Coastal-Trapped Wave Mode Fitting: Reanalysis of the Australian Coastal Experiment

PETER C. MCINTOSH AND RICHARD B. SCHAHINGER

CSIRO Division of Oceanography, Hobart, Tasmania, Australia

(Manuscript received 12 May 1992, in final form 3 May 1993)

ABSTRACT

The original time domain analysis of data from the Australian Coastal Experiment involved fitting coastal-trapped wave modes to an array of velocity time series using a truncated singular value decomposition. While the truncation was necessary for noise reduction, it is shown that important information concerning the separation of mode 1 and mode 2 was discarded. A weighted least-squares mode-fitting technique is introduced that uses the data to estimate both the signal-to-noise ratio and the relative weighting of the fitted modes. In addition, the velocity data are augmented by sea-level data.

Findings from the present analysis differ in several important respects from the original results. It is found that mode 1 has approximately twice the energy flux of mode 2 and that mode 3 is statistically insignificant at the southern end of the East Australian waveguide. In addition, mode 1 is not highly correlated with mode 2. These differences are primarily due to changes in mode 1; mode 2 remains essentially unchanged from the original analysis. These revised modes, when used as boundary conditions to a wind-forced coastal-trapped wave model that predicts velocity and sea level along the coast, lead to a small but significant increase in prediction skill over the original modes. The reanalysis raises questions regarding the energy source for the coastal-trapped wave modes.

The difference between the original and present analyses is reduced by the inclusion of sea-level data. The ability of the instrument array to resolve coastal-trapped wave modes is discussed, and the problems associated with nonorthogonality of the theoretical modal structures as sampled by the array are highlighted. It is noted that the small number of degrees of freedom in the data leads to 95% confidence limits on modal energy fluxes that are as large as 69% of the estimated values.

1. Introduction

Much of our present knowledge concerning the velocity field on the southeastern Australian shelf has been derived from the Australian Coastal Experiment (ACE), carried out between September 1983 and March 1984 (Freeland et al. 1986). ACE was designed to enable the alongshore evolution of the coastal-trapped wave (CTW) signal to be identified, and to be compared with predictions from a numerical CTW model forced by local winds and remote CTWs. In the original time domain analysis (Church et al. 1986a,b), time series of three CTW modes were estimated from current meter data at three locations spanning about 500 km in the alongshore direction (see Fig. 1). A statistically based "eddy" mode was also fitted in an attempt to minimize contamination from the East Australian Current.

Results from ACE showed that northward propagating CTWs may account for a significant part of the subtidal-frequency alongshore current and coastal sea-level variability in the vicinity of Sydney. The southward flowing East Australian Current and its associated eddies were observed to have an episodic impact upon the velocity field in this region, with the most pronounced activity occurring on the outer shelf and slope region (Freeland et al. 1986; Church et al. 1986a,b; Huyer et al. 1988).

The CTW modes obtained at the southernmost line (line 1) were used as boundary conditions for a forced-CTW model, and the predicted modal amplitudes at the two northern lines compared with the observed data. On the basis of these comparisons Church et al. (1986b) deduced that about three-quarters of the CTW energy flux observed at line 2 was the result of free CTW propagation from line 1, with the remainder due to wind forcing in the intervening region.

Our aim in this study is to improve the resolution of the CTW modes in the ACE region through a number of changes to the fitting procedure used by Church et al. (1986a,b). The need for a reappraisal of the ACE modes was triggered by the work of Schahinger and Church (1994), in which the predictive capabilities of the forced-CTW model were extended beyond the ACE period to times when extensive moored arrays were not available for the definition of CTW modes. To achieve this, Schahinger and Church (1994) identified easily measurable proxies for the observed modes at line 1 such as coastal sea level in northeastern Bass

Corresponding author address: Dr. Peter C. McIntosh, CSIRO, Division of Oceanography, GPO Box 1538, Hobart 7001, Australia.

