International Journal of Offshore and Polar Engineering
Vol. 1, No. 1, March 1991 (ISSN 1053-5381)

Real Time Estimation of Waves and Drift Forces

J. Knoop
Department of Mechanical Engineering, University College London, London, United Kingdom

R. Eatock Taylor
Department of Engineering Science, University of Oxford, Oxford, United Kingdom

ABSTRACT

It is demonstrated how information about the undisturbed wave at all frequencies of interest can be retrieved by constrained deconvolution of the outputs of several sensors with different known linear response characteristics, placed on or around a floating structure. A model of the sea state which retains information about the amplitude, phase and direction can then be used to calculate the resulting non-linear drift forces by means of known quadratic transfer functions.

It is shown how, with certain limitations, the wave and drift force calculation can be implemented in real time using frequency domain techniques. The resulting force time history could be used as feed forward information by the controller of a dynamic positioning system, to improve performance and efficiency. This could be further enhanced by applying linear prediction techniques (e.g. an autoregressive model) to the force time history.

INTRODUCTION

The dynamic position control problem arises because wave frequency forces are too large to be counteracted, and varying the demanded thrust at these frequencies would cause needless wear and tear on the thrusters. Simple low pass filtering would introduce a lag that would tend to degrade the performance of the system. The most widely used technique to overcome this is Kalman filtering, which was presented in the context of dynamic positioning by Balchen, Jensen and Saelid (1976). The method uses a state space model of the first order dynamics of the vessel to separate out the second order responses and predict them for one or more steps ahead.

What is more desirable, however, is to estimate the instantaneous forcing, so that the appropriate preventative action can be taken before the vessel starts moving off station. Some systems already include wind and current feed forward information. The concept of wave feed forward was originally proposed by Pinkster (1978), who demonstrated the effectiveness of using probes around the vessel to estimate the integral of the relative wave height around the waterline. He showed that this related to one component of the mean drift force and it was used successfully to improve position keeping.

The aim of this work is to extend the idea so that all components of the drift force can be included, and to leave open the possibility of taking into account the effects of wave directionality. The procedure falls into three distinct stages: Reconstructing the incident wave from measurements made on board the vessel, estimating the drift force using the full quadratic transfer function, and finally predicting the force itself into the near future to maximize the value of the control information. All three stages need to be carried out simultaneously and in real time.

Reconstruction of the incident wave from onboard measurements can generally be classed as inverse filtering or deconvolution, and the associated problems relate primarily to the loss of information at frequencies where the vessel does not respond; the unscrambling of the diffracted and radiated waves if wave height measurements are to be used; and the implementation in real time. As described in the subsection on Inverse Filtering under the Theory section, these have been overcome by using constrained frequency domain deconvolution on a set of sensors with different response characteristics which cover the full spectrum of wave energy. These characteristics are specified by means of frequency domain transfer functions, which can be fully determined by diffraction analysis. As well as vessel responses, the total wave elevation (including scattered and radiated waves) at any point around the vessel can be evaluated for any incident wave direction (Eatock Taylor and Sincock, 1989). Due attention has been given to coping with the effects of using finite length data samples and non-causality of the inverse filter. The formulation provides for the design of a set of stable frequency or time domain filters, the outputs of which can be summed to give a full estimate of the incident wave.

By using a weighted sum of several sensor outputs which are spatially separated or sensitive to wave direction, it is possible to identify the predominant wave direction. Although not presented here, the formulation can also be extended to reconstruct waves from different directions, by conventional or adaptive beamforming techniques (Knoop, 1990).

Having developed the means of measuring the incident wave, it is then possible to calculate the drift force using quadratic transfer functions. The theoretical approach is briefly described in the subsection on Calculation of Drift Forces under the Theory section, whilst the Real Time Implementation section examines the implications of using it in real time, and ways of establishing the range of time from the past up to "time now" over which the estimate of the drift force would be valid. The quadratic transfer function may be smoothed to extend this range, but at the expense of reduced frequency resolution. In the limit, if the quadratic transfer function is entirely flat (independent of frequency), the bifrequency calculation reduces to a squaring process which operates only on the current input sample. In many applications, the filtered square of

Received May 3, 1990: revised manuscript received by the editors July 26, 1990.
KEY WORDS: Inverse filtering, non-linear forcing, linear prediction, wave feed forward.

the wave amplitude may be sufficient. The scale factor or mean drift operator could then be chosen to suit the frequency content of the incident wave. Using the same programs, it is therefore possible to vary the level of sophistication of the drift force calculation to suit the accuracy with which the transfer functions can be determined.

