What Does a Wave Radar Actually Measure?

Kevin Ewans[1], Philip Jonathan[2], and Graham Feld[3]

[1]Sarawak Shell Bhd, 50450 Kuala Lumpur, Malaysia.
[2]Shell Research Ltd, Chester, CH1 3SH, UK.
[3]Shell Global Solutions, Aberdeen, UK.

Abstract

The SAAB REX WaveRadar sensor has arguably become the gold standard in platform-based wave measurement systems in the offshore oil and gas industry. It offers in-situ surface elevation wave measurements at relatively low operational costs. Furthermore, there is adequate flexibility in sampling rates, allowing in principle sampling frequencies from 1Hz to 10Hz, but with an angular microwave beam width of 10 degrees and an implied ocean surface footprint in the order of metres, significant limitations on the spatial and temporal resolution might be expected. Indeed there are reports that the accuracy of the measurements from wave radars may not be as good as expected. We review the functionality of a wave radar using simulations to better understand how wave radar estimates compare with known surface elevations. In addition, we review recent field measurements made with a WaveRadar set at the maximum sampling frequency, in the light of the expected functionality and the numerical simulations, and we include inter-comparisons between SAAB radars and buoy measurements for locations in the North Sea.

1 Introduction

Wave measurements are required for the operation of offshore oil and gas facilities, and wave measurement devices that can be mounted on the platform itself are attractive to operators, because of their relatively low operational costs – they are easy to maintain and service and do not require expensive ship time needed for deployment and recovery of wave buoys. In addition, they offer the potential to measure absolute surface elevation, when mounted on a fixed platform, providing the opportunity to investigate more fundamental aspects of the wave field.

Among the instrument types that have been used are wave staffs, lasers, and radars; but by far the most widely used by the offshore oil and gas industry is the SAAB WaveRadar. Shell alone has 12 facilities in the North Sea instrumented with REX WaveRadars and 10 in the South China Sea. These data provided the vast majority of the platform-based wave data analysed as a part of the CresT JIP project (Christou and Ewans, 2011).

The WaveRadar is a downward-looking unit and can be set to sample continuously at frequencies up to 10Hz, apparently providing excellent temporal resolution for the study of wave processes; but with a specified angular microwave beam width of 10 degrees and an implied ocean surface footprint in the order of metres, significant limitations in the spatial resolution and consequently the effective temporal resolution might be expected. Noreika *et al*. (2011) report an inter-comparison between measurements made with a Datawell Directional Waverider buoy (DWR) and a SAAB Rex

WaveRadar off the Australian Northwest Shelf. The WaveRadar was mounted at an elevation of 26 metres above the water surface on Woodside Energy Limited's North Rankin A (NRA) platform, 135 km north-northwest of Dampier, and the DWR buoy was deployed 3km away in 125 metres of water.

Data were compared from parallel measurements made over a two year period, covering a wide range of conditions. The WaveRadar sea states were found to produce significant wave heights 4% to 10% lower than the DWR, with the biggest departure occurring during tropical cyclone conditions. At storm peak sea states, the WaveRadar significant wave heights were 16% less than those from the DWR. However, the significant wave height calculated from the spectrum corresponding to wave periods longer than nine seconds and wave periods were in good agreement. We present results in this paper that do not show such a bias between wave buoy measurements and WaveRadar measurements in the North Sea.

Nevertheless, the uncertainty of the effect of the resolution of the wave radar on the measurements remains, and it was the objective of the work reported in this paper, to gain an understanding of this. Accordingly, we review the functionality of a WaveRadar using simulations to better understand how the radar estimates compare with known surface elevation. In addition, we review recent field measurements made with a Rex WaveRadar set at the maximum sampling frequency in the light of the expected functionality and the numerical simulations, and we present comparisons between WaveRadars and Datawell wave buoys at several locations in the North Sea.

A description of the WaveRadar is given in Section 2; the simulations undertaken to assess the functionality of the radar are presented in Sections 3 and 4; the field measurements are given in Section 5; and conclusions are given in Section 6.

2 The WaveRadar

The WaveRadar REX is now manufactured by Rosemount Tank Radar AB, Gothenburg, Sweden, and is a derivative of their TankRadar system. (The company was SAAB until 2001 when it was bought by Emerson, and the radar unit is now part of the Rosemount group within Emerson). The first WaveRadar was a land-based TRL unit installed in the early 90s by Rijkwaterstaat in The Netherlands. The first TRL/2 unit was installed in May 1993, and late in 1999 the first WaveRadar Rex was tested at the Shell Cormorant Alpha platform. There are more than 500 WaveRadars installed worldwide, most of which are WaveRadar Rex gauges.

The WaveRadar consists of a transmitter head housing the electronics and a 0.44 m diameter parabolic dish antenna (Plate 1). It employs the Frequency Modulated Continuous Wave (FMCW) method for determining the range to the target. The signal is a low power (< 0.5 mW) frequency modulated 9.7 to 10.3 GHz (linear sweep) continuous microwave signal. Each measurement cycle consists of a linearly-increasing frequency sweep, followed by a linearly-decreasing frequency sweep – producing a triangular frequency modulation signal. The transmitted signal changes frequency during the time taken for the reflected signal to return. The received signal reflected from the water surface is mixed with the signal that is being transmitted at that moment, and the result is a signal with a beat frequency proportional to the distance to the surface. Various processing of the mixed

signal is performed, including digital filtering, spectral analysis, and “peak picking”, to convert the beat frequency to a distance. During the measurement cycle of 10.3 Hz, a number of measurements are produced and an average distance recorded. The measurement cycle includes both an increasing frequency sweep, followed by a decreasing frequency sweep, allowing Doppler shifts to be averaged out.

The manufacturer specifications (RS Aqua, 2009) quote an accuracy: Range < 50 m = ±6 mm, Range > 50 m = ±12 mm; Maximum deviation: Range < 50 m = ±1 mm; Clear downward view of the sea surface within at least a 10 angle conical beam; Range 3 m to 65 m from the adaptor (the measurement datum or origin). The manufacturer calibrates every WaveRadar REX in a special calibration facility; and due to the design and construction of the electronic and microwave unit, the WaveRadar REX calibration is extremely stable and periodic re-calibration is not required.

Plate 1: WaveRadar REX mounted on a South China Sea platform.

3 Simulations

3.1 Method

3.1.1 Surface Wave Simulations

A simple plane sinusoidal wave of frequency 0.2 Hz, long crested in one horizontal direction, was simulated first, to build some understanding of the radar simulations with this simple case before examining the irregular wave case. The amplitude of this sinusoid was set to 1 metre.

Then, a linear simulation of the ocean surface was employed, to provide a more realistic surface for the radar simulations. The simulations were produced with a JONSWAP frequency spectrum, with a peak frequency of 0.10 Hz, and a bimodal directional distribution (Ewans, 1998). The frequency

spectrum was defined with a resolution of 6.3×10^{-4} Hz, over the range 0 to 10 Hz. The resolution and range was chosen to provide a time series of 32,768 points with a time step of 0.0485 s. The directional resolution was 1 degree. A contour plot of the frequency-direction spectrum is given in Figure 1. The units of the colour contours are dB re. 1 $m^2Hz^{-1}deg^{-1}$.

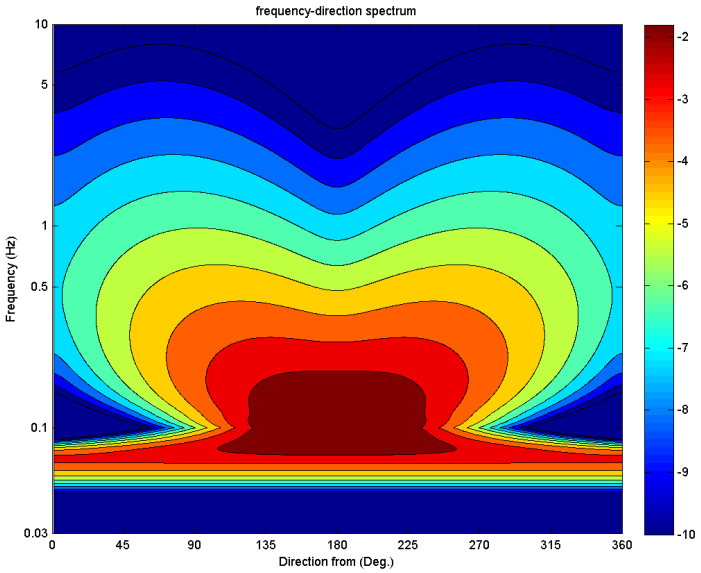

Figure 1: JONSWAP-bimodal frequency direction spectrum. Colour scale in dB re. 1 $m^2Hz^{-1}deg^{-1}$.

At each time step, the surface was generated over a 5 m square at a resolution of 0.01 m, to provide spatial resolution for surface waves with frequencies up to 10 Hz.

3.1.2 WaveRadar simulations

The simulation of the WaveRadar involves a number of modelling steps – specification of the radar antenna pattern, determining the range to many points on the sea surface at an instant of time, modelling the attenuation as a function of range, modelling the reflection from the sea surface at each point, accumulation of the echoes from all of the points, and performing analysis of the accumulated signal strength from all of the echoes to determine the range corresponding to the measured surface elevation.

A complete simulation of the WaveRadar functionality would involve duplicating the FMCW method, entailing modelling the FMCW signal over its triangular sweep through one cycle, and gathering the returned signal at a specific time, associated with the echoes from all the ocean surface grid points. The frequency sweep process during the increasing frequency stage is depicted in Figure 2. The figure shows the radar signal sweep beginning at frequency f_0, at time t_0, and linearly increasing to frequency f_e, at time t_e, over an elapsed time of around 50 ms. At time t_e, the signal frequency starts to decrease returning to f_0, over the ensuing 50ms; this is not shown in the figure. The parallelogram with the thick borders depicts the frequencies of the received signal from the ocean surface echoes. The time delay between the radar transmission and the first return at each frequency is very small, corresponding to the point on the ocean surface closest to radar. The last

return at the given frequency occurs a short time later, corresponding to the most distant point on the ocean surface. The diagram is relevant for a surface wave field that is stationary in time.