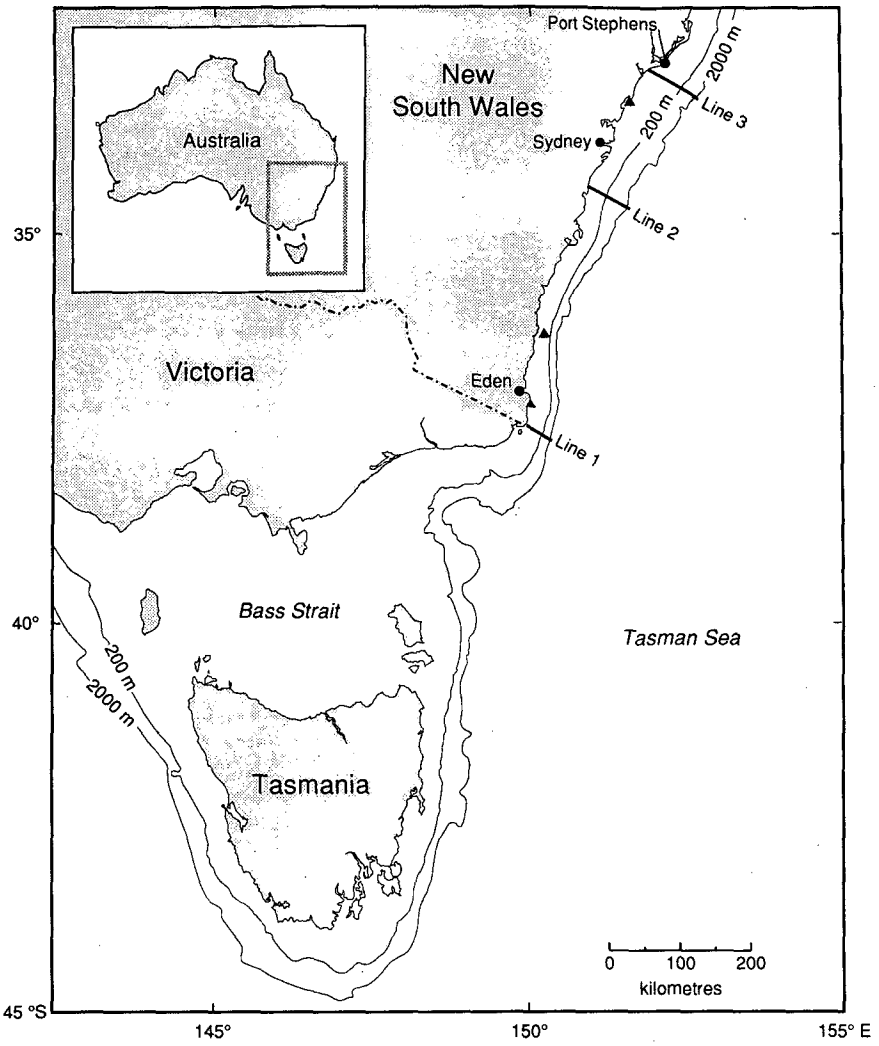

FIG. 1. Map of southeastern Australia showing the location of the three lines of ACE current meters. Triangles denote locations from which wind-stress measurements are used to drive the numerical forced-CTW model, and dots indicate the site of sea-level measurements.

Strait and wind stress within the strait. These proxy modes were then used as the upstream boundary conditions of the model. As the feasibility of this method relied almost wholly on previous estimates of modal amplitudes at line 1, a careful reexamination of the mode-fitting technique was necessary.

Our reanalysis involves three major changes. First we add all available sea-level data to the velocity data used in the original mode-fitting procedure. Second, we replace the truncated singular value decomposition technique, used by Church et al. (1986a,b) to stabilize what turned out to be an ill-conditioned problem, by a weighted least-squares technique in which marginally useful information is down weighted rather than discarded. One of the weighting factors is related to the signal-to-noise ratio of the dataset, which is determined from the data itself by optimizing the prediction of withheld data. Third, the relative importance of the fitted modes, which influences the solution of an ill-conditioned problem, is determined iteratively, so that the eventual solution is consistent with the prior specification of these modal weights. This latter feature is perhaps the most novel aspect of our method, and is possible only because we can estimate average weights from the sequence of solutions at individual times. We take some care to try and isolate the effect of incorporating additional data from the effect of our analysis method.

In comparing our revised modes with the original modes for the ACE period we use a number of diagnostics. These include the residual velocity and sea-level variance averaged over the entire time series as a percentage of the observed variance, the absolute and relative energies of the modes, and the degree of cor-

relation between modes. The worth of the new mode-fitting procedure can be gauged (in part) by the quality of the predictions from the forced-CTW model that was employed in the ACE studies. That is, if the revised modes at line 1 are used as the upstream boundary condition of the model, are the hindcast time series of sea level and current along the New South Wales coast better matched with the observations?