The final stage is then to take the most recent reliable estimate of the drift force which is available and try to predict its variation up to the present, and hopefully some time into the future as well. The theory involves fitting an autoregressive model to all the previous estimates of the drift force time history. Once a given set of linear prediction coefficients has been established by autocorrelation matching techniques (Kay and Marple, 1981), it can continue to be used at subsequent steps to rapidly update the prediction on the basis of the latest estimates. The term linear prediction relates to the linear combination of previous samples and is substantially more powerful than simple linear extrapolation, with wide applications in Geophysics and Speech Synthesis (Makhoul, 1975). The time consuming task of fitting the model parameters need only be carried out at appropriate intervals as the sea state changes. Details are given by Mitchell et al. (1989). Some results of applying this technique are described in the Results section. These include both experimental data and simulations, which indicate the feasibility of the suggested approach.

THEORY

Inverse Filtering

The response of a vessel to a unidirectional irregular wave can be considered as the output of a filter operating on a single random input. The vessel is characterized by its impulse response function. The impulse response has the effect of convolving or smearing the input such that the response is effectively the linear superposition of the individual responses to a train of impulses of varying size. For a discretely sampled system this is written as:

$$y_n = \sum_{j=-\infty}^{\infty} h_j x_{n-j} \tag{1}$$

where x_n is the sequence of the input samples at a time interval of Δ seconds, y_n the output, and h_j the discrete time impulse response function.

The inverse operation of solving for all the values of x_n which give a known set of y_n, is not immediately possible. It can, however, be achieved by transforming the signals into the frequency domain such that the above equation becomes:

$$Y(\omega) = H(\omega)X(\omega) \tag{2}$$

The sequence x_n can then be retrieved by inverse Fourier transformation.

$$x_n = \int_{-\omega_N}^{\omega_N} \frac{Y(\omega)}{H(\omega)} e^{i\omega n\Delta}\, d\omega \quad H(\omega) \neq 0 \tag{3}$$

where ω_N is the Nyquist frequency.

There are two major problems with this equation. The first is that infinite frequency resolution is not generally available, because one has to work with finite length data sequences. In practice, the integration becomes a summation over a discrete set of frequencies. The more fundamental problem, however, is that the transfer function is very likely to be small if not zero at some frequencies. This has the effect of amplifying any noise in the measured responses.

This can be overcome by imposing a constraint on the input to be estimated, that is to assume that the signal only contains frequencies at which the transfer function is non-zero (or greater than a threshold δ).

The constraint is an operator such that, for the frequencies of interest the constrained signal is equal to the original, but outside this range the constrained signal $X'(\omega)$ is set to zero:

$$X'(\omega) = B(\omega)X(\omega) \text{ where } B(\omega) = 1 \text{ if } |H(\omega)| > 6, \quad \text{otherwise } B(\omega) = 0 \tag{4}$$

Thus, with one response measurement, it is only possible to reconstruct the constrained input, rather than any possible input. However, if several different outputs to the same input are available, each with its own frequency response characteristics, it is possible to combine them to build up the complete signal. In the text the m x 1 vector of outputs or responses will be referred to as $\{Y(\omega)\}$. The symbol ~ will be used to denote a matrix in the equations. Hence:

$$\underset{\sim}{Y}(\omega) = \underset{\sim}{H}(\omega)X(\omega) \tag{5}$$

where $\{H(\omega)\}$ is the vector of transfer functions that relate the incident wave to each of the respective responses to be considered.

In the frequency domain, the new estimate of the constrained input could be given by:

$$X'(\omega) = \underset{\sim}{A}(\omega)^T X(\omega) \tag{6}$$

where $\{A(\omega)\}$ is the corresponding vector of finite amplitude constrained inverse filters such that:

$$\underset{\sim}{A}(\omega) = \begin{bmatrix} B_1(\omega)/H_1(\omega) \\ \vdots \\ B_m(\omega)/H_m(\omega) \end{bmatrix} \tag{7}$$

The estimate $X'(\omega)$ will be correct for any frequency where

$$\underset{\sim}{A}(\omega)^T \underset{\sim}{H}(\omega) = 1 \tag{8}$$

If this is not possible, $\{A(\omega)\}$ is chosen such that the right hand side of Eq. 8 is zero, in which case the estimate will not contain any components at that frequency. Rather than setting $B_j(\omega)$ to either one or zero, there is in fact an infinite number of combinations, including complex numbers which can satisfy the above constraint equation. This possibility can be used to extend the method to filter waves from different directions, using the same set of sensor outputs (Knoop, 1990).

