocity can be measured up close to the surface then translated to a surface displacement spectrum. This method provides a much better frequency response because the measurements can be made farther up in the water column where the exponential attenuation of wave energy with depth, has not reduced the signal much (*Fig. 13*).

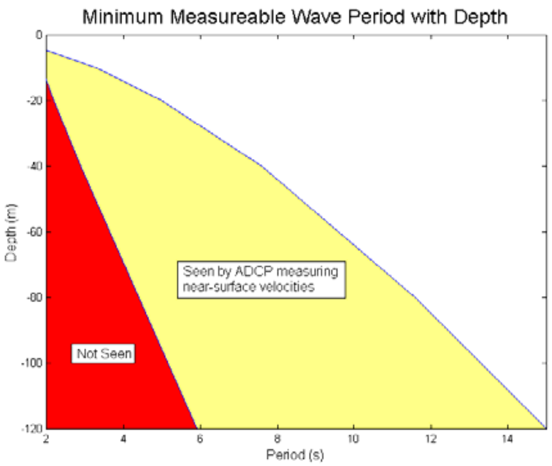

Fig. 13 The approximate decay of wave energy with depth that is resolvable by an ADCP using near-surface velocity measurements. The reported depth is the mounting depth of the ADCP, the yellow area is the limit of measurements for the pressure sensor which is located at the mounting depth (note: the upper bound of this area corresponds to Fig. 3). However, the ADCP measures velocity profiles to near the surface, which allows it to measure much higher frequency waves than can be seen at the mounting depth. Only the red area is excluded from the near-surface velocity measurements.

2. The surface track is direct measurement of the surface and is not frequency dependent except for the resolution of the echolocation of the surface.
3. The pressure sensor derived spectrum is a traditional technique, but is limited because the measurement must be made on the bottom.

Orbital Velocity Height Spectrum

The velocity height spectrum is the most reliable and versatile technique. The velocity height spectrum is calculated from time series of orbital velocities measured for each bin and beam in the profile. The velocity measurement does not readily suffer from the depth/frequency limitations of a pressure sensor, because velocity data is measured right up to the surface. Because of the many independent sensors (bins), a special algorithm is used to capitalize on the abundance of data. By assuming a linear relationship between the measured velocity spectrum and the translation to surface displacement equation for a given frequency, a least squares fit can be applied to the data. The least squares algorithm calculates a surface spectrum and a noise spectrum from spectra for the entire velocity profile. This reduces the measurement noise enormously and allows higher frequencies to be measured. For example, a spectrum calculated from 25 bins, 4 beams, and band averaged by 4, is 20 times quieter than a single time series spectrum. With measurement error so small, one can be sure that any fluctuation in the spectrum is actual environmental statistics.

Surface Track Height Spectrum

The surface track height spectrum is calculated by transforming time series of echo-location ranges for each beam. The time series are screened for wild points then transformed. After power spectra are calculated for each beam, the spectra can be averaged. Bias is removed and band averaging is done to complete the calculation.

Surface track spectrum does not suffer the depth limitations of pressure and velocity measurements so it is ideal for measurement of higher frequency energy. If the surface is within the profiling range of the instrument, it is possible to measure wave energy higher than 0.6 Hz using surface track spectra.

There are two main sources of error for surface track. First, the surface can be difficult to locate under glassy surface conditions. When the ocean surface is very smooth, it acts like a mirror, reflecting the ADCP's slant beams away. While the surface is wind roughened enough to surface track, most of the time you will want to use velocity as a backup for those times when it is not. The second source of error is created by the bin size used for profiling. The size of the depth cell indicates the resolution of the raw surface return. To improve this, a linear interpolation to the peak range is done. This interpolation improves the quantization error by square root 12.

Pressure Height Spectrum

The pressure derived height spectrum is available as a reference for those who would like to see a familiar technique for comparison. The pressure height spectrum is limited in its frequency response because of the frequency dependent attenuation of waves with depth. To calculate pressure height spectrum, time series of pressure sensor depths are transformed into the frequency domain. The noise floor is calculated and removed to avoid biasing the wave height spectrum. To account for the attenuation with depth the spectrum is then translated to surface displacement by cosh(kh).