At an instant of time, the received signal from all the echoes corresponds to frequencies defined by a vertical slice through the parallelogram, as for example that given at t_e. For a random surface, we make the assumption that each frequency (corresponding to a given range to a surface point) is unique and the phase of the echo can be ignored. Taking the example of Figure 2, at t_e, the ranges are determined from the difference in the frequencies between f_1 and f_2 and f_e. At any given time, the Fourier transform of the received signal mixed with the emitted signal is calculated to determine signal power as a function of delta f. In practice, this is done over the 50 ms sweep. The Fourier transform provides estimates of the amplitudes of each frequency difference component, and it has a peak corresponding to a particular frequency difference or range of interest.

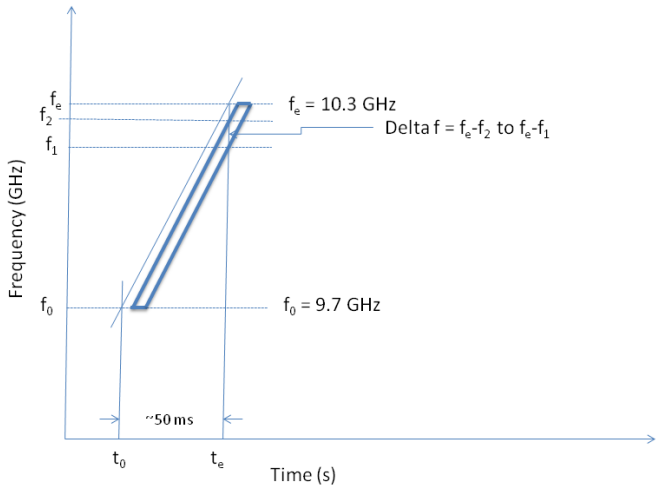

Figure 2: Schematic of the FMCW transmission and reflected signal frequencies over the increasing frequency stage.

Our approach to the simulations is based on the assumption that the signal processing to convert the frequency differences to ranges is done perfectly by the WaveRadar, to keep simulation times to a manageable level, and also because the actual processing performed is proprietary and unknown. Accordingly, we jump directly to what would be the output of the FMCW frequency analysis – the reflected signal intensity (or gain) as a function of range, which is available directly from the simulations.

Similarly, we minimise simulation times by treating the radar as a point source (located at the origin of a rectangular Cartesian coordinate system with z positive upwards). This was considered justified, due to the relatively small size of the antenna, and comparison of results from randomly choosing a point on the surface of the radar as the source confirmed this.

Accordingly, it has been assumed that our signal is the summation of all the reflected signals from the surface of the water at an instant of time. The ray diagram given in Figure 3 outlines the scenario for a specific point, j, on the surface.

The signal strength, $E_j'(x_j, y_j, z_j)$, of the received echo from point, j with vector location $\boldsymbol{x_j}$ and coordinates x_j, y_j, z_j, on the ocean surface can then be expressed as:

$$E_j'(x_j, y_j, z_j) = E(\theta_{j0})A(2r_j)R(\theta_{jr})$$

where $E(\theta_{j0})$ is the antenna signal strength at the angle θ_{j0} subtended by the point x_j, y_j, z_j to the downward vertical at the antenna location of $(0,0,0)$, where $z_j = \eta(x_j, y_j)$ is the surface elevation, and θ_{j0} is also a function of x_j, y_j, z_j.

$A(2r_j)$ is the attenuation associated with path-loss over the range, $r_j = \sqrt{x_j^2 + y_j^2 + z_j^2}$, and

$R(\theta_{jr})$ is the reflection coefficient corresponding to the radar signal reflection angle θ_{jr} at location j, where θ_{jr} is also a function of x_j, y_j, z_j and equal to the angle between the incoming ray from the radar and the local surface normal $\boldsymbol{n}_j$ at x_j, y_j, z_j.

The total received signal at our instant of time is then simply the sum, $\sum_j E_j'(x_j, y_j, z_j)$.

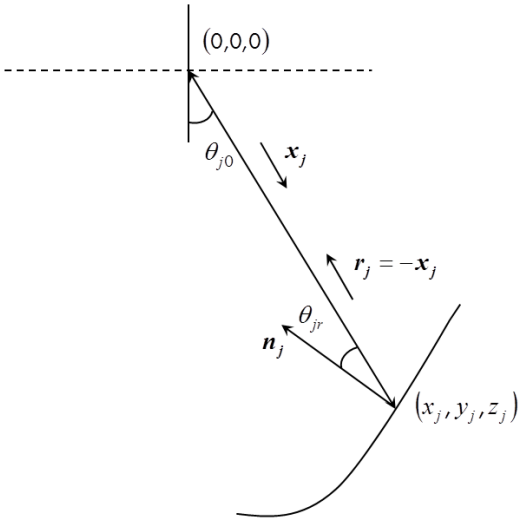

Figure 3: Ray diagram for the point j on the surface

The signal strength, $E(\theta)$, was obtained from the antenna beam pattern provided by RS-Aqua, the suppliers of the Rosemount WaveRadar systems, and is given in Figure 4, with Swedish labelling.

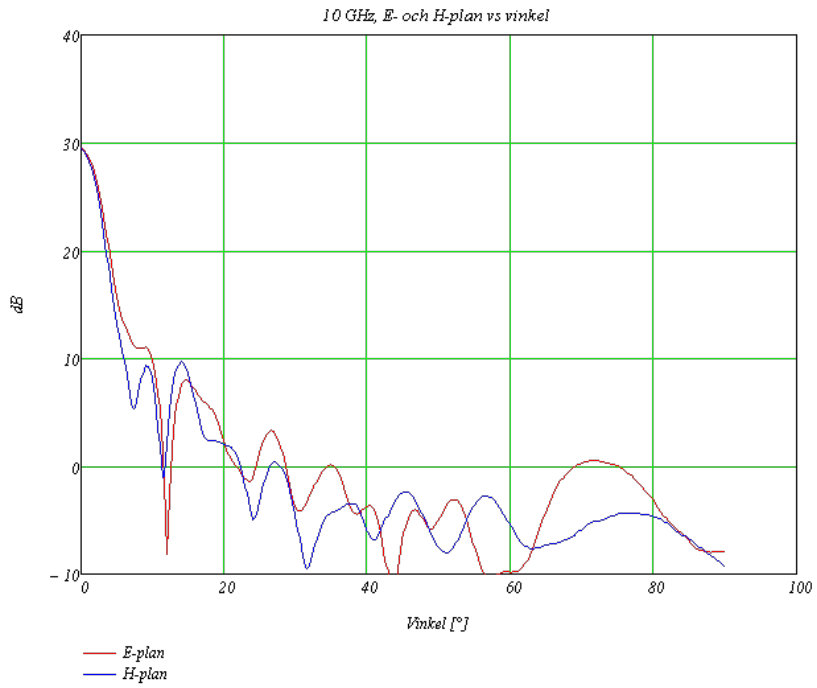

Figure 4: WaveRadar antenna beam pattern.

We have used the Friis expression for the attenuation associated with propagation of the radar signal, $A(r)$, given (in dB) by

$$A_{dB}(r) = -20\log\left(\frac{4\pi 2r}{\lambda}\right)$$

where λ is the wavelength of the radar signal.

The backscatter from the ocean surface is known to depend on the mean-square slope of the surface (Liu *et al*, 2000), which can take on a wide range of values, depending on the sea state and wind forcing. We investigate the backscatter effect on the simulated radar surface elevation estimates for three diffuse reflector forms:

1. An ideal diffuse reflector, in which the intensity of reflection is the same for all reflection angles. This type of reflectance is known as Lambertian, and we refer to is as “Lambertian” in this paper.
2. A combined diffuse-specular reflector, defined by

   $$R(\theta_r) = \cos^{50}\theta_r$$

   The reflection in this case is symmetric about the reflected angle, but reduces (relatively slowly) with increasing angle. We refer to this as “Diffuse” in this paper.

3. A combined diffuse-specular reflector, which is more specular than diffuse, defined by

   $$R(\theta_r) = \cos^{1000}\theta_r$$

   We refer to this as “Specular” in this paper.

The three reflection models are plotted in Figure 5. A gain of one signifies 100% reflection or a perfect reflector. In our modelling, the Lambertian reflector is a perfect reflector for all angles, and the Diffuse and Specular reflectors are perfect reflectors at the specular reflection angle. A more physically realistic approach would be to turn these reflection forms into densities to better reflect the scattering, but this would involve multiplying each by a constant, and the result would simply be a change in the received signal gain, with the reflection gain from each point on the surface changed by the same amount. As it is only the relative gain that is important, we chose to keep the reflection forms simple.

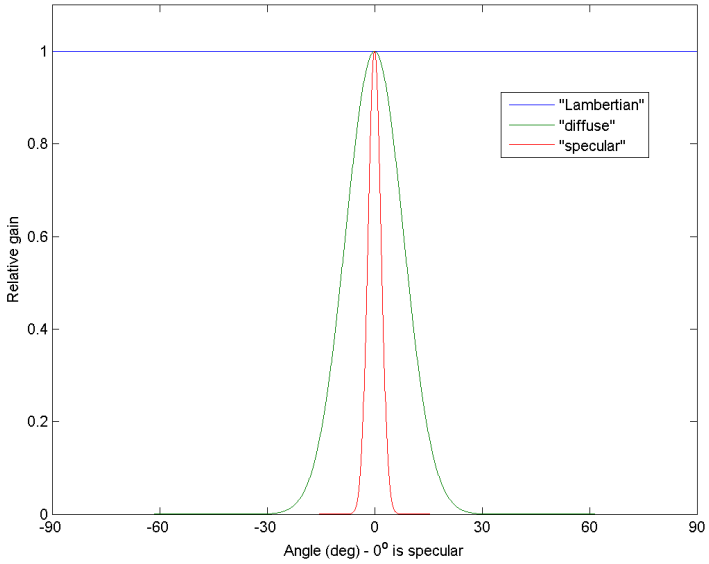

Figure 5: Diffuse reflection models

With reference to Figure 3, the ray from the radar source, at point (0,0,0), to each surface point must be determined. This ray then defines the range to the surface point, allowing the attenuation associated with path loss to be determined, and the angle that the surface point subtends to the vertical at $(0,0,0)$, allowing the appropriate attenuation antenna transmitter gain to be selected.