While the focus of the comparisons will be the modes obtained at line 1, the modes at line 2 will also be reexamined. We are principally interested in whether the alongshore phase speeds of the revised first and second CTW modes are any closer to the theoretical values than were the original modes. The line 3 modes will not be pursued in this study; Church et al. (1986a,b) found CTW theory to be less applicable at this line due to the dominance of the East Australian Current.

The structure of the paper is as follows: new features of our method are introduced in section 2, and the ability of the ACE array to resolve CTW modes is assessed in section 3. Results are presented in section 4 and compared with the previous analysis. Section 5 contains a discussion of our results, particularly in relation to the energy source of CTWs. The paper concludes with section 6.

2. Data and method

a. Sea-level data

The original coastal-trapped wave mode fitting by Church et al. (1986a) used only the alongshore component of velocity from current meters deployed over the continental margin. There were nine velocity time series at line 1, and 12 at line 2. We have augmented these data with sea-level data corrected for atmospheric pressure and otherwise processed in the same way as the velocity data (Forbes 1985). That is, the hourly time series were first low-pass filtered (half-power point at 40 h) and decimated to 12 h values. The twice-daily data were then high-pass filtered to remove signals with periods longer than about 20 days.

At line 1, the available data consisted of adjusted sea level at the coast, and bottom pressure measured at the 135-m and 500-m isobaths. Line 2 (and line 3) had only nearby coastal data. For ease of notation, we shall refer to both adjusted coastal sea level and bottom pressure as simply sea level.

The sea-level time series span the period from 21 September 1983 to 20 March 1984 (363 12 h records), though the first 24 days of the 500-m sea-level time series at line 1 were discarded due to an apparent drift. The duration of this drift was too close to the low-frequency end of the high-pass filter to be completely removed.

b. Stable mode fitting

The original mode-fitting procedure involves expressing the data as a sum of three coastal-trapped wave modes plus an empirically determined "eddy" mode. At each alongshore location and at each time, the ith current measurement d_i is represented as a linear combination of the velocity modal structures $G_j(x_i, z_i)$, where x_i and z_i are the offshore location and depth, respectively, of the ith current meter. The CTW velocity modal structures are obtained by differentiating the pressure modal structures, $F_j(x_i, z_i)$, which in turn are obtained as eigenvectors of the low-frequency, low-wavenumber CTW eigenvalue problem (Church et al. 1986a). The eddy mode is determined by correlating each time series with one deep offshore current record where it is assumed that the signal is entirely due to the East Australian Current (Freeland et al. 1986; Church et al. 1986a). The amplitude of the jth mode is denoted by ϕ_j, and is determined by solving the following equations:

$$\sum_{j=1}^{4} G(x_i, z_i)\phi_j = d_i, \quad i = 1, \cdots, n, \qquad (1)$$

where there are n velocity measurements at that time and alongshore position. If p sea-level measurements, h_k, are available, then an additional set of equations is obtained:

$$\sum_{j=1}^{4} F(x_k, z_k)\phi_j = h_k, \quad k = 1, \cdots, p. \qquad (2)$$

In matrix notation, Eqs. (1) and (2) are combined and written

$$\mathbf{G\Phi} = \mathbf{d}. \qquad (3)$$

The matrix $\mathbf{G}$ has more rows than columns because there are more measurements than fitted modes at each line and at each time. Hence the matrix system (3) cannot, in general, be solved exactly, but rather a least-squares solution is sought. In addition, the system was determined by Church et al. (1986a) to be rank deficient, requiring the use of a technique such as the singular value decomposition (SVD) to obtain a stable solution.