Standard non-directional spectra plots as output from the RDI ADCP Waves Array Technique software is shown below. Note that all three measurements agree very well, and that the high frequency cutoff decreases with the depth of the measurements, as expected.

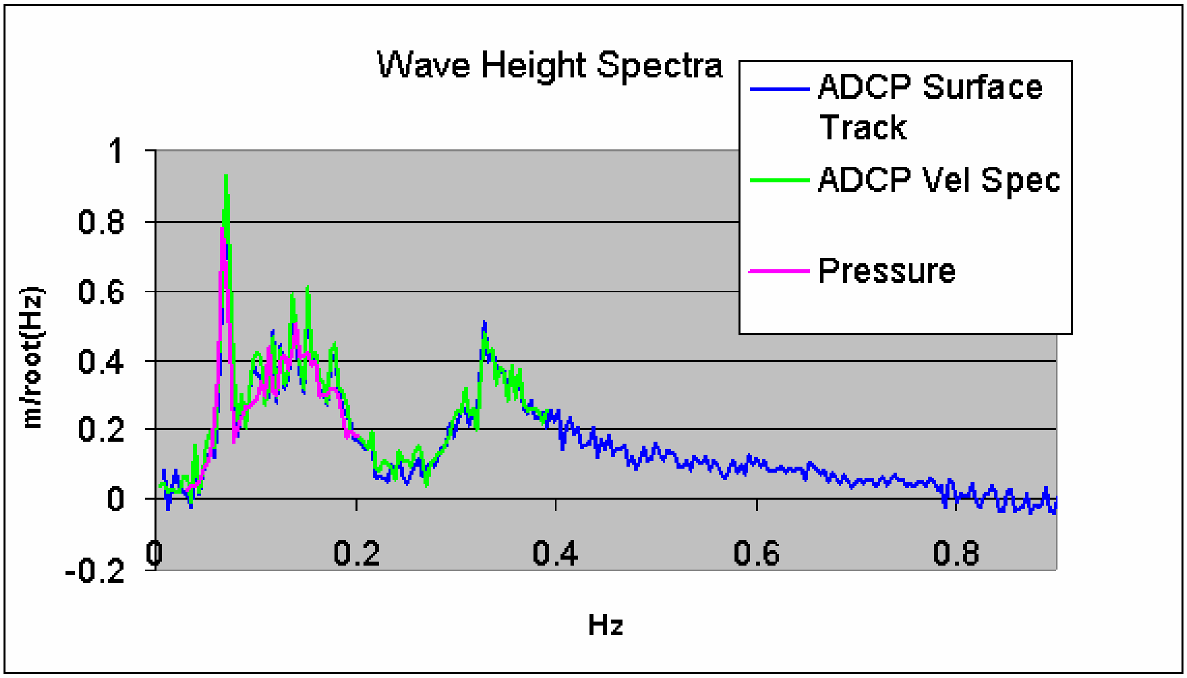

Fig. 14:A non-directional spectrum as output by the RDI ADCP Waves Array Technique. The spectra calculated by each of the three independent techniques agree very well. The high frequency limit decreases with the depth of the measurements, with the pressure sensor at the instrument head seeing the most significant limitations while the surface tracking measurements are limited only by the rate of data measurement.

Bias Removal

Each of the techniques for measuring wave height spectra (pressure, surface-track, and velocity) has its own set of error sources and its own measurement noise. By Parseval's theorem and conservation of energy, the energy in the frequency domain is the same as the energy in the time domain. We assume that the measured time series is a superposition of the signal (waves) and the measurement noise and that the noise of measurement is white. The Fourier transformation of white noise is white noise. This white noise has a zero mean distribution, however when we square the frequency spectrum to get power, the noise becomes always positive. This biases the power spectrum by adding a positive offset. To measure this bias the software calculates the mean of the power spectrum at frequencies too high for environmental wave energy. By finding the noise floor of the power spectrum, it can then be removed. This process is done for each of the independent height spectra calculations and is required in order to get good agreement among the three. Because the bias has been subtracted from the overall spectrum, it is possible to have negative values at frequencies with no energy in them. These negative values are retained because with future averaging they will tend to zero. If the negative values were not retained, future averaging of spectra would bias the measurement.