The reason for going to these lengths is that we have now described the inverse filtering problem in terms of a linear combination of the outputs of the filters $A_j(\omega)$, acting on the responses $Y_j(\omega)$. Using the inverse Fourier transform, the discrete time history of the constrained signal is now given by:

$$x'_n = \int_{-\omega_N}^{\omega_N} \underset{\sim}{A}(\omega)^T \underset{\sim}{Y}(\omega)\, e^{i\omega n\Delta}\, d\omega \tag{9}$$

In the time domain the above equation can be written as a convolution sum.

$$x'_n = \sum_{j=-\infty}^{\infty} \underset{\sim}{a}^T_j \underset{\sim}{y}_{n-j} \tag{10}$$

where $\{y\}_n$ is the vector of the set of measured responses at time interval n, and $\{a\}_j$ is the set of impulse responses filters obtained by inverse Fourier transformation of each element of $\{A(\omega)\}$.

There are therefore two routes to obtaining the original input or wave: either by inverse transformation of the frequency domain estimates, or by the linear superposition of the time histories obtained from the convolution of each output with its constrained time domain inverse filter.

Calculation of Drift Forces

Given a time history of the incident wave, several ways of estimating the drift force on the vessel have been proposed. The simplest method is to assume that the drift force is proportional to the square of the wave height elevation. Depending on the level of sophistication, the constant of proportionality can be made dependent on the instantaneous frequency of the wave envelope (Hsu and Blenkarn, 1970).

The more general approach, however, is to use the second order Volterra series model to characterize the response of a non-linear system in terms of a quadratic impulse response function, which is convolved with the product of the input at different time lags. This was used by Dalzell (1977) for calculation of added wave resistance for ships with forward speed. Other authors — e.g., Sincock (1990) — have applied similar techniques in the frequency domain to calculate slow drift behavior of compliant structures. In this case, the quadratic transfer function (QTF) describes how different frequencies in the input(s) interact to produce a response at the sum or difference of the two frequencies. When this approach is adopted, it is convenient to associate low frequency drift behavior with the sum frequency components (based on positive and negative frequencies).

The equation for the second order drift force is obtained by first defining the two-dimensional function which is the product of the wave inputs at discrete time lags:

$$x(n_1, n_2) = x_{n_1} x_{n_2} \tag{11}$$

The second order drift force is then given by the two-dimensional convolution of the input with the quadratic impulse response function $h(n_1, n_2)$:

$$f(n_1, n_2) = \sum_{j=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} h(i, j)\, x(n_1 - i,\, n_2 - j) \tag{12}$$

The resulting function has a value at every point in the n_1, n_2 plane, but the points of interest are where $n_1 = n_2$.

The above equation can then be transformed into the two-dimensional frequency domain so that:

$$F(\omega_1, \omega_2) = H(\omega_1, \omega_2)\, X(\omega_1, \omega_2) \tag{13}$$

where $H(\omega_1, \omega_2)$ is the quadratic transfer function, which could be obtained by the two-dimensional Fourier transform of $h(n_1, n_2)$. Thus the full expression for the drift force in the time domain is the double inverse Fourier transform:

$$f(n_1, n_2) = \int_{-\omega_N}^{\omega_N} \int_{-\omega_N}^{\omega_N} F(\omega_1, \omega_2) e^{i\omega_1 n_{1\Delta}} e^{i\omega_2 n_2 \Delta}\, d\omega_1\, d\omega_2 \tag{14}$$

Since this only needs to be evaluated along the line $n_1 = n_2$, the calculation can be reduced to a single Fourier transformation along the sum frequency line $\omega = \omega_1 + \omega_2$ and a summation over ω_2:

$$f(n) = \int_{-\omega_N}^{\omega_N} \int_{-\omega_N}^{\omega_N} H(\omega_1, \omega_2)\, X(\omega_1, \omega_2) e^{i(\omega_1 + \omega_2)n\Delta}\, d\omega_1\, d\omega_2 \tag{15}$$

If $\omega = \omega_1 + \omega_2$ and $\zeta = \omega_2$, we can define the sum frequency component of the drift force F(ω) as:

$$F(\omega) = \int_{-\omega_N}^{\omega_N} H(\omega - \zeta, \zeta) X(\omega - \zeta, \zeta)\, d\zeta \tag{16}$$

which gives the simple expression for the drift force in the time domain as a single inverse Fourier transform:

$$f(n) = \int_{-2\omega_N}^{2\omega_N} F(\omega) e^{i\omega n\Delta}\, d\omega \tag{17}$$

Together, the discrete forms of Eqs. 16 and 17 represent the frequency domain method for calculating the drift force at sample instant *n*.