Accordingly, the position vector, $\boldsymbol{x_j}$, of the surface point, j, relative to $(0,0,0)$, is simply given by $\boldsymbol{x_j} = (x_j, y_j, z_j)$, noting that $z_j = \eta(x_j, y_j) < 0$. The set of all position vectors defines the surface elevation $\eta(x_j, y_j)$.

The magnitude of $\boldsymbol{x_j}$ is the range to the surface point, and the angle that the surface point subtends the vertical, θ_{j0} at the source can be derived from the z-component, z_j, of the direction cosines of $\boldsymbol{x_j}$, as

$$\theta_{j0} = \cos^{-1} \frac{\boldsymbol{x_j} \cdot (0,0,z_j)}{x_j |z_j|}$$

The diffuse-specular reflection component back toward the radar at point j, is the reflection vector, $\boldsymbol{r_j} = -\boldsymbol{x_j}$, from which the reflection angle, θ_{jr}, is determined from

$$\theta_{jr} = \cos^{-1}\frac{\boldsymbol{r_j}\cdot\boldsymbol{n_j}}{r_j n_j}$$

where $\boldsymbol{n_j}$ is the vector normal to the surface at $\boldsymbol{x_j}$, and $\boldsymbol{n_j}$ is calculated from the cross product of the gradient of the surface defined by the position vectors, $\boldsymbol{x}$, of all the surface points, as

$$\boldsymbol{n_j} = -(1,0,\nabla\boldsymbol{x}_x)\times\left(0,1,\nabla\boldsymbol{x}_y\right)$$

where $\nabla\boldsymbol{x}_x = \frac{\partial\eta}{\partial x}$ and $\nabla\boldsymbol{x}_y = \frac{\partial\eta}{\partial y}$

At each time step of 0.0485s, the estimated surface elevation from the simulations, involves accumulating the reflected signals from all of the surface points – some 251,000 points – each signal consisting of a gain and range. For practical purposes we express our signal intensities in dB and refer to them as gain. The reflected signals are ordered in terms of range, and a cumulative sum of the gains calculated. The cumulative sum is smoothed, converted into a distribution and the density function derived. The estimated surface elevation is then taken as the maximum of the density function. Figure 6 is an example. The green curve is the smoothed distribution, and the blue line is the density. The vertical dashed line is the actual range from the radar to the point on the ocean surface directly beneath the radar, with coordinates $\left(0,0,\eta(0,0)\right)$, which we abbreviate to surface location $(0,0)$ henceforth. In this example, the estimated range from the simulation is equal to the actual range.

The WaveRadar takes many samples over each 0.0971 s cycle and averages the result. This cycle includes both the increasing and decreasing sweeps. For practical reasons, we simply take the average of pairs of consecutive 0.0485 s samples.

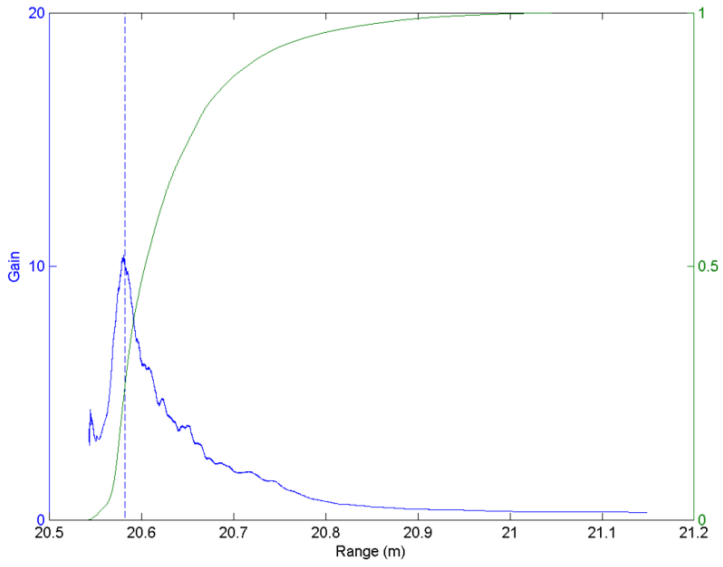

Figure 6: An example of the received signal from the ocean surface at an instant of time. The green line is the distribution of gain, the blue line is the density, and the vertical dashed line is the actual range from the radar to the point $(0,0)$ directly below the radar.

4 Results

4.1 Long-crested Regular Wave

The time series of the simulated WaveRadar surface elevations, together with the actual surface elevations, are plotted in Figure 7. The agreement is good but not perfect – the simulated WaveRadar slightly under-predicts the peaks and over-predicts the troughs. Thus, for this simple plane wave, the simulations produce a small positive bias in the range, which would affect estimates of the still water level but not the wave height (or crest height relative to the mean).

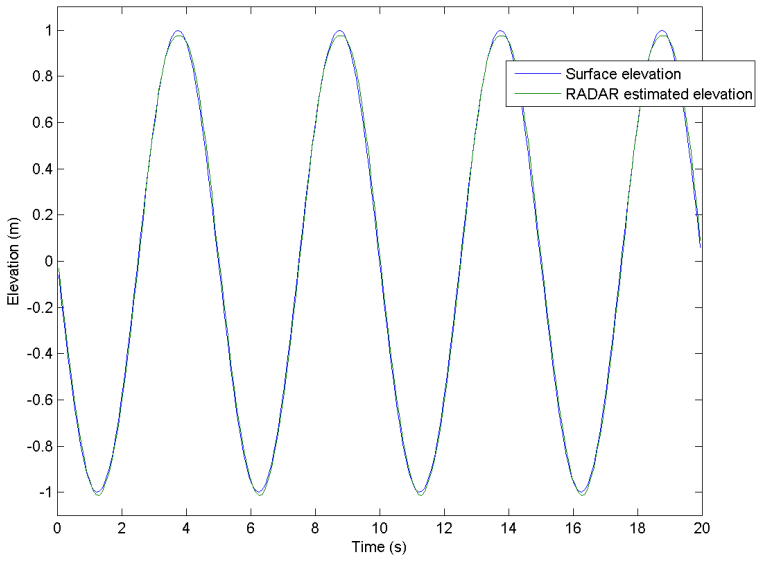

Figure 7: Time series of the surface elevation of a simple plane wave with frequency 0.2 Hz and amplitude of 1 m. The blue curve is the actual surface elevation directly beneath the radar at the point (0,0), and the green curve is the simulated radar signal.

The spectra of the plane wave signals are given in Figure 8. The blue curve is the spectrum of the actual surface, and the green is the radar spectrum. The left hand plot has a linear scale, and the right hand plot has a log scale. No discernible difference is apparent in the linear plot, but the log plot shows a noise floor in the radar spectrum some six orders of magnitude higher than the spectrum of the actual surface. Peaks corresponding to harmonics of the signal are also apparent in the radar spectrum, but the largest of these is still two orders of magnitude less than the signal peak. The side lobes of the spectrum of the actual surface are associated with the Blackman-Harris window used in the spectral analysis.

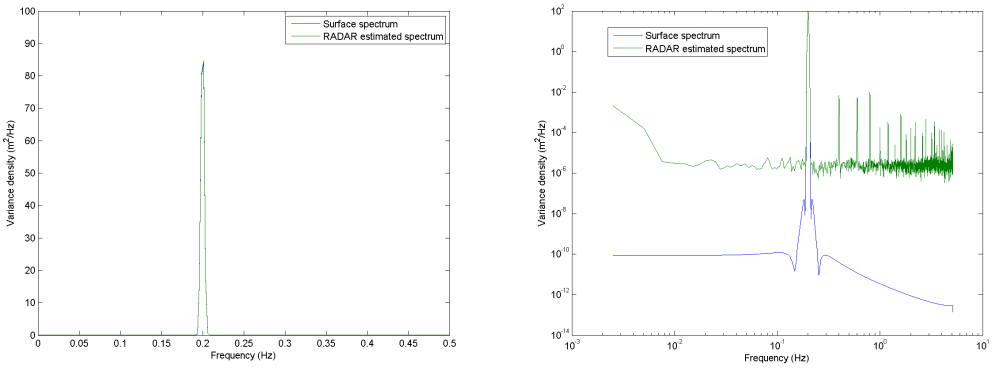

Figure 8: The spectra of the plane wave. The blue curve is the spectrum of the actual surface, and the green is the radar spectrum. The left hand plot has a linear vertical scale, and the right hand plot has a log vertical scale.

4.2 Random Linear Surface Elevation

An example of the results at an instant of time is given in Figure 9. The figure shows contour plots of the surface elevation (top left), reflected signal gain (top right), range of the surface (lower left), and a plot of the gain-range distribution and density (lower right), at 1.36 seconds (corresponding to the first crest in the time series - see Figure 18) after the start of the simulations. The contours show respective values at points on the surface given by the x and y coordinates. The units of the surface elevation are given in z-coordinate metres (where the WaveRadar is at position $(0,0,0)$), the units of the gain are in dB and are negative, and the units of the ranges (absolute distance from the WaveRadar to the point on the surface) in metres.

The surface elevation contour plot shows that the distance to the surface at the radar source (0,0,0) is around 18.6m. The actual value is plotted as the vertical dashed line on the gain-range plots (lower right). The contour plot of the reflected signal gain shows that the most intense signals are those that are reflected directly beneath the radar source, even though the range is less to locations at grid points around (-0.6,-1), as given in the range contour plot lower left.