Cut-Off Frequency Calculation

A highest reasonable frequency or cutoff frequency is calculated for each of the wave height spectra. The cutoff frequency is important for the calculation of significant wave height, determining the frequency range over which to search for a peak and plotting. For the pressure spectrum, and the velocity spectrum, at a particular depth, there is some frequency range at which the measurable wave energy drops below the noise floor. For example, a pressure sensor in 150 meters depth cannot possibly measure 3-second interval waves. To calculate the cutoff frequency *WavesMon* uses linear wave theory to determine a rough starting point based on the depth and frequency. This relation is essentially the inversion of the translation to surface displacement equation, for a specific depth and frequency. The algorithm then leaps forward to a higher frequency, and slides a filter back wards through the spectrum data until the filtered spectrum exceeds the noise floor.

Having a cutoff frequency is also valuable for scaling velocity spectra to surface displacement. The scaling equation has a frequency and depth dependent gain. At high frequencies, this gain is large. Once the wave energy has dropped below the noise floor at high frequencies, the large gain is multiplied by the noise floor, not the vanishing small wave energy. This causes the tail end of the velocity and pressure spectra to grow exponentially. There is an exponentially growing gain and a fixed noise floor. By having a valid cutoff frequency this region can be avoided by automated (computer) analysis.

Correction for Background Currents:

In the presence of strong background currents, the subsurface measurements will result in an overestimate of the wave height. This can be corrected using the profiled velocity data to create the weighted background velocity field to use in equation (4). *Fig. 15* below shows non-directional spectra that were gathered in the presence of a large background current, and demonstrates the differing spectra that are obtained when that background current is neglected or included. Notice that the spectrum calculated from the pressure sensor is the most affected, which is due to the fact that the pressure sensor measures the deepest, and therefore has the most amplification applied to it.

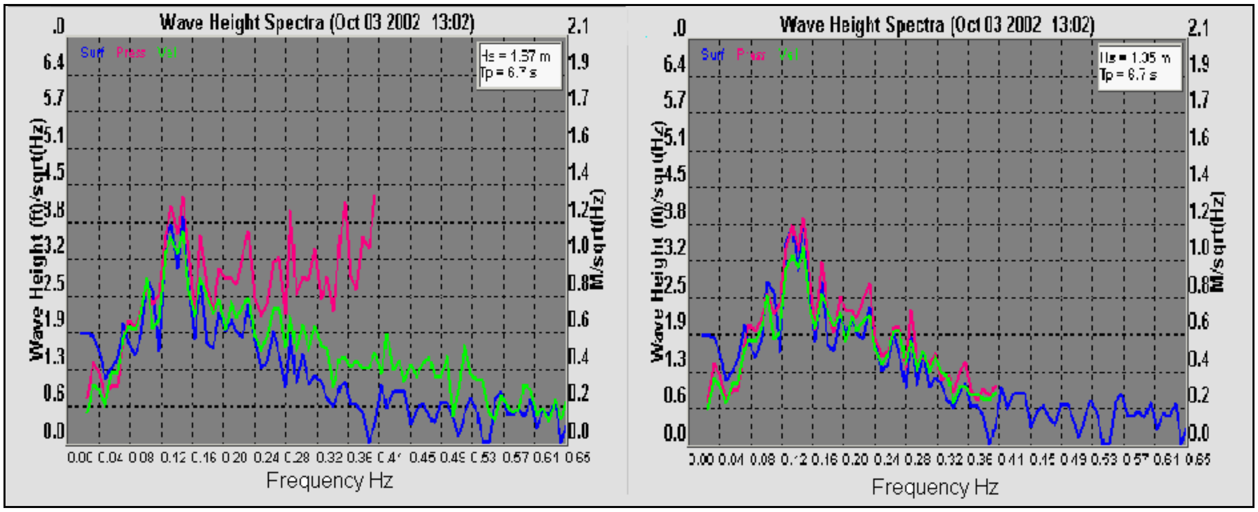

Fig. 15: Non-directional wave spectra gathered during a large background current. The panel on the left transferred the data to the surface values without correcting for the background currents, the panel on the right is the same data, but the transfer properly included the background currents.