REAL TIME IMPLEMENTATION

It is necessary to examine the impulse response function of the inverse filter more closely in order to ascertain the conditions in which real time implementation is possible. In practice, this does not have infinite extent, neither into the past nor into the future. The limits of the convolution sum can therefore be changed with no loss of accuracy:

$$x_n = \sum_{j=-f}^{p} a_j y_{n-j} \tag{18}$$

if $\quad |a_j| < \varepsilon$ for $j < -f_a$ and $j > p_a$

where ε is a practical threshold below which the function is assumed to be zero. Thus f_a and p_a can be considered as the future and past extent respectively of a given inverse filter. Note that the future extent f_a should not be confused with the drift force time history $f(n)$.

The overall extent or memory of the impulse response will depend on the frequency resolution of the system, manifested as the steepest gradient in the frequency domain transfer function. Hence the rectangular window $B(\omega)$ implied by Eq. 4 will theoretically have infinite extent because of the step change in the frequency response. The method of overcoming this problem, used in the design of finite impulse response (FIR) digital filters, is to smooth the sharp edges of the frequency domain transfer function, so that there is a gradual transition from zero to unity gain. Frequency resolution is compromised, but with a significant reduction in the length of the convolution filter. The same technique can be applied to the inverse filtering problem using several outputs because, as mentioned above, the value of $B_j(\omega)$ is not limited to either 1 or zero. Instead of taking only one of the responses to estimate a particular frequency band in the input, a weighted sum of different responses is used, so that the non-zero portions of $B_j(\omega)$ overlap, and each element of $\{A(\omega)\}$ varies smoothly with frequency.

As well as the factors mentioned above, the future extent of the inverse filter will depend on the phase lag of the measured responses relative to the incident wave. Hence non-causal effects can be minimized by using responses which lead the wave elevation as measured at the center of gravity of the vessel, e.g., wave properties near the bow in head seas.

In order to carry out frequency domain convolution, the input is assumed to be periodic, with a cycle of *N* samples, and the impulse response function is extended to *N* samples and arranged in wrap-around order such that $h_{-j} = h_{N+1-j}$. Both functions are then transformed by means of the fast Fourier transform, and the resulting Fourier coefficients are multiplied term by term. The inverse transformation of the resulting Fourier series is the circular convolution of the two input functions. The circular convolution differs from the desired linear convolution in two respects. The first p_a samples of the output are corrupted because the past extent of the impulse response function includes the end of the previous cycle of the assumed periodic signal, rather than the true signal in the convolution sum. Similarly, the last f samples are corrupted because the future extent of the impulse response includes the start of the next cycle of the periodic signal, rather than the true signal. It can be shown mathematically that, for a finite duration impulse response, circular convolution in the frequency domain is identical to linear convolution in the time domain in the region $p_a < n < N\text{-}f_a$ (Defatta et al., 1986).

If this procedure is to be carried out in real time, the sampled input data points are stored in a fixed length buffer. At each sample instant, the current contents of the buffer are extended to a total of N points by adding L zeroes into the future. The length of the added zeroes should be equal to the duration of the impulse response function to minimize end effects. The extended block is Fast Fourier Transformed, each Fourier coefficient multiplied by the corresponding transfer function, and the resultant series inverse transformed. Relative to "time now," the resultant output time history runs over $N\text{-}L$ past points to L points into the future. As described above, the start of the block will be inaccurate because the convolution will include the null effect of the added zeroes, which gets progressively less significant until the pth point, after which the result is unaffected by the assumption of periodicity in the input.

If the impulse response is causal — i.e., $f_a = 0$ — the wave time history will be correct right up to "time now." The extension into the future will generally be non-zero, which represents the slowly decaying response to the impulses received up to "time now." In the case of the constrained inverse filter which imposes a bandwidth limitation on the estimated input, this is quite an effective prediction into the near future (half a cycle). If the impulse response is non-causal, the implication is that future points are required for the correct result. Hence the given output for which future inputs were set to zero, is only accurate up to f_a points before "time now."

As far as the drift force calculation is concerned, rather than the true linear convolution of Eq. 12, the frequency domain method (using finite blocks of data) actually calculates the two-dimensional block convolution of the periodic input function which is analogous to the circular convolution of the one-dimensional case (Dudgeon and Mersereau, 1984). This is realistic only if the extent of the quadratic impulse response function is finite. As for a one-dimensional filter, we can define the past and future extent of the quadratic impulse response function, so that frequency domain block convolution yields the correct result in the range

$$p_q < n_1 - f_q \quad \text{and} \quad p_q < n_2 < N - f_q \tag{19}$$

where N is the block length used in the finite discrete Fourier transform, p_q the longest of the past extents or decay times of the impulse response in each dimension, and f_q the longest of the future extents or greatest non-causality.

The extent of the quadratic impulse response function can be examined by inverse transformation of the quadratic transfer function, taking into account the change in the axes to sum and difference frequency as described in Mitchell et al. (1989). As in the one-dimensional case, it is possible to reduce the extent of the quadratic impulse response by smoothing the quadratic transfer function in the frequency domain.