The SVD solution is given by (e.g., Wunsch 1978)

$$\mathbf{\Phi} = \sum_{i=1}^{R} \frac{\mathbf{u}_i^{\mathrm{T}}\mathbf{d}}{s_i}\mathbf{v}_i, \qquad (4)$$

where $\mathbf{u}_i$ and $\mathbf{v}_i$ are the singular vectors of $\mathbf{G}$ (eigenvectors of $\mathbf{GG}^{\mathrm{T}}$ and $\mathbf{G}^{\mathrm{T}}\mathbf{G}$, respectively). The associated singular values, s_i, are obtained as the square root of the eigenvalues of the smaller of the two matrices $\mathbf{GG}^{\mathrm{T}}$ and $\mathbf{G}^{\mathrm{T}}\mathbf{G}$, and are ordered from largest ($i = 1$) to smallest ($i = 4$). The sum is taken up to the rank, R, of the system, which Church et al. (1986a) determined to be 3. Hence information contained in the fourth singular vector is discarded. It will be shown in a later section that this amounts to discarding information regarding the separation of mode 1 from mode 2. However, trun-

cating the SVD sum has the advantage of reducing the impact on the solution of small singular values in the denominator of (4). In particular, Church et al. (1986a) found that truncating the SVD considerably reduced the solution squared amplitude $\mathbf{\Phi}^T\mathbf{\Phi}$.

An alternative method for stabilizing the solution to a rank deficient system is described by, among others, Wunsch (1978), and is called variously "ridge analysis," "tapered least squares," or "smoothed least squares." The smoothed least-squares solution is obtained by minimizing a penalty function consisting of a linear combination of the equation residual variance and the solution "energy" (squared amplitude)

$$\mathbf{e}^T\mathbf{e} + \sigma^2\mathbf{\Phi}^T\mathbf{\Phi}, \tag{5}$$

where $\mathbf{e} = \mathbf{G\Phi} - \mathbf{d}$ is the vector of equation residuals. The smoothing parameter σ determines the trade-off between having a large solution with a small residual variance (small σ), or a less energetic solution with increased residual variance (large σ).

The solution that minimizes (5) is

$$\mathbf{\Phi} = \sum_{i=1}^{m} \frac{\mathbf{u}_i^T\mathbf{d}}{\mathbf{s}_i + \sigma^2/\mathbf{s}_i}\mathbf{v}_i, \tag{6}$$

where the sum is now over the total number of fitted modes, $m = 4$. The impact of small singular values on the solution can be reduced by a suitable choice of the smoothing parameter σ. The advantage of this approach is that information contained in all the singular vectors is retained, although it may be downweighted if σ is comparable to or larger than s_i.

c. *Weighting*

The importance of row and column weighting of $\mathbf{G}$ has been discussed by many authors (e.g., Wunsch 1978; Menke 1984). In its simplest form, row weighting amounts to ensuring that no one equation dominates the residual term $\mathbf{e}^T\mathbf{e}$ by virtue of the units in which the equation is expressed. Here, we will be fitting modes to both velocity data and sea-level data, so that some equations will be in meters per second, while others will be in meters. To compensate, the first term in the penalty function (5) is replaced by a weighted version, $\mathbf{e}^T\mathbf{W}_e^{-2}\mathbf{e}$, where the diagonal matrix $\mathbf{W}_e$ contains the inverse of the row weights. The choice of these weights is discussed in the next section.

Column weighting amounts to rescaling the contribution of the modal amplitudes ϕ_i to the penalty function. In the absence of column scaling, if one mode has considerably less energy than the others, relatively large variations in this mode due to data or model errors will not be sufficiently penalized. Hence, we define a diagonal matrix of square roots of estimated modal energies, $\mathbf{W}_\phi$. We assume that all weight matrices are independent of time; there is insufficient information to do otherwise.

The modified penalty function becomes

$$\mathbf{e}^T\mathbf{W}_e^{-2}\mathbf{e} + \sigma^2\mathbf{\Phi}^T\mathbf{W}_\phi^{-2}\mathbf{\Phi}. \tag{7}$$

The solution to (3) that minimizes this modified penalty function is

$$\mathbf{\Phi} = [(\mathbf{G}^T\mathbf{W}_e^{-2}\mathbf{G} + \sigma^2\mathbf{W}_\phi^{-2})^{-1}\mathbf{G}^T\mathbf{W}_e^{-2}]\mathbf{d} \equiv \mathbf{G}^{-g}\mathbf{d}, \tag{8}$$

which reduces to the smoothed least-squares solution (6) when the weight matrices are set to the identity matrix and the singular value decomposition of $\mathbf{G}$ is made. The notation $\mathbf{G}^{-g}$ represents the generalized inverse of $\mathbf{G}$.