Directional Spectrum

The ADCP uses its bins and beams as an array of sensors. By computing spectra from time series at each sensor, wave phase information is obtained. Next the cross spectrum is calculated between every sensor and every other sensor. This cross-spectral matrix has information about the relative phase between every sensor and every other sensor at each frequency. The forward relation assumes a linear relationship between the cross-spectral matrix and the directional spectrum.

There are a number of ways to constrain this problem. The Maximum Likelihood Method (MLM) is used for RD Instruments wave directional spectra because it is a good tradeoff between narrow directional accuracy and false side-lobes. The directional spectrum calculation works best if three depth cells are used. A slight improvement is gained if four or more are used, however the time of processing is increased dramatically. The placement of these bins should be higher in the water column but within the profiling range of the instrument and conservatively below the surface. The further out they are, the more aperture the array has. The directional spectrum also benefits from band averaging. By band averaging adjacent frequency bands, the data is smoothed before inversion and the degrees of freedom of the raw data are increased.

The directional spectrum algorithm has several features that are important:

- The ADCP bins and beams are used as a virtual array of sensors. This provides the many degrees of freedom required describing a potentially complex or sharp directional distribution. In contrast, a PUV or triplet type algorithm can only provide 3 degrees of freedom and only the first three terms in the Fourier series are actually derived from measured data.
- The array of sensors has a substantially greater aperture than a single point measurement (PUV, Buoy). A larger antenna or array aperture is required to measure finer directional resolution and accurately reproduce narrow spectra.
- The Maximum Likelihood Method (MLM) used to construct the directional spectrum allows us to measure multiple directions arriving at the same frequency. The MLM process smears the directional distribution, however by iteratively applying the technique (IMLM), this smearing can be undone. The IMLM technique converges to a directional spectrum that matches the data for a reasonable number of iterations. Experience has shown one, two, or three iterations sweeps up

most of energy back into the peak. The difference between a directional spectrum that was processed with three iterations and 20 iterations is small.

The real power of the RDI ADCP Waves Array Technique is that, as an array, it is capable of measuring waves of similar frequency propagating in different directions. *Fig. 16* shows a Cartesian plot of the energy on direction vs. frequency axes. Note particularly in the center of the plot that the ADCP Waves Array Techniques is resolving two different wave trains with 0.17 Hz frequency (about six second period): one is incident from about 295 degrees and the other is incident from about 225 degrees. A PUV gauge deployed in the same seas reported a single wave direction of 260 degrees – where there was in fact no energy of any significance.

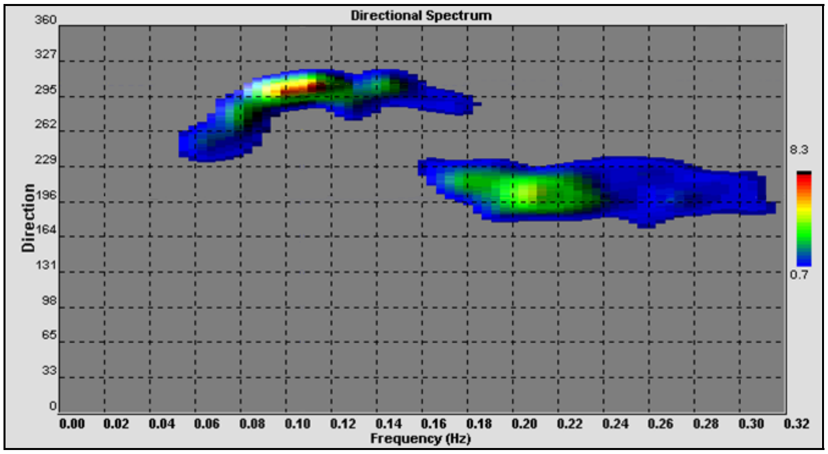

Fig. 16:. Note the presence of two waves at 0.17 Hz propagating in different directions.