The combined effect of reflected signal gain and range, from combining the gain and range in the two contour plots, is presented in the lower right plot. The plot shows that most of the reflected signal intensity falls within a narrow range, and that the maximum of the density curve, which is taken as the radar estimate is, in this case, very close to the actual range (dashed line). This result is clearly a consequence of the narrow antenna beam width of the WaveRadar, but the maximum of the gain-range plot can be expected to vary, depending on the wave field at a given time, and we have indeed observed less accurate and more accurate examples. Examples of less accurate cases are given below.

These results in Figure 9 are for a Lambertian surface reflector. The gain contour plots for the Diffuse and Specular reflectors are given in Figure 10. By comparison with the results for the Lambertian case (Figure 9), the Diffuse and Specular contours are noisier, with the degree of noise increasing with increasing specularity. However, characteristics of the pattern are consistent, with maximum gain occurring in the vicinity of the point (0,0). Apparently, this is sufficient to still give an

unambiguous maximum in the gain-range plot (Figure 11), coinciding with the point (0,0) for both the Diffuse and Specular cases.

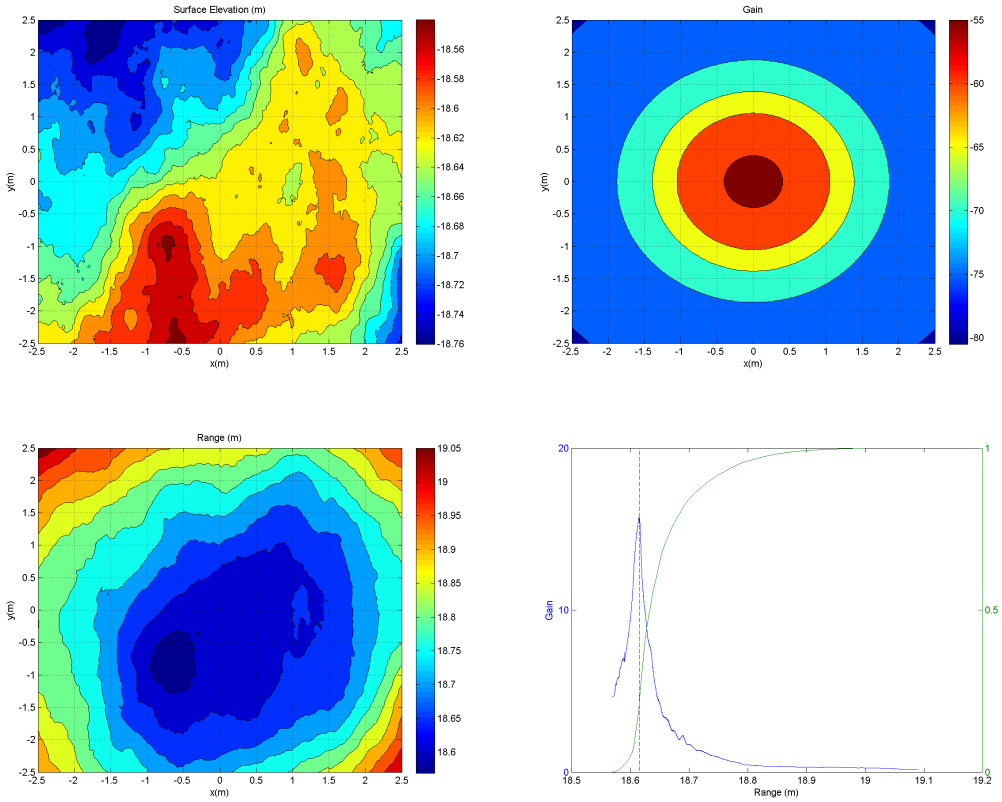

Figure 9: The surface elevation (top left), reflected signal gain (top right), range of the surface (lower left), and gain-range distribution and density (lower right), at around 1.36 seconds.

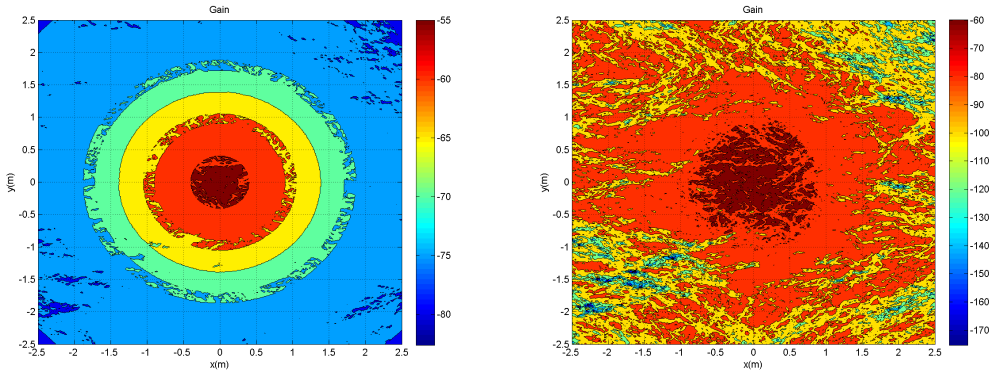

Figure 10: Gain of the Diffuse and Specular reflections for the surface wave field in Figure 9

Figure 11 gives the gain-range plots for the Diffuse (left), and specular (right) reflectors, for the surface elevation wave field in Figure 9. Both are very similar to that for the Lambertian reflector and lead to accurate estimates of the water level at (0,0).

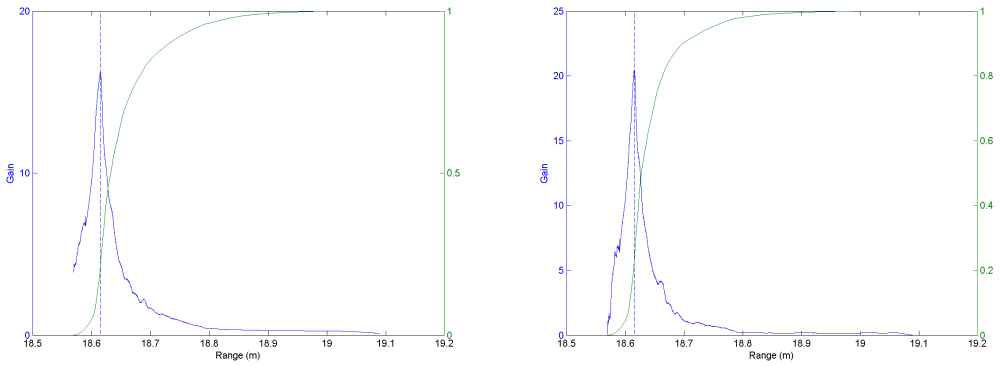

Figure 11: Gain-range plots for the Diffuse (left) and Specular (right) reflectors, for the surface elevation wave field of Figure 9.

The next example is the wave field at 52.48 seconds after the start (see Figure 18), for a case when the wave height is small and of short period. Figure 12 gives the results for the Lambertian reflector. In this case, the actual range to the point (0,0) corresponds to a local maximum in the gain-range plot. However the maximum in the plot is at a shorter range, approximately 0.1m shorter, corresponding to a point on the surface nearby but at a higher point on the wave, which is receding from the point (0,0).

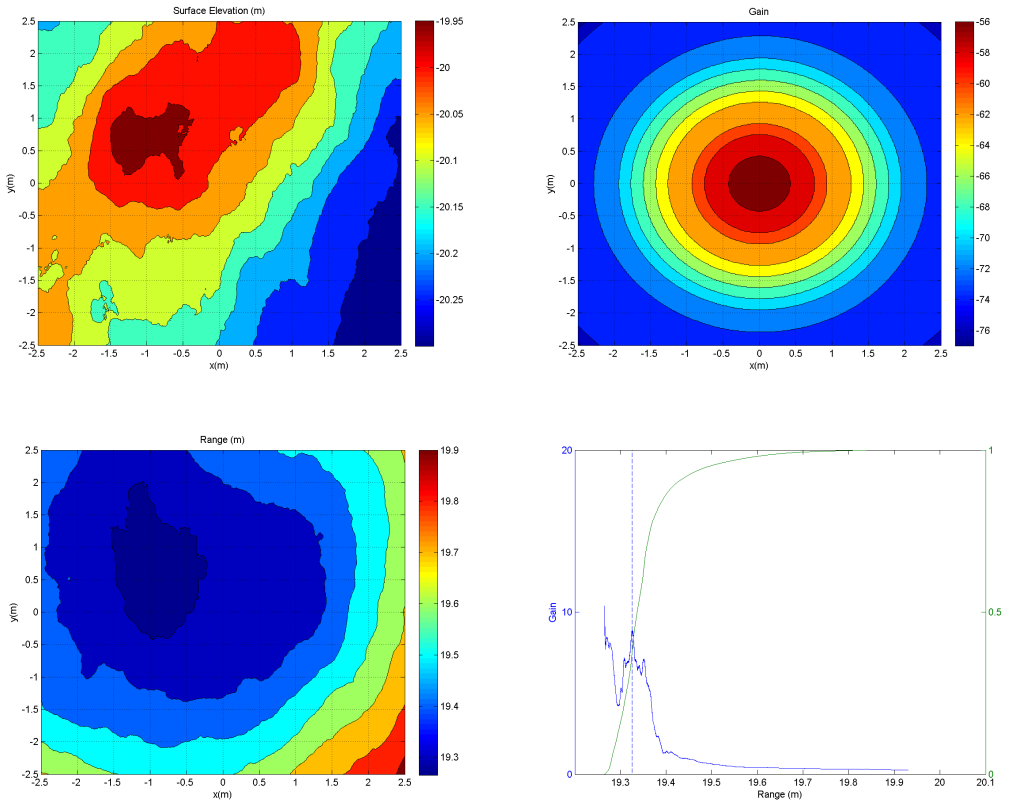

Figure 12: The surface elevation (top left), reflected signal gain (top right), range of the surface (lower left), and gain-range distribution and density (lower right), at around 52.48 seconds.

The gain and gain-range plots for the Diffuse and Specular reflected signals are given in Figures 13 and 14 respectively. The gain contour plots are qualitatively the same as for the case at 1.36 seconds after the start.