An alternative way of expressing the weighted problem and its solution in terms of the nontruncated SVD will be useful later. By defining scaled variables, we can write the scaled equations, penalty function, and solution in forms directly analogous to the unscaled versions (3), (5), and (6), respectively; that is,

$$\hat{\mathbf{G}}\hat{\mathbf{\Phi}} = \hat{\mathbf{d}} \tag{9}$$

$$\hat{\mathbf{e}}^T\hat{\mathbf{e}} + \sigma^2\hat{\mathbf{\Phi}}^T\hat{\mathbf{\Phi}} \tag{10}$$

$$\hat{\mathbf{\Phi}} = \sum_{i=1}^{m} \frac{\hat{\mathbf{u}}_i^T\hat{\mathbf{d}}}{\hat{\mathbf{s}}_i + \sigma^2/\hat{\mathbf{s}}_i}\hat{v}_i, \tag{11}$$

where $\hat{\mathbf{G}} = \mathbf{W}_e^{-1}\mathbf{G}\mathbf{W}_\phi$; $\hat{\mathbf{\Phi}} = \mathbf{W}_\phi^{-1}\mathbf{\Phi}$; $\hat{d} = \mathbf{W}_e^{-1}\mathbf{d}$; $\hat{e} = \mathbf{W}_e^{-1}e$; and $\hat{\mathbf{u}}_i$, $\hat{\mathbf{v}}_i$, and $\hat{\mathbf{s}}_i$ are the singular vectors and singular values of $\hat{\mathbf{G}}$.

d. *Estimating parameters*

1) SMOOTHING PARAMETER

There are a number of methods available for choosing the smoothing parameter σ (e.g., Lawson and Hanson 1974; Wahba 1990, chapter 4). Some are based on reducing the solution sensitivity to errors in the data and model to an acceptable level, while others rely on prior knowledge of the expected signal-to-noise ratio (σ can be interpreted as the reciprocal of this ratio). We have chosen a technique called generalized cross validation (GCV) (Wahba and Wendelberger 1980; Wahba 1990; McIntosh and Veronis 1993), which seems to require the least additional information or assumptions. The principle behind GCV is to use the data itself to validate the choice of σ by optimizing the prediction of withheld data.

To be more precise, a value of σ is chosen, and the solution (8) is obtained after leaving one equation out. The right-hand side of the withheld equation (either a velocity or sea-level measurement) is then predicted from this solution, and the discrepancy stored. This procedure is repeated, leaving out each equation in turn, to obtain an overall measure of the predictive skill of the model with the particular value of σ chosen. The optimum σ is found by maximizing this predictive skill using a standard optimization technique. In practice, the computational tedium of solving many $n - 1$

problems has been circumvented by some clever linear algebra (see, e.g., Golub et al. 1979), and the optimum σ may be found by minimizing the expression

$$V(\sigma) = \frac{n^{-1}\mathbf{e}^T\mathbf{W}_e^{-2}\mathbf{e}}{[n^{-1}\text{trace}(\mathbf{I} - \mathbf{G}\mathbf{G}^{-g})]^2}. \quad (12)$$

Our initial idea was to find the optimum σ at each time for which data were available, that is, every 12 h. However, there are only about 10 measurements at each time, and GCV seems to need a larger number (generally upwards of 50) to work properly (Wahba 1990, p. 65). We have experimented with combining a number of individual matrix problems of the form (3) together in a large block matrix formulation and estimating σ for this larger problem. The values of σ obtained stabilized when more than about 15 individual matrix problems were combined. We have decided to combine 21 individual times together (i.e., 10 days of data) in order to estimate a value of σ applicable to the central time. This process is repeated for each 12 hourly time.

As a check on GCV, we have also performed Monte Carlo experiments to determine the variation in singular values of the scaled matrix $\hat{\mathbf{G}}$ due to errors in the computed modal structures (matrix elements). We assume errors in modes 1 to 3 of 10%, 15%, and 20%, respectively, and errors of 50% in the eddy mode. Normally distributed errors with these standard deviations are added to the matrix, and the standard deviations of the singular values determined over 200 realizations. With this level of error, it is surprising that the smallest singular values have a standard deviation of less than 15%. A lower bound on σ, denoted by $\sigma_{\min}$, is chosen so that the smallest singular value in the denominator of the SVD (11) is increased by at least twice its standard deviation. This is done by defining $\sigma_{\min}$ by

$$\frac{\sigma_{\min}^2}{\hat{\mathbf{s}}_m} = 2\text{sd}(\hat{\mathbf{s}}_m), \quad (13)$$

where sd stands for standard deviation. Hence we are 95% confident that the smallest singular value of $\hat{\mathbf{G}}$ is no smaller than the true (unknown) singular value for the analysis. In practice, $\sigma_{\min}$ is always smaller than the smallest σ value obtained using GCV, which indicates that the GCV procedure is stabilizing the analysis adequately.