Maximum Likelihood Method (MLM)

The Maximum Likelihood Method was originally used as a directional estimator with seismic data to verify the nuclear test ban treaty. MLM does not provide the narrowest possible directional estimator, however it is a good tradeoff. Extremely narrow directional estimators (Eigen Vector Method) have larger side-lobes in the antenna function. Larger side-lobes increase the odds false detection in a particular direction. The MLM is not prone to false indications and can be made as narrow as the data will support with the Iterative MLM process. In the process of inversion, the MLM spreads or smears the directional distribution at each frequency, much like a convolution filter. The IMLM process corrects for this spreading.

Iterative Maximum Likelihood Method (MLM)

The MLM solution for the wave directional spectrum is a good reliable first estimate of direction. However, the inversion process spreads or smears the directional distribution. To correct for the spreading a new cross-spectral matrix C is calculated using the forward relation on the MLM directional spectrum. Then a twice-smeared directional spectrum is calculated from this new C.

By taking the difference between the once smeared and the twice-smeared estimates and applying some small gain (over relaxation) a correction is calculated. This correction is then applied to the once smeared or original MLM estimate to try to get back to the real environmental directional distribution. This process can be applied iteratively. Each successive iteration has to correct less, and the directional spectrum eventually converges to the best estimate that agrees with the data.

In practice, only a few iterations are required to achieve a beautiful spectrum and the point of diminishing returns. Compare spectra with and without iteration. Observe how the power is swept up into a narrower, higher peak with iteration. The directional distribution will not get narrower than the true environmental distribution represented by the data. Because of this and the fact that the algorithm is designed to measure multiple directional peaks, it will provide more true to life directional distributions than many other techniques.