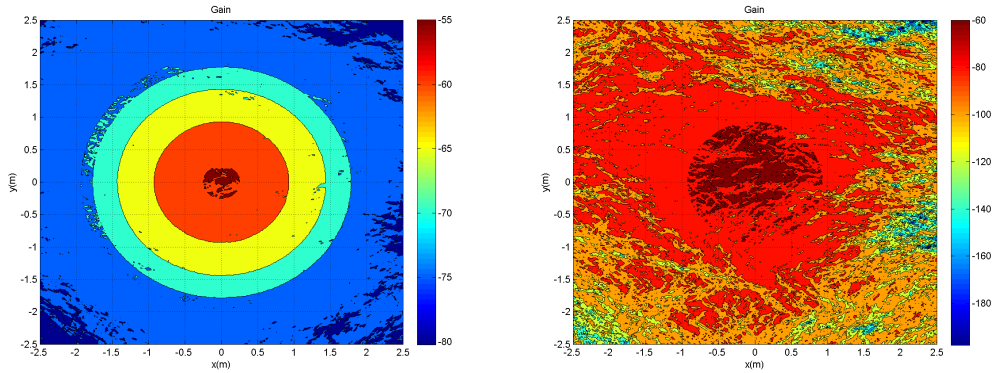

Figure 13: Gain of the Diffuse and Specular reflections for the surface wave field in Figure 12.

Similarly, the gain-range plots given in Figure 14 are qualitatively the same as for the Lambertian reflector (Figure 13). In the same way as the Lambertian case, the Diffuse and Specular reflections show a strong reflector at the point (0,0), but a stronger reflector is occurring at a higher point on the receding wave, again approximately 0.1 m closer to the radar, resulting in this point on the surface being selected as the measured range. The inaccuracy in the measured range here is a direct consequence of the method of choosing the peak in the gain-range plot. For the 1.36 s instant, simply taking the maximum was sufficient, but for this case, it introduces an error. Nevertheless, it is clear that due to the narrowness of the gain-range density function, it is not possible to get a grossly inaccurate result, irrespective of the method to pick the peaks, and this is primarily due to the narrowness of the radar antenna beam pattern. However, as we shall see later, this error introduces bias at the high frequencies.

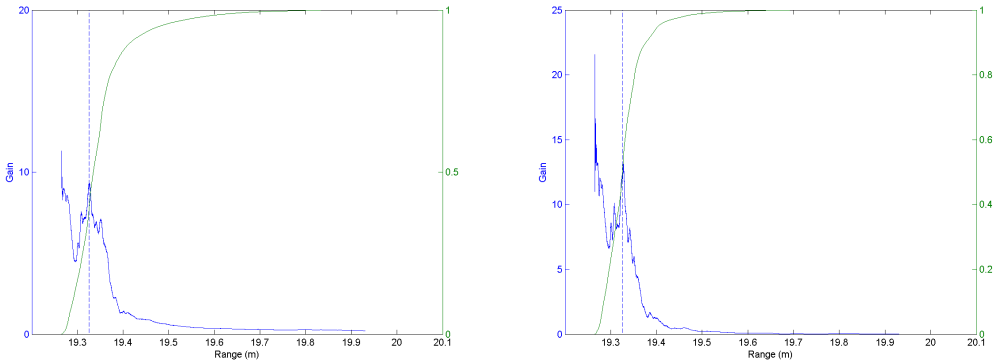

.

Figure 14: Gain-range plots for the Diffuse and Specular reflections for the surface wave field in Figure 12.

Figures 15, 16, and 17 give the results for the wave field at 64.13 seconds after the start of the record. This corresponds to the large peak in the first 100 s of the time series (Figure 18). The surface

elevation contour plot (Figure 15 upper left), shows a crest directly below the radar, but there is a more elevated local peak at (0.5,0), near to (0,0), the point of interest. The local peak is closer to the radar and provides a slightly more intense reflected signal. This is the maximum in the gain-range plots for all three reflector types (Figure 15 – bottom right, Figure 17). As for the previous case, the error is in the order of 0.1 metres.

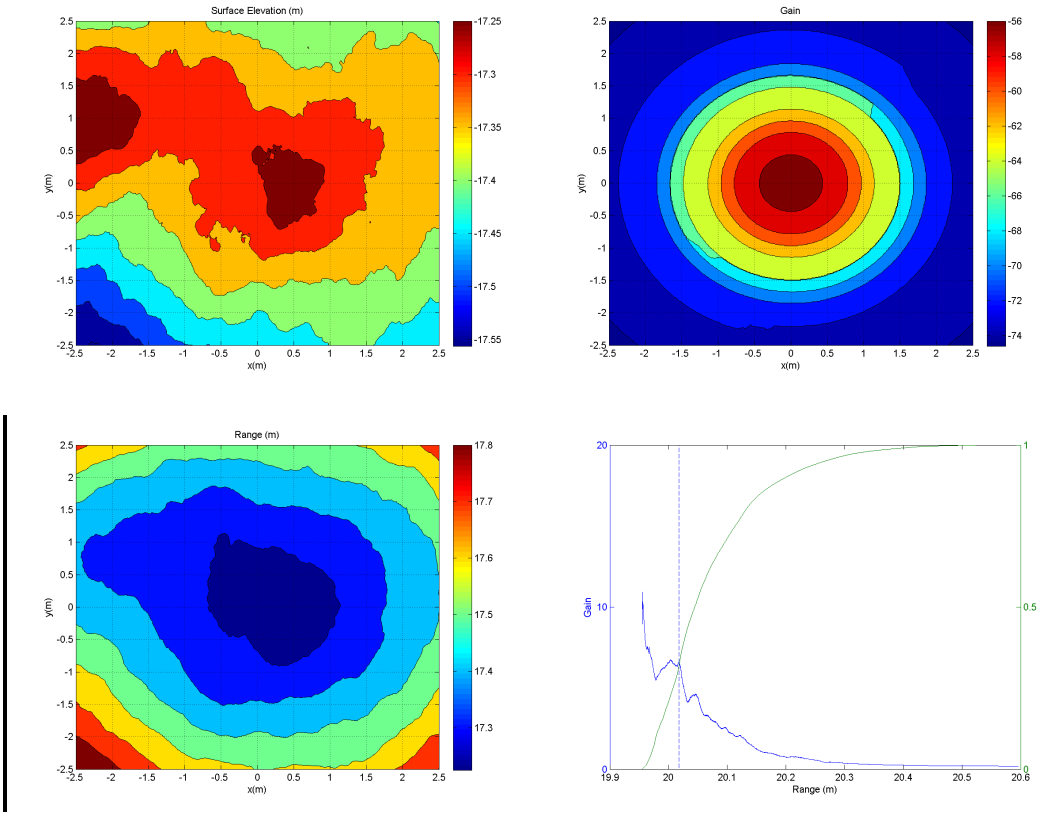

Figure 15: The surface elevation (top left), reflected signal gain (top right), range of the surface (lower left), and gain-range distribution and density (lower right), at around 64.13 seconds.

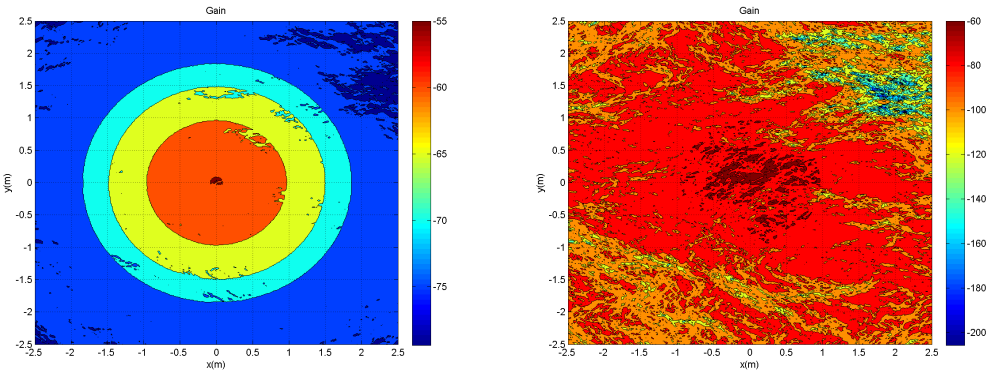

Figure 16: Gain of the Diffuse and Specular reflections for the surface wave field in Figure 15.

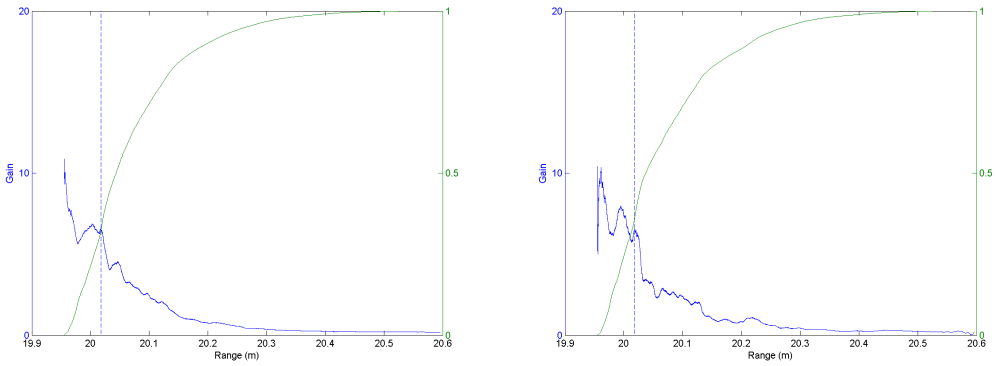

Figure 17: Gain-range plots for the Diffuse and Specular reflections for the surface wave field in Figure 15.