Thus the drift force time history will be incorrect, due to having the wrong input to the convolution for the first p_a points, and it will take the next p_q points for that error to be "forgotten." The result should be correct from there up to the end of the finite block, i.e., "time now," if both the inverse filter and the QTF are causal. If not, the force estimate will suffer from having the wrong future input for $f_a + f_q$ points before "time now."

At each cycle, the portion of the force time history that is not influenced by the end effects could then be retained and combined with the estimates from previous cycles to build up the historical time history which is used to fit the autoregressive model. Once established, this can be used as a linear predictor at each cycle so that the estimate can catch up with real time.

RESULTS

In order to demonstrate the above procedure, experiments were carried out in the wave basin at Heriot Watt University with a 1:81 scale series 60 tanker model. A wave probe was mounted on the bow of the moored vessel, and the response to a unidirectional random sea state was measured for an hour. A load cell in the front mooring line was used to give a measure of the actual forces on the vessel, and by filtering this, the slowly varying drift force was obtained as described by Mitchell et al. (1989).

With the vessel removed, the same sea state was subsequently measured by a wave probe fixed at the position where the center of gravity of the vessel had been. Offline cross spectral analysis of the measured input with the various vessel responses yielded directly the transfer functions of interest, which could be used to determine the appropriate inverse filter.

Fig. 1 shows the experimentally measured linear transfer function of the bow-mounted wave probe, plotted against normalized frequency (i.e., a fraction of the Nyquist frequency, which in this case is 0.555 Hz). Nearly all the power in the irregular wave spectrum was in range from 0.0625 to 0.125 of the Nyquist frequency. The corresponding constrained inverse filter is shown in both frequency and time domain representations in Figs. 2 and 3, respectively.

The quadratic transfer function was obtained by cross bi-spectral analysis of the measured wave time history and the filter drift force time history as described by Sincock (1989) and is shown (amplitude only) in Fig. 4. The bifrequency axes are expressed relative to the Nyquist frequency.