Flow Diagram for Wavesmon Processing

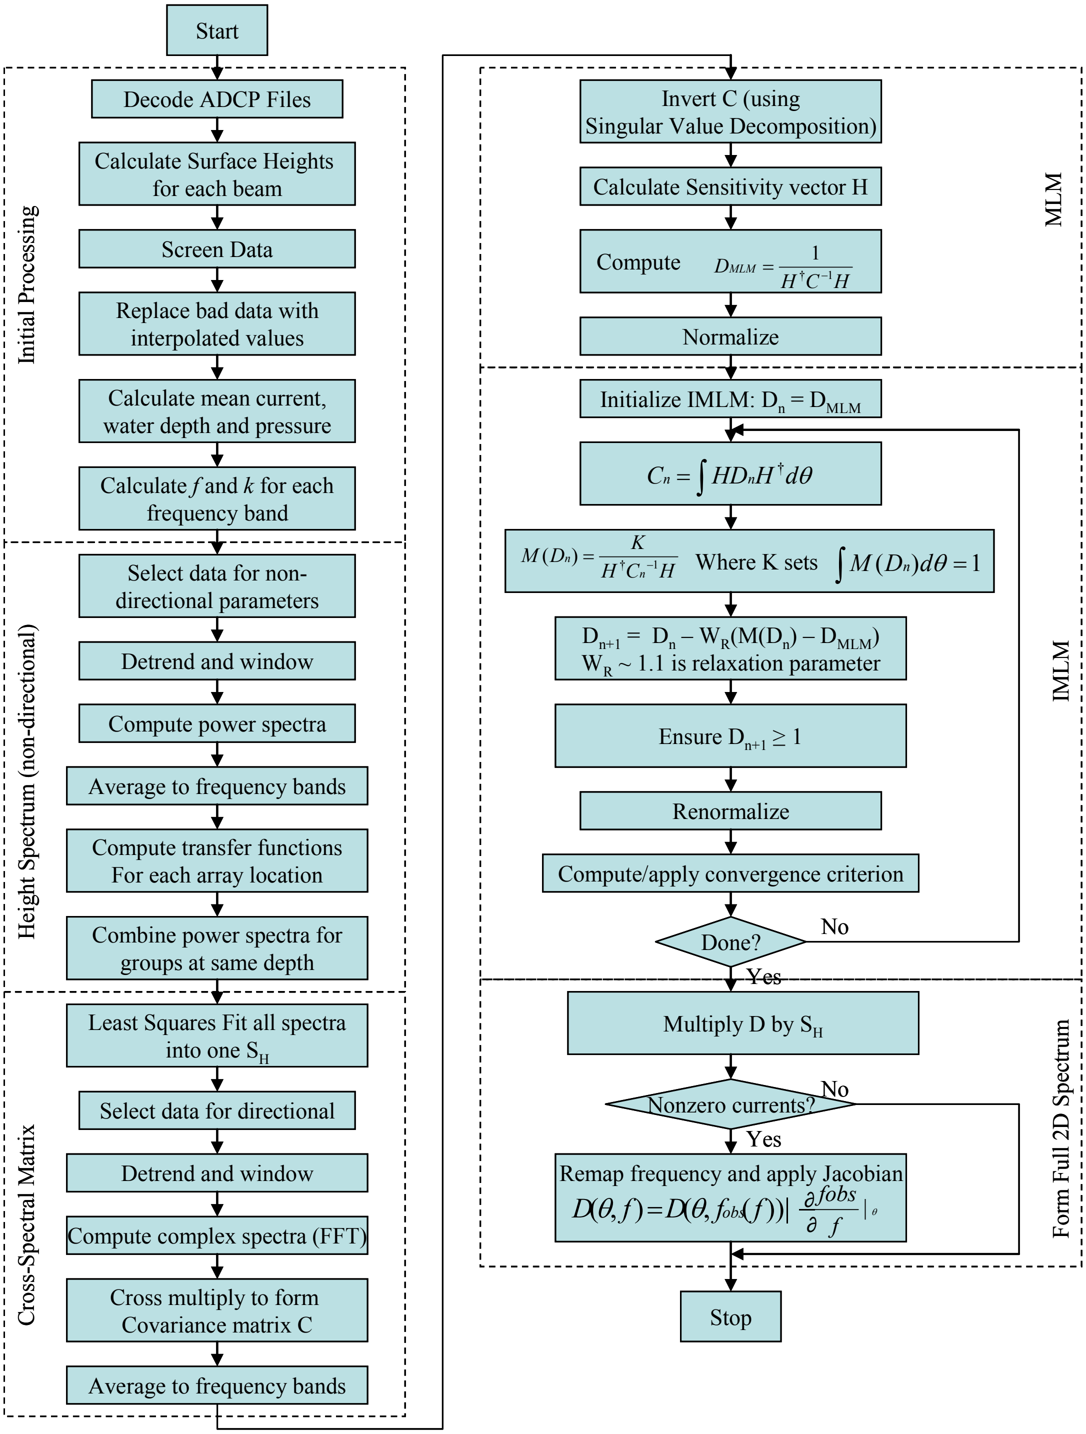

Fig. 17: Complete flow diagram describing Wavesmon processing

ADCP Waves Performance Specifications

System Frequency	**1200 kHz**	**1200 kHz**	**600 kHz**	**600 kHz**	**300 kHz**	**150 kHz**	**75 kHz**
Deployment Depth	5 m	10 m	20 m	40 m	80 m	120 m	200 m
Bin Size	0.35 m	0.35 m	0.5 m	1.0 m	2.0 m	2.0 m	4.0 m
Non-Directional Spectrum							
Surface Track Cut-Off Freq (Hz)	0.99	0.99	0.99	0.74	0.54	0.43	0.34
Surface Track Cut-Off Period (s)	1.01	1.01	1.01	1.35	1.85	2.33	2.94
Surface Track Min Wave Ht (m)	0.1	0.1	0.14	0.29	0.57	0.57	1.10
Velocity Cut-Off Freq (Hz)	0.60	0.56	0.50	0.38	0.36	0.17	0.18
Velocity Cut-Off Period (s)	1.67	1.79	2.00	2.63	2.78	5.88	5.56
Velocity Min Wave Ht. (m)	0.05	0.05	0.05	0.15	0.10	0.20	0.50
Pressure Cut-Off Freq (Hz)	0.39	0.30	0.23	0.16	0.10	0.07	0.06
Pressure Cut-Off Period (s)	2.56	3.33	4.35	6.25	10.00	14.29	16.67
Pressure Min Wave Ht (m)	0.03	0.03	0.03	0.05	0.1	0.1	0.10
Directional Spectrum							
Directional Cut-Off Freq (Hz)	0.58	0.42	0.30	0.21	0.14	0.12	0.10
Directional Cut-Off Period (s)	1.72	2.38	3.33	4.76	7.14	8.33	10.00
Directional Min Wave Ht. (m)	0.05	0.05	0.05	0.15	0.10	0.20	0.50