The surface elevation of the first 100 seconds of the simulated time series is presented in the left plot of Figure 18. The discrepancies between the actual surface and the radar simulations are barely discernible in the time series plot, but this is more apparent from the scatter plot in Figure 18 (right), which shows some spread in the simulations but virtually no overall bias. The scatter is largest for the small elevations, which tend to be biased a little high. This is apparently due to situations at times when the surface is elevated close to but not at the location (0,0), which produces a larger gain at a shorter range in the simulations. It is interesting to note from the scatter plot, that there is less scatter in the peaks and troughs of the time series, indicating the radar estimates are closer to the actual values at those times. This is important for studying the individual crest and heights in the time series.

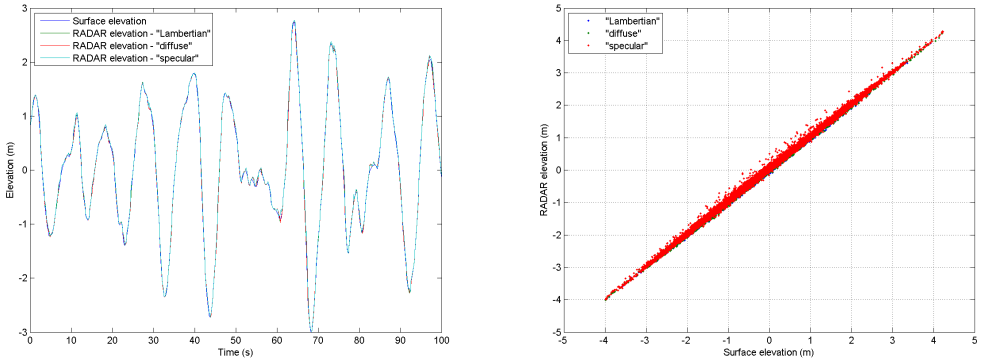

Figure 18: Plots of the time series (left) and scatter plot (right) of actual and simulations of the surface elevation for the three reflector types.

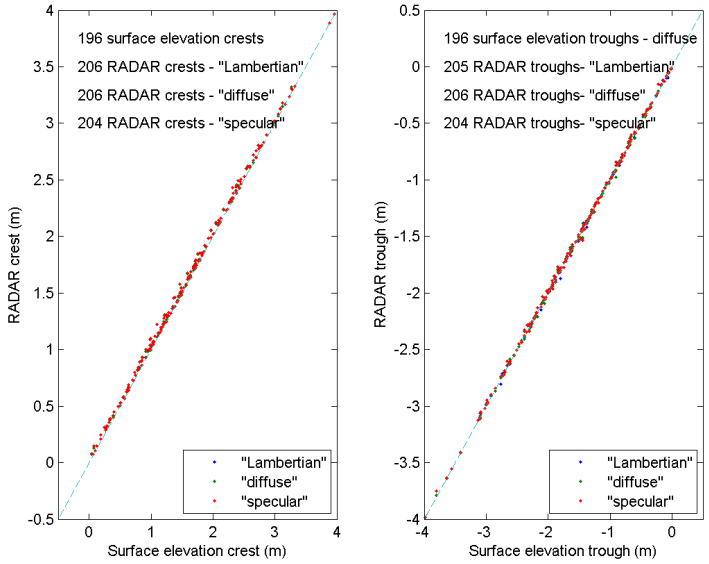

Figure 19: Scatter plots of the zero-crossing crest and trough elevations for the three reflector types.

Figure 19 shows scatter plots of the crest and trough elevations for this time series, for the three reflector types. Clearly there is some scatter in the radar estimates of the crests and troughs. There is a discernible bias in the crests but not in the troughs. The bias in the crests is likely due to cases where a local maximum in the crest occurs at a point close to (0,0) that produces a larger gain and shorter range in the simulations. However, there is generally better agreement between the actual and estimated values for the largest crest and troughs. It is also interesting to note the difference between the number of crests and troughs, given near the top of each plots. The radar estimates have less zero-crossing crests and troughs than in the actual time series, and the number for each reflector is different. This is almost certainly due to the occurrences of small crests and troughs that are not well resolved in the simulations.

The spectra are given in Figure 20. The blue curve is the spectrum of the actual surface, and the coloured are the radar spectra, for the three different reflector types, plotted on linear scales on the left and log scales on the right. As for the plane wave results, there is no discernible difference apparent in the linear plot, but the log plot shows increased levels in the radar spectrum at high and low frequencies, with departures becoming apparent for frequencies below 0.06 Hz and above 0.6 Hz. The spectral parameters are given in the left plot. The significant wave heights are identical to three significant figures, the radar first moment periods, T1, are biased up to 1% short, and the second moment periods, T2, biased 4-9% short. The departures of the parameters T1 and T2 from the simulations to those of the original surface elevations is solely due to the increased high frequency spectral levels of the radar simulations.

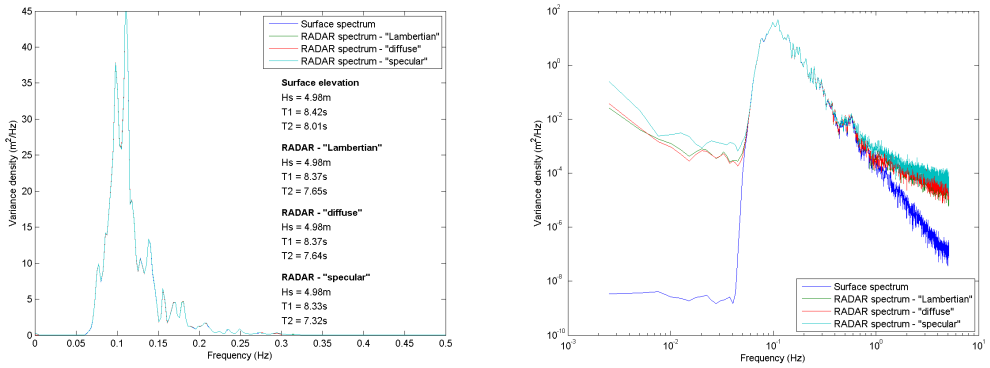

Figure 20: The spectra of the time series – linear left and log right. The blue curve is the spectrum of the actual surface, and the green, red and magenta are radar spectra for the different reflection types.

Overall the simulations show that the radar measurements of the sea surface are excellent in the vicinity of the spectral peak. At frequencies well below, and well above the spectral peak (approximately five times the spectral peak in the example presented here), the simulations are biased high, but as these are for elevations with spectral densities more than two orders of magnitude less than the spectral peak, they have little practical consequence, on the wave spectrum and the significant wave height, and their effect is relatively small on estimates of the crests and troughs. Periods calculated from the moments of the spectra are biased short on account of the elevated spectral levels of the simulations at high frequencies, but the degree of bias will depend on upper frequency used to calculate the spectral moments – if this is constrained to 0.50 Hz, little bias will occur.

The inaccuracies in the simulations result from the approach to peak fitting the gain-range density. We have simply chosen to take the maximum of the density. This introduces error when there is a more elevated section of water close to the point directly below the radar at (0,0). A more sophisticated approach will likely yield improved accuracies.

The WaveRadar simulations were repeated for a sensor height of 30 metres above mean sea level, to assess the effect of the increase in the antenna gain spread, but no noticeable change in the results occurred.

5 Field Measurements

This section involves analysis of wave data measured with WaveRadars. In Section 5.1, we present a wave spectrum for measurements acquired on a platform in the South China Sea at the maximum sampling rate of 10 Hz, and in Section 5.2 we present comparisons between measurements made with a WaveRadar and a nearby wave buoy at several locations in the North Sea.

5.1 10 Hz Sample Measurements

The 10 Hz data were recorded with a Rex WaveRadar mounted at 13 metres above mean sea level on the St Joseph Platform offshore Sabah in Malaysia. The water depth at the St Joseph Platform is 94 metres.

On 10th June, 2012, the WaveRadar was set to record at the maximum sampling rate of 10 Hz. The variance density spectrum was computed for one hour of data, starting at 16:00 hours, at which time a local wind-sea was active in the region, responding to a local wind of around 15 m/s. The spectrum is plotted on log axes in Figure 21.

The main features of this spectrum are the well defined spectral peak, and the rapid fall off in spectral level at frequencies below the spectral peak and above the spectral peak. Below the spectral peak, the spectral levels plateau to levels some three orders of magnitude below that of the spectral peak. Above the spectral peak, the spectral levels fall to levels nearly six orders of magnitude below the spectral peak. Over the frequency range 0.2 Hz to 0.8 Hz ($\approx 1.6f_p$ to $6.6f_p$) the spectral levels fall off with $f^{-4.7}$, which is consistent with observations that typically fall between f^{-4} and f^{-5} (Babanin, 2010). Accordingly, these appear to be reasonable measurements of the equilibrium range of the spectrum of wind waves. Above 0.8 Hz, the spectrum continues to roll off at a steady rate of $f^{-2.9}$, until 4 Hz ($\approx 33f_p$), when the spectral levels "kick up" to the Nyquist frequency of 5 Hz. The "kick up" is symptomatic of aliasing, but for frequencies below 4 Hz, there is no reason to suspect the spectral levels are erroneous. This being the case, the WaveRadar appears to be resolving signals at high frequency that are nearly six orders magnitude smaller than the those at the spectral peak. This is substantially better than achieved with the simulations described in Section 4.2, suggesting the WaveRadar performs better in practice than modelled in the work reported in this paper.

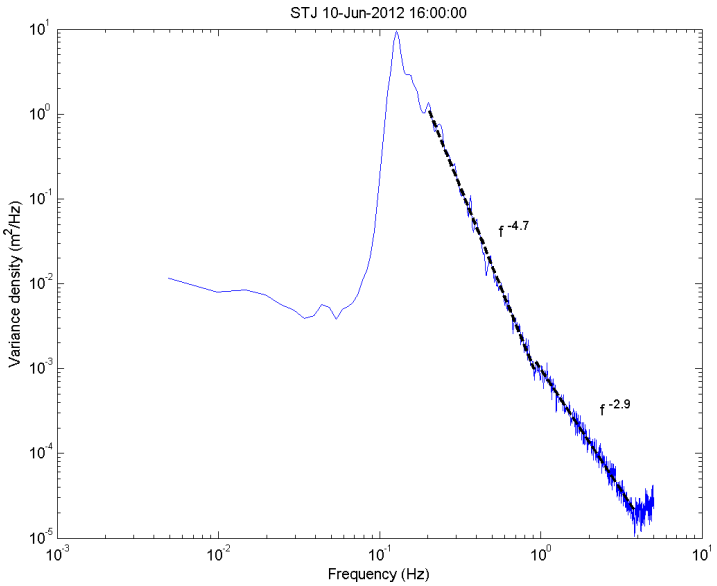

Figure 21: The wave spectrum for wave measurements at St Joseph Platform at 16:00 hrs on 10th June, 2012 local time.