2) ROW WEIGHTS

If the statistical view of solving least-squares problems is taken (e.g., Menke 1984), then the weighted penalty function (7) can be thought of as a χ^2 variable provided that the weight matrices $\mathbf{W}_e^2$ and $\mathbf{W}_\phi^2$ are the covariance matrices for the equation residual and solution, respectively, and provided also that the equation residual and the solution have zero mean. Minimizing (7) then amounts to obtaining the most likely solution if Gaussian statistics are assumed. In practice, full covariance information is rarely available, and the weight matrices are often assumed to be diagonal. We have chosen a diagonal equation residual weight matrix with elements proportional to the standard deviation of the entire velocity or pressure dataset at the appropriate line. In other words, we divide each equation involving velocity data by a single estimate of the true velocity standard deviation, and each equation involving sea level by an estimate of the sea-level standard deviation. A more sophisticated weighting would be to divide each equation by an estimate of the expected error in that equation. This estimated error would probably be dominated by eddy noise not accounted for by the eddy mode but would also contain contributions arising from approximating the CTW mode structures and from instrument noise.

One way of deciding whether the sea-level and velocity equations are appropriately scaled is to look at the scaled data resolution matrix (Menke 1984)

$$\mathbf{N} = \hat{\mathbf{G}}\hat{\mathbf{G}}^{-g}. \quad (14)$$

The diagonal elements of this matrix indicate the extent to which individual equations are resolved and are contributing to the solution. Experiments with different row scales for the velocity and sea-level equations show that the diagonal of $\mathbf{N}$ is more nearly uniform for our choice of scales than if either of the velocity or sea-level equation scales are varied by a factor of 2.

We have investigated the sensitivity of our solutions to the choice of equation scaling. In particular, the velocity and sea-level standard deviations used to scale the equations were independently increased and decreased by a factor of 2. The three major findings of this study were not altered. Preempting the results somewhat, for all choices of row weights, mode 1 contained more energy than mode 2, mode 3 contained negligible energy, and modes 1 and 2 were not correlated.

3) COLUMN WEIGHTS

In the absence of full covariance information, the squared column weight matrix $\mathbf{W}_\phi^2$ is assumed to be diagonal. It is desirable that the diagonal elements consist of the solution variance (time-averaged modal energy fluxes), so that a relatively high energy mode does not dominate the penalty function. In the case of the eddy mode, squared amplitude would be used, although it may not be interpreted as an energy flux. One of the reasons for conducting a mode-fitting exercise such as this is to estimate the time-averaged energy flux in each fitted mode. In the absence of any other information, the original analysis by Church et al. (1986a) could only assume a priori equal energy in each mode. However, a further analysis could now make use of this preliminary estimate of energy fluxes to choose more appropriate column weights.

Our iterative weighting procedure starts with uniform column weights, although the final weighting does not depend on this starting point. A first estimate of the modal amplitudes is obtained from (8) using GCV to choose the smoothing parameter. Modal energy fluxes for this solution are calculated, the column weights are redefined and a new solution obtained. This procedure is iterated until the energy flux changes are insignificant (less than 0.25%), which typically takes seven iterations. Note that the column weights are not a function of time; they change only between iterations and between lines. We have not combined column weighting with the SVD technique although this can be done.

We have experimented with using GCV to determine the column weights as well as the smoothing parameter. The results obtained were quite similar to those obtained using the iterative scheme, but we found that at some times the solution amplitudes dropped to almost zero. This was caused by GCV choosing overly large column weights, which we attribute to the multidimensional minimization algorithm finding a local rather than a global minimum of the GCV function (12). This did not seem to be a problem when using GCV to estimate only the smoothing parameter, in which case the minimization is one-dimensional.