Frequently Asked Questions

How do I Measure Very Small Waves?

Very small waves have very small orbital velocities and require a very quiet measurement. The best instrument for small waves is the 1200 kHz in shallower water. Choose a slightly larger bin size even as much as a meter. If the resulting Velocity Spectra do not match the Pressure-Derived Spectra very well at long wavelengths, it may be beneficial to use the **Small Wave Screening Frequency**. This allows the pressure sensor to aid the velocity spectrum calculation for very long, very small waves.

The Small Wave Screening frequency can be set on the Data Screening page of Wavesmon.

How do I Measure High Frequency Waves in Deeper Water?

In deeper water, the surface track is more reliable because the beam footprint on the surface is larger. Choosing a medium bin size will optimize the setup for the surface track. For example, a 600 kHz system deployed in 40 meters of water depth using 0.75-meter bins will do an exceptional job of surface tracking out to 0.8Hz.

How do I Measure Very Large Waves?

Use the default setup, then make sure that enough bins are being collected (WN command) to profile three meters past the surface even when a large wave passes. High tides should be considered as well.

How do I Export Screen Images in Real-Time?

The Graphical Outputs page of the setup wizard allows selection of the current profile or direction view (left side of the screen) and the waves view (right side of the screen) to PNG (Portable Network Graphic) format image files. These files have a filename that includes the date-time stamp for the images. The *.png files are compatible with browsers, and Windows software (Word, PowerPoint etc). They have good resolution and are relatively small (memory size).

How do I Save Text Files?

The Advanced File Outputs page of the setup wizard allows the selection of text outputs in various supported formats (described by selecting the format and clicking the “Show” button).

How do I Correct Waves for the Effects of Currents?

The ADCP Environment page of the setup wizard allows the selection of Correct for Currents under Other Options.

How do I Check for Data Quality?

Look at the processed data (*.wvs file) in WavesView. Observe how well the three different wave height spectra match. If the Surface Track Spectrum seems quiet and reasonable, it can be used as a reference. If the Velocity and the Pressure Spectra are offset, even at lower frequencies, then the mean water depth may be in error or the altitude (instrument height above the bottom) may be in error. See the **Processing** tab for corrections of the depth measurement.

How does the Magnetic Variation Setting Affect the Data Collected?

This depends on whether the magnetic variation has been applied when the data was initially gathered or if it is being added in post-processing. If the magnetic variation was entered at the beginning of data collection in real time (which is equivalent to sending the EB command for experienced ADCP users) then both the array processing and UVW processing (accessed by selecting “Moored (Dynamic) Mounting”) will be correct. All displays and all log files will also be correct. However, if packets data is collected and stored, the individual packet headings are not corrected.

If the magnetic variation is to be added or corrected in post-processing then the array processing and the UVW processing will be corrected. However, the current ensembles, packets (if collected) and log files will not be corrected in the reprocessing. It is worth noting that the magnetic variation applied from the advanced menu in Wavesmon during reprocessing is added to the magnetic variation read from the header in the stored data. If, for example, a magnetic variation of 13° is desired but the data was collected with 11° applied, then 2° should be entered as the magnetic variation correction to be applied in reprocessing.