5.2 Comparisons with Other Sensors

The WaveRadar measurements used for the inter-comparisons presented here were made on the North Cormorant, Gannet, and Auk platforms, at elevations of 28.7 m, 22.5 m, and 24.3 m respectively. Plan views of these and the locations of the WaveRadar are plotted in Figure 22. The North Cormorant WaveRadar location might be expected to result in measurements that have interference from the platform, for sea states with predominant directions from the north to east sector. In the case of the Gannet and Auk platforms these sectors would be southwest to north and east to south respectively.

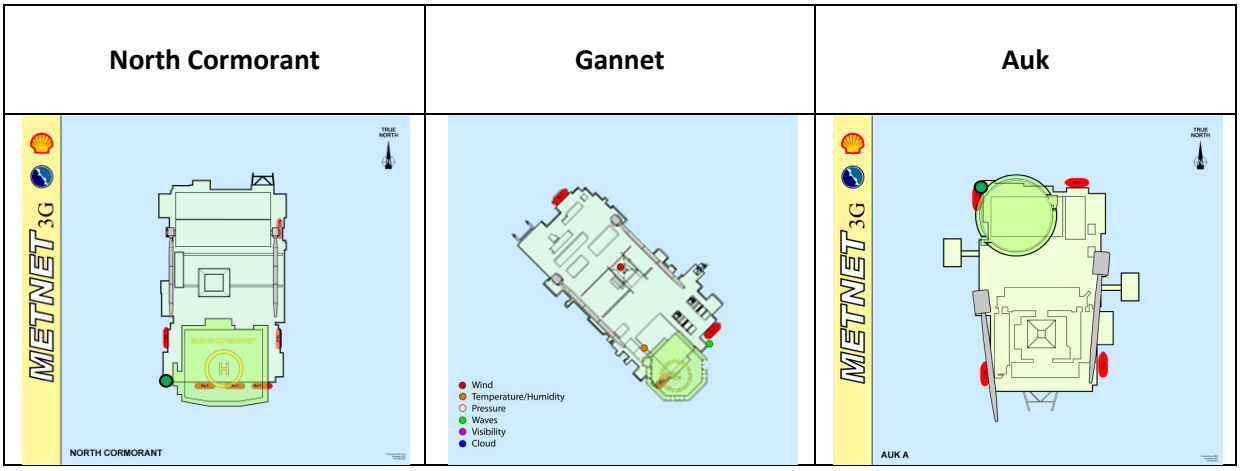

Figure 22: The locations of the WaveRadar (green dot) on the three North Sea platforms.

The wave buoy used for the comparisons at North Cormorant was located approximately seven kilometres north-northeast of the platform, and that for Auk was located about four kilometres west of the platform; but as there was no buoy located at the Gannet platform, the data from a wave buoy located at the Anasuria site, some 15 km to the northwest, were used for comparison with the Gannet platform WaveRadar measurements.

Figures 23 presents scatter plots of the significant wave height derived from measurements made at the North Cormorant facility with a Datawell Wavec buoy and SAAB WaveRadar, and Figure 24 presents scatter plots for the significant wave height from the Wavec buoy and a SAAB Rex WaveRadar. The scatter plots are stratified by the mean wave direction derived from the Wavec buoy, and the result of a linear regression for each data set is given at the top of each plot.

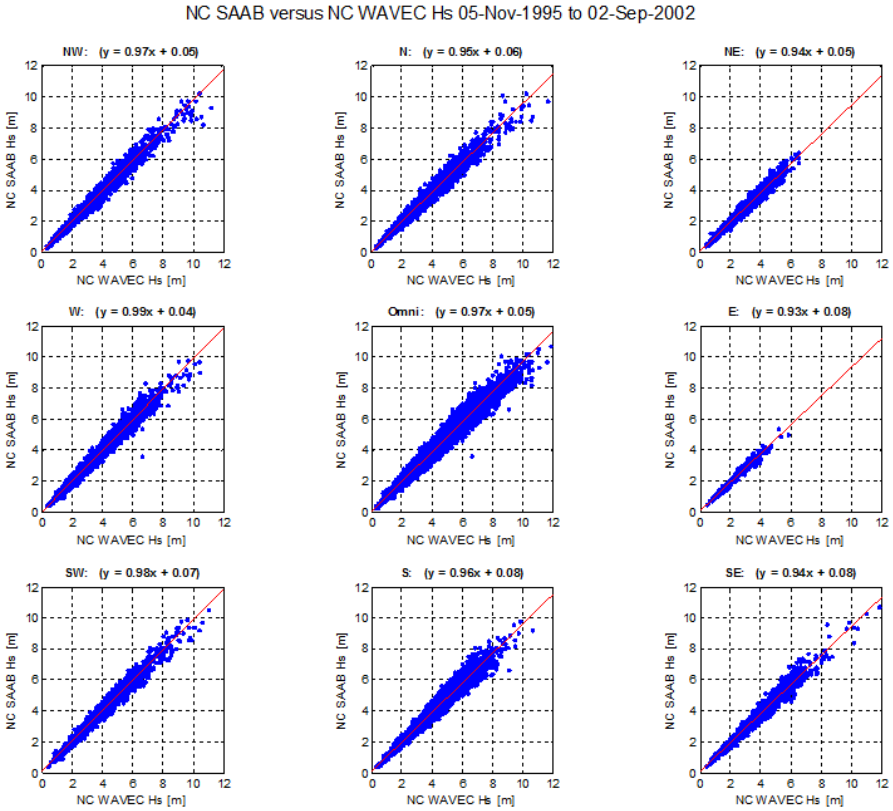

Figure 23: Scatter plots of the significant wave height measured by a Datawell Wavec buoy and a SAAB WaveRadar, at the North Cormorant platform, stratified by mean direction of the waves at the Wavec buoy.

Figure 23 shows that the regression slope varies between 0.99 and 0.93, indicating that the WaveRadar significant wave heights are less than those from the Wavec on average. This effect is smallest when the waves are from the West, the side of the platform where the radar was installed and the side most likely to be least affected by the presence of the platform. Accordingly, these measurements suggest virtually no difference between the significant wave heights derived from the Wavec and WaveRadar when platform interference of the wave field is expected to be small.

Remarkably, the comparison between the SAAB Rex WaveRadar and the Wavec show the WaveRadar to be producing significant wave heights 7% to 11% higher than Wavec, with a similar trend with wave direction as observed in Figure 23. The WaveRadar on the North Cormorant platform that produced the data presented in Figure 23 was a hybrid unit that had a SAAB processor with a SAAB Rex antenna. This is understood to be the only difference between the set up between the installations associated with the data sets presented in Figures 23 and 24. Nevertheless, the differences between the two sets of comparisons are clearly due to the specific set up of the system rather than anything related to the actual ability of the sensor to measure the wave heights.

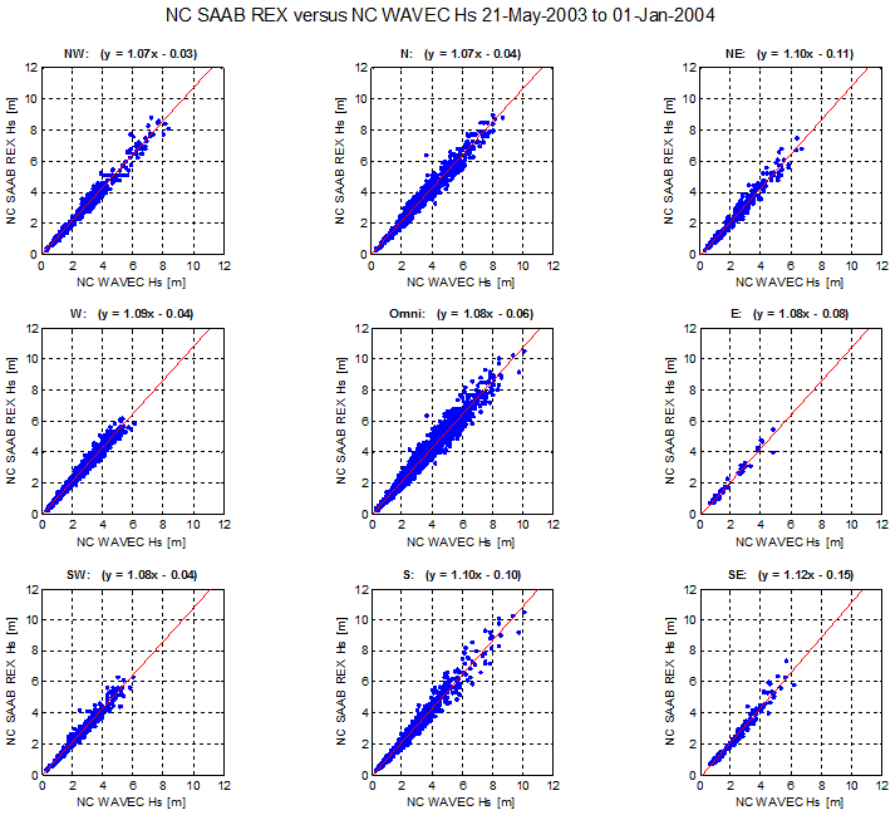

Figure 24: Scatter plots of the significant wave height measured by a Datawell Wavec buoy and a SAAB Rex WaveRadar, at the North Cormorant platform, stratified by mean direction of the waves at the Wavec buoy.