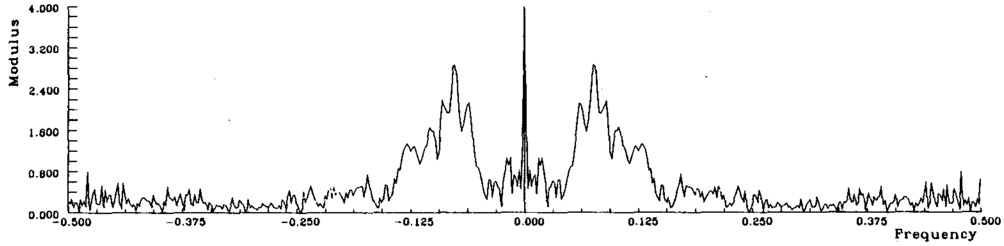

Fig. 1 Experimental bow probe transfer function

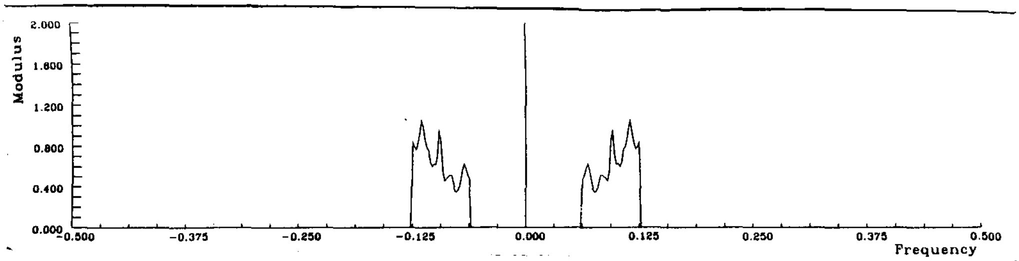

Fig. 2 Bow probe inverse filter transfer function

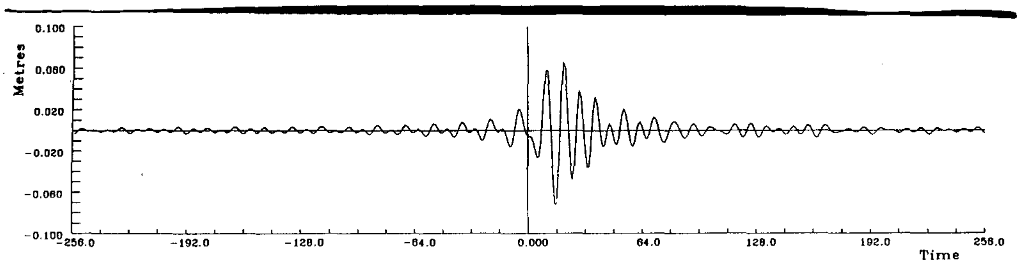

Fig. 3 Bow probe inverse filter impulse response

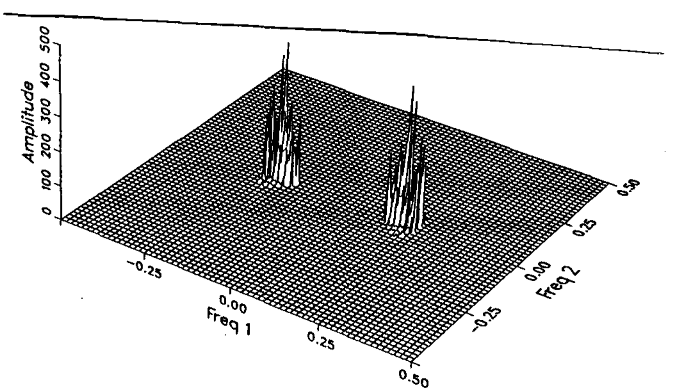

Fig. 4 Surge force quadratic transfer function (kN/m^2)

A computer simulation was implemented to read the stored records of the vessel responses and carry out the inverse filtering operation as if it were happening in real time. The constrained deconvolution of the bow probe response was carried out in the frequency domain in blocks of 512 points with 128 points of zero padding into the future.

Fig. 5 shows a snap-shot of the real time wave and drift force estimation in a unidirectional random sea state. The wave has been estimated from the measured response of the bow-mounted wave probe and is plotted relative to "time now," together with the wave as it was measured without the model in the tank. It can be seen that the wave estimate has in itself certain prediction properties. These arise in part because the probe was located forward of the center of the vessel, but also because the inverse filter decays sinusoidally, which prevents a step change in the wave at "time now" to the future wave, which is assumed to be zero. Also plotted are the measured and calculated drift forces relative to "time now"; the latter has been calculated using the wave estimate shown and the experimentally derived quadratic transfer function.

The point-by-point agreement is quite sensitive to the upper cutoff frequency for the drift force, which in this case corresponds to 33 seconds drift force period. Nevertheless, it can be seen that the estimate closely follows the form of the measured force.

Fig. 6 shows a similar plot of real time drift force calculation, except that the QTF has been determined numerically by a second order diffraction analysis. The force estimate is plotted together with the drift force that has been calculated offline in a 32000 point segment, using the same QTF and a simulated wave time history. It can be seen that in this case, calculating the force in real time does not introduce significant errors and the estimate is effectively valid up to "time now" and for some seconds into the future, due to the smoothness of the given QTF.

For the above cases, a single cycle of reading new data, Fourier transformation, inverse filtering, drift force calculation using the full QTF, inverse transformation, and plotting of the time histories was accomplished in well under 1 second on a Microvax II computer. A large portion of this time is in fact spent performing the animation of the time histories on the screen.

The above results do not include autoregressive prediction. The effectiveness of this technique has been demonstrated separately however, using both measured and artificially generated drift force time histories. In Fig. 7, a 40 pole linear prediction filter has been fitted to the first 4980 points of a 32000 point artificially generated drift force time series. The circles show the prediction of the 40 points immediately following the last point used in the model fitting process, plotted over the curve of the actual values. The process works by first calculating the one step ahead prediction as a sum of the previous 40 points, each weighted by their respective linear prediction coefficient. This is then repeated with the input points shifted along and using the previous one step ahead prediction. Thus the model can run forward to predict any number of points ahead. As can be seen from the plot, the prediction starts to diverge from the true value after about 40 points, i.e., the stage at which it is being fed by speculated points alone.

Fig. 8 gives a similar example, except that the drift force time series has been measured experimentally in the wave tank. The autoregressive model has again been fitted to the previous 4000 points, but has 100 poles. It can be seen that the prediction is effective for 100 points which, at a full scale sample interval of 0.9, is 90 seconds into the future. The same set of prediction coeffi-