e. *Solution error estimates*

Apart from calculating the modal amplitudes, it is also possible to calculate an estimate of the standard error in these amplitudes. The full covariance matrix of modal amplitudes is given by (Lawson and Hanson 1974; Menke 1984; Tarantola 1987)

$$\mathrm{cov}(\mathbf{\Phi}) = \frac{\mathbf{e}^{\mathrm{T}}\mathbf{W}_e^{-2}e}{n-m}(\mathbf{G}^{\mathrm{T}}\mathbf{W}_e^{-2}\mathbf{G} + \sigma^2\mathbf{W}_\phi^{-2})^{-1}. \quad (15)$$

It may be shown that this expression, and indeed the solution (8), is independent of a uniform scaling of either weight matrix. In other words, one can double the size of all elements of $\mathbf{W}_\phi$, for example, and σ will double also to compensate. Alternatively, one may double all elements in $\mathbf{W}_e$, and σ will decrease by a factor of 2, leaving the solution and its covariance unchanged. Hence it is only the relative weighting expressed by differences between the elements within either weight matrix that affects the solution and its error estimate. It is the GCV procedure that determines the appropriate trade-off between errors in the data and the solution size, in other words, the signal-to-noise ratio.

We use the solution covariance matrix in a statistical test of whether a modal amplitude is significantly different from zero relative to its standard error. The standard error is just the square root of the appropriate diagonal element of the covariance matrix.

3. Mode orthogonality, separation, and array design

It is desirable that the columns of the matrix **G** are orthogonal, so that the solution for one mode is unaffected by either leaving another mode out of the calculation, or by changing the weighting assigned to another mode. The theoretical CTW cross-shelf structures form an orthogonal set (Wang and Mooers 1976; Clarke 1977), but the orthogonality relation is not directly related to the structure of **G**. Hence it may not be possible to design an instrument array to ensure that the columns of **G** are orthogonal. Even if this could be achieved for the CTW modes, the eddy mode does not obey any theoretical orthogonality relation and so the corresponding column of **G** may be dependent on the columns associated with CTW modes.

In the various experiments we conducted, it was clear that leaving one mode out of the fitting procedure changed the amount of energy in the remaining modes. This led us to examine the normalized inner product of the columns of **G** at line 1, both with and without the sea-level data (see Table 1). The diagonal column weight matrix $\mathbf{W}_\phi$ will not change these inner products, and so the analysis applies to both the original SVD solution and the present GCV solution. Columns that are orthogonal will have an inner product of 0, while identical columns will have an inner product of 1. The most striking feature of this table is that for either analysis, mode 1 is strongly coupled with both mode 2 and the eddy mode. Hence there will be some ambiguity in apportioning energy between these modes. However, leaving the eddy mode out, for example, does not mean that exactly 64% of its energy will flow into mode 1. The redistribution of energy is more complex because all the modes are coupled to various degrees.

A more rigorous way to examine the separation of modes due to a particular instrument array is to use the singular value decomposition. The SVD of a matrix gives far more information than is necessary to solve the matrix system. In particular, the $\mathbf{v}_i$ singular vectors specify the linear combination of modes associated with each singular value. If a singular value is discarded by truncating the SVD sum (4), then the associated linear combination of modes is not available to the solution.

The singular vectors and ratio of singular values to the largest for the original truncated SVD solution ob-

TABLE 1. Normalized inner products of columns of **G** corresponding to various CTW modes at line 1 using all data (velocity and sea level) and velocity data alone.

Inner product	All data	v data only
1 and 2	0.80	0.80
1 and 3	0.09	0.07
1 and eddy	0.64	0.66
2 and 3	0.52	0.51
2 and eddy	0.22	0.23
3 and eddy	0.30	0.30

TABLE 2. The v singular vectors for the original (SVD) analysis at line 1. The vectors presented are for the longest period with a fixed number of instruments. Numerical values only change slightly with differing numbers of instruments. The singular values, s_i are normalized by the largest singular value, s_1.

Mode	v_1	v_2	v_3	v_4
1	0.51	0.01	0.28	0.81
2	0.58	0.45	0.45	−0.52
3	0.11	0.72	−0.67	0.16
eddy	−0.63	0.53	0.52	0.21
s_i/s_1	1.00	0.89	0.42	0.14

tained by Church et al. (1986a) are shown in Table 2. The most striking feature of these singular vectors is that information about the separation between mode 1 and mode 2 is contained mainly in v_4, whereas v_1 and v_3 have essentially no ability to separate modes 1 and 2. Here v_2 has some ability to resolve mode 2, although it mainly resolves mode