The results for the Auk platform are given in Figure 25. These results substantiate those observed in Figure 23 – that the WaveRadar and Wavec significant wave heights are essentially the same for wave directions not expected to result in significant platform interference with the wave field.

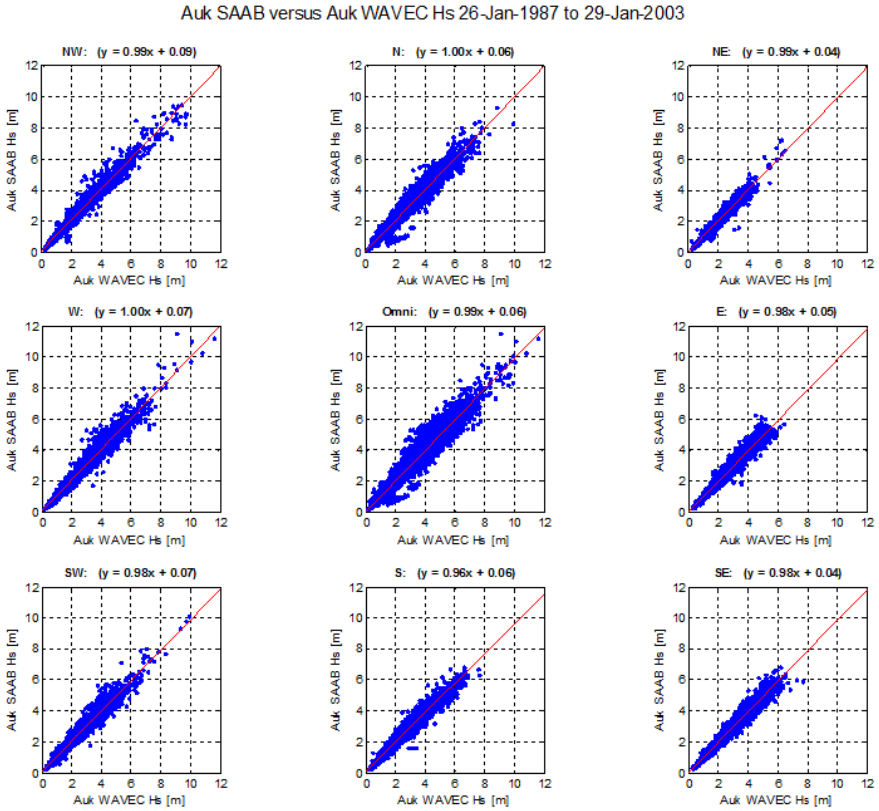

Figure 25: Scatter plots of the significant wave height measured by a Datawell Wavec buoy and a SAAB WaveRadar, at the Auk platform, stratified by mean direction of the waves at the Wavec buoy.

The comparisons between the Gannet WaveRadar and the Anasuria Directional Waverider buoy are given in Figure 26. Interestingly, the results show the significant wave heights to be lower than the Directional Waverider buoy values by 5%, fairly consistently for all directions. This may be due to actual differences between the wave fields at Gannet and Anasuria but probably not as a consequence of water depth, which is around 90 metres at each location.

Finally, it is important point to note that while the differences between the relationships from the various comparison data sets remains to be explained, we do not observe a trend for the relationships between the significant wave heights from the WaveRadars and buoys to change with the value of the significant wave height. It would however be an interesting extension to restrict the regression analysis to say the largest 10% of the sea states for each direction to confirm this.

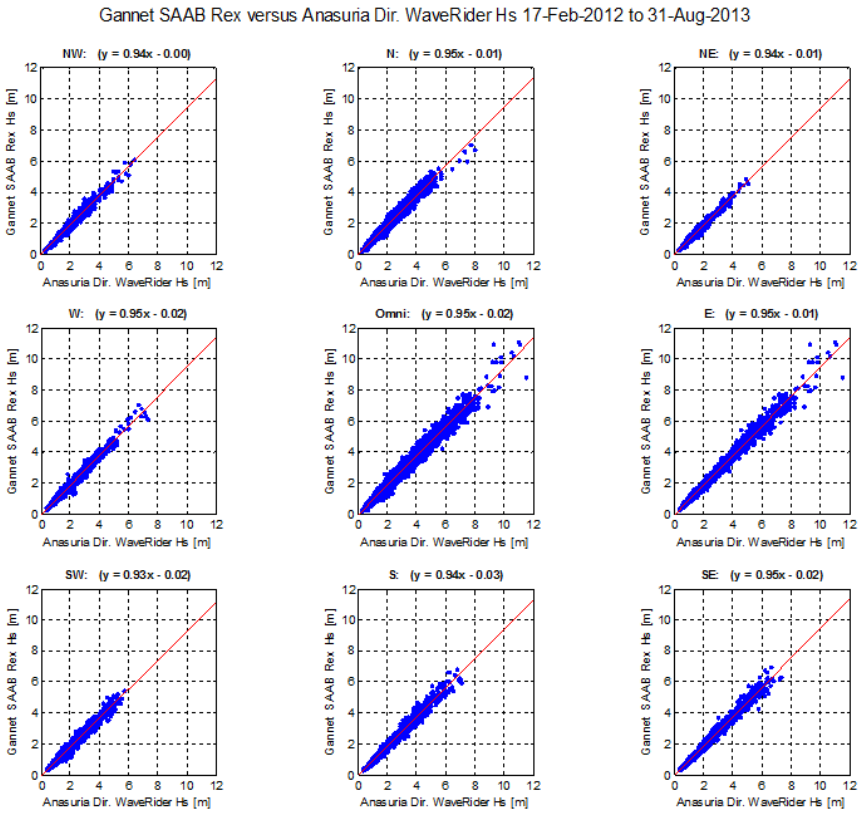

Figure 26: Scatter plots of the significant wave height measured by a Datawell Wavec buoy and a SAAB WaveRadar, at the Auk platform, stratified by mean direction of the waves at the Wavec buoy.

6 Conclusions

Our simulations of WaveRadar measurements of a random linear surface wave field indicate that the WaveRadar should faithfully follow the surface elevation at a point directly below the radar at frequencies between 0.06 Hz and 0.6 Hz. The main cause for the departures in the simulated measurements outside that frequency band is due to the particular method we have employed for processing the reflected radar signals, and especially the peak-picking method. However, this has no effect on the significant wave height, but the elevated spectral levels above 0.6 Hz can bias the spectral moment periods high by a few percent, if the calculation of the spectral moments includes frequencies above 0.6 Hz.

Similarly, the departures in the simulated measurements outside that frequency band have no appreciable effect on the calculated zero-crossing crests and troughs, though a small spread is seen for small values of those parameters.

The field measurements made at a sampling frequency of 10 Hz indicate that the WaveRadar performs much better than our simulations suggest, with the expected roll-off in spectral density continuing to much higher frequencies than the simulations, and the low-frequency plateau occurring an order of magnitude lower relative to the spectral peak.

The comparisons of the significant wave height of WaveRadar measurements against Datawell wave buoy measurements made in the North Sea generally show fairly good agreement. Comparisons between the earlier WaveRadar units and the Wavec buoy are in very good agreement for wave directions not expected to be affected by the platform and small reductions in the WaveRadar values compared with the buoy values for directions expected to be affected by the platform. Reduced significant wave heights due to platform interference on the wave field are expected.

Comparisons of the Rex WaveRadars against the wave buoys show systematic differences in the significant wave height in some cases, though the differences are relatively small (~10% at worst). This cannot be explained by platform interference, but appears to be more related to the specific setup of the instrumentation.

The comparisons do not show any trend with increasing significant wave height.

Finally, we acknowledge that further investigation of the field measurement setups is needed, to explain the differences between the results of the various inter-comparisons, and improvements in the WaveRadar "peak-picking" algorithm might lead to improved fidelity of the simulations; but the results of this study support the conclusion that the WaveRadar provides good measurements of the surface wave elevation, not only for supporting offshore operational activities and engineering requirements but also for investigations into fundamental aspects of ocean surface waves.

Acknowledgements

The work described in this paper would not have been possible without the assistance of staff from Emerson, who made time available to meet with one of us in Linköping for open discussion on the functionality of the WaveRadar; and in particularly, we would like to thank Hakan Bertling and Jan Westerling for providing additional information and responding to further questions following the meeting.

References

Babanin, A.V., 2010: Wind input, nonlinear interactions and wave breaking at the spectrum tail of wind-generated waves, transition from f^{-4} to f^{-5} behaviour. *Ecological safety of coastal and shelf zones and comprehensive use of shelf resources. Proceedings "To the 30th anniversary of the oceanographic platform in Katsiveli",* Marine Hydrophysical Institute, Institute of Geological Sciences, Odessa Branch of Institute of Biology of Southern Seas.- Sevastopol, 21, 173-187

Christou, M., and K. C. Ewans, 2011: Examining a comprehensive data set containing thousands of freak wave events. Part 1 – description of the data and quality control procedure. *Proceedings of OOAE2011 30th International Conference on Ocean, Offshore and Arctic Engineering*. 19 – 24 June, 2011, Rotterdam, The Netherlands.

Ewans, K. C., 1998. Observations of the directional spectrum of fetch-limited waves. Journal of Physical Oceanography, 28, pp. 495–512.

Liu, Y., Su, M., Yan, X., and W. Liu, 2000: The Mean-Square Slope of Ocean Surface Waves and Its Effects on Radar Backscatter. *J. Atmos. and Ocean. Technology*, 17, 1092-1105.

Noreika, S., Beardsley, M., Lodder, L., Brown, S., and D. Duncalf, 2011: Comparison of Contemporaneous Wave Measurements with a Rosemount Waveradar REX and a Datawell Directional Waverider Buoy. *12th International Workshop on Wave Hindcasting and Forecasting & 3rd Coastal Hazard Symposium*. 31 October – 4 November, 2011, Kohala Coast, Hawaii.

R.S. Aqua Ltd, 2009: WaveRadar Rex Operating Manual – issue J.