Summary

This primer was written with two goals in mind: the first was to simply describe the field of directional wave measurements. To recap: directional wave measurements are an attempt to statistically characterize wave energy in terms of its magnitude, frequency and direction. The magnitude and frequency are relatively easy to characterize from a single instrument, but care must be taken if subsurface measurements are used. Determining direction requires more than one measurement, and the greater the number of independent measurements the greater the resolution of the technique.

The second goal was to present the RDI ADCP Waves Array Technique. To summarize this technique it:

- Uses the beam geometry of a single, bottom-mounted ADCP to create a multi-element array of surface measurements.
- Uses three independent techniques to determine the non-directional wave parameters. The spectra obtained from the orbital velocities are considered to be the primary measurement. The spectra obtained from the pressure sensor and from the range to surface measurements are also provided, but for redundancy and quality assurance.
- Applies a Maximum Likelihood Method to a twelve element array of orbital velocities (after correcting to the surface) to determine wave direction. This allows the resolution of waves of similar frequency propagating in different directions, which ordinarily would require deploying several instruments.
- Uses the vertical profile of velocities to determine the properly weighted background current and exclude any such current from the transfer of the measured subsurface parameters to the surface.

The RDI Waves Array Technique is as easy to deploy as any single point instrument deployed to measure PUV. But through the use of BroadBand processing and the creation of an array, it is a far more capable instrument.

Further Reading

An excellent qualitative discussion of ocean waves, along with a number of entertaining sea stories, is provided by:

Bascom, Willard, 1980: *Waves and Beaches,* Anchor Books, Doubleday, New York, USA.

For the more mathematically inclined there are a number of excellent texts, including:

Kundu, Pijush K., 1990: *Fluid Mechanics*, Academic Press, San Diego, USA.
Lighthill, James, 1978: *Waves in Fluids*, Cambridge University Press, Cambridge, UK

For the mathematical details of the statistics presented:

Priestly, M.B., 1981: *Spectral Analysis and Time Series*, Academic Press Ltd., Great Yarmouth, UK

For more mathematical treatments of directional wave measurements there are number of resources including, but by no means limited to:

Dean, Robert G. and Robert A. Dalrymple, 1991: *Water Wave Mechanics for Engineers and Scientists*, World Scientific, Singapore
Longuet-Higgins, M.S., D.E. Cartwright and N.D. Smith, 1963: Observations of the Directional Spectrum of Sea Waves using the Motion of a Floating Buoy. *Ocean Wave Spectra*, Prentice-Hall, 111-136
Capon, J., R.J. Greenfield, R.J. Kolker, 1967: Multidimensional Maximum-Likelihood Processing of Large Aperture Seismic Arrays, *Proc. IEEE*, **55**, 192-211
Pawka, S.S., 1983: Island Shadows in Wave Directional Spectra. *J. Geophys. Res.,* **88**, 2579-2591
Oltman-Shay, J. and R. Guza, 1984: A Data Adaptive Ocean Wave Directional Spectrum Estimator for Pitch/Roll Type Measurements. *J. Phys. Oceanogr.*, **14**, 1800-1810

For more detailed exposition of the RDI Waves Array Technique:

Terray, EA, BH Brumley and B Strong, 1999: Measuring Waves and Currents with an Upward-Looking ADCP, *Proc. IEEE 6th Working Conference on Current Measurement,* IEEE, New York, 66-71
Strong, B, BH Brumley, EA Terray and GW Stone, 2000: Performance of ADCP-Derived Directional Wave Spectra and Comparison with Other Independent Measurements, *Proc. Oceans 00,* IEEE, New York.

Some comparative studies carried specifically on the RDI Waves Array Technique include:

Terray, EA, RL Gordon and BH Brumley, 1997: Measuring Wave Height and Direction Using Upward-Looking ADCPs, *Proc. Oceans 97,* IEEE, New York, 287-290.
Rørbæk, K and H Andersen, 2000: Evaluation of Wave Measurements with an Acoustic Doppler Current Profiler, *Proc. Oceans 00,* IEEE, New York
Shih, HH and B Strong, 2002: Laboratory Study of ADCP Wave Measurements, *Proc. OMAE 02,* ASME, New York