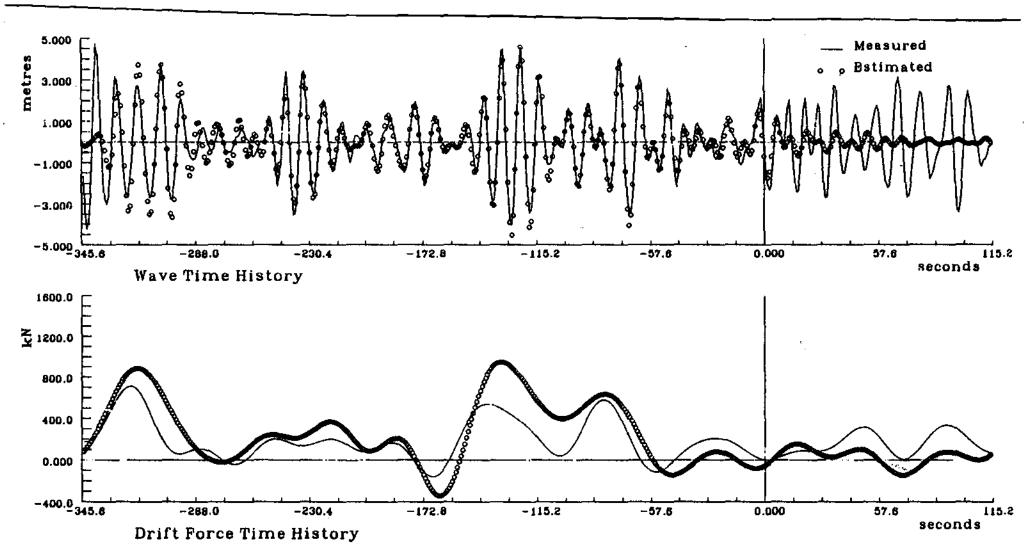

Fig. 5 Real time estimation of experimental wave and drift force

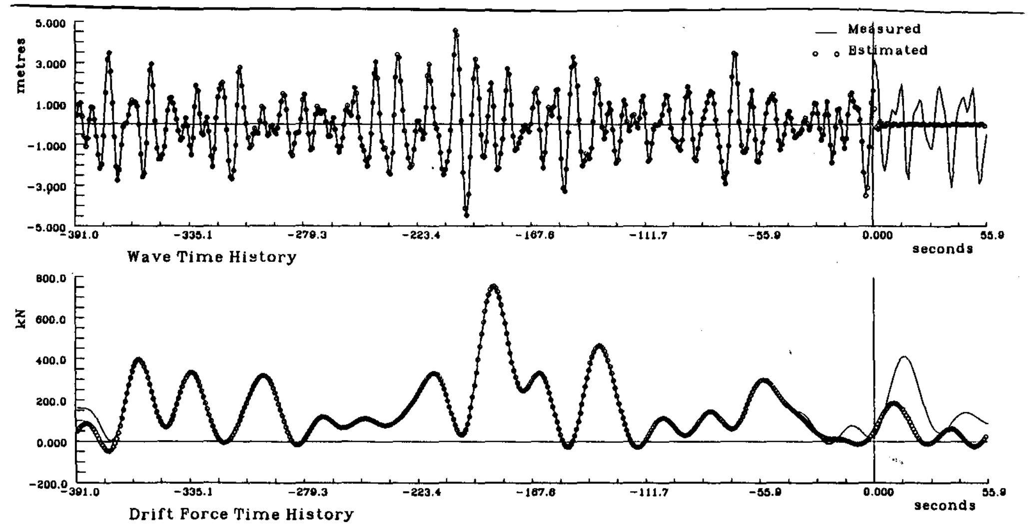

Fig. 6 Real time estimation of simulated wave and drift force

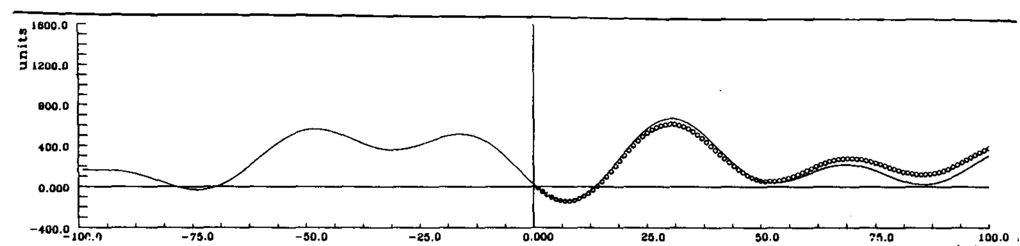

Fig. 7 Linear prediction of simulated drift force

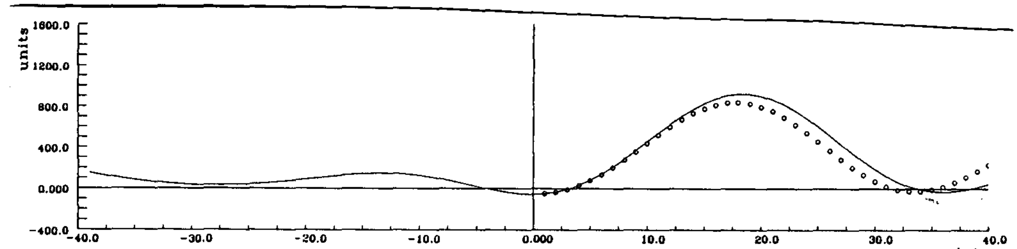

Fig. 8 Linear prediction of experimentally measured drift force

cients can be used to start the prediction at any point in the time series, either during or after the period over which the model was fitted. The time taken to fit the 100 pole model was approximately 20 seconds on the MicrovaxII computer, whilst the time required to step through the 100 point prediction was negligible compared with the full-scale sample interval.

CONCLUSIONS

The estimation of the undisturbed incident wave from a floating structure has been shown to be feasible in real time. By using constrained inverse filtering of several readily available measurements of relative wave height and vessel responses, it is possible to retrieve amplitude and phase information at all frequencies of interest.

The method of modelling the incident wave would depend on the sophistication of the motion measurement system, but in the first instance would be a long crested sea with varying direction. The formulation is sufficiently general, however, to allow for any combination of sensors and the possibility of using beamforming techniques to break down the incident wave into several discrete directions, although this technique is not yet fully proven.

The estimation of second order forces from the unidirectional wave estimate has also been shown to be feasible. The same algorithm can rapidly produce a simple wave squared approximation or, given a theoretical model of the vessel hydrodynamics in regular bichromatic waves, a more accurate estimate using the quadratic transfer function.

A causal dynamic system would generally yield a non-causal inverse filter. In addition, the quadratic impulse response may also be non-causal, such that the drift force estimate is only valid after a certain delay. By analyzing the impulse response representations of the transfer functions, it is possible to assess the extent of this delay. The indications are that it can be compensated for by fitting an autoregressive model to the slowly varying force time obtained from previous estimates, and then using this as a linear predictor.

Applying the linear prediction techniques to the measured (as opposed to estimated) force time history showed that it is possible to obtain good predictions of the force up to 1.5 cycles ahead (as much as 90 seconds).

Design of a system for a particular application would require an extensive first order diffraction analysis to establish the vessel response characteristics and the total wave elevation at selected points for all frequencies and directions of interest. A second order bi-frequency analysis would be necessary to establish the wave elevation to drift force quadratic transfer functions.

ACKNOWLEDGMENTS

This work was carried out as part of the managed program of research into Floating Production Systems (1987-89) which was jointly funded by the Science and Engineering Research Council (SERC grant GR/E/62639) and a number of industrial sponsors.

The authors would like to thank the sponsors and also to thank their colleagues at Heriot Watt University for the use of their experimental facilities.

REFERENCES

Balchen, JG, Jensen, NA, and Saelid, S (1976). "Dynamic Positioning Using Kalman Filtering and Optimal Control Theory," *Computer Applications in Shipping and Shipbuilding*, Vol 3 - Automation in Offshore Oil Field Operation, North Holland Publishing Co.

Dalzell, JF (1977). "Application of the Functional Polynomial Model to the Ship Added Resistance Problem," *Proc 11th Symp on Naval Hydrodynamics*, University College, London, pp 279-302.

Defatta, DJ, Lucas, JG, and Hodgkiss, WS (1988). *Digital Signal Processing: A System Design Approach*, John Wiley.

Dudgeon, DE, and Mersereau, RM (1984). *Multidimensional Digital Signal Processing*, Prentice Hall.

Eatock Taylor, R, and Sincock, P (1989). "Wave Upwelling Effects in TLP and Semisubmersible Structures," *Ocean Engineering* 16, pp 281-306.

Hsu, FH, and Blenkarn, KA (1970). "Analysis of Peak Mooring Force Caused by Slow Vessel Drift Oscillation in Random Seas," *Proc Offshore Tech Conf*, 1159.

Kay, M, and Marple, SL (1981). "Spectrum Analysis - A Modern Perspective," *Proc the IEEE*, Vol 69, No 11, pp 1380-1419.

Knoop, J (1990). "Real Time Estimation of Drift Forces on Floating Bodies in Directional Seas," *Thesis* to be submitted, University College London.

Makhoul, J (1976). "Linear Prediction: A Tutorial Review," *Proc the IEEE*, Vol 63, No 4, pp 561-580.

Mitchell, KL, Knoop, J, and Eatock Taylor, R (1989). "Prediction of Low Frequency Responses in Directional Seas," SERC managed program on Floating Production Systems Structures, 1987-1989.

Pinkster, JA (1978). "Wave Feed-forward as a Means to Improve Dynamic Positioning," *Proc Offshore Tech Conf*, 3057.

Sincock, P (1989). "Non-linear Compliant Systems in Irregular Seas," *PhD Thesis*, University College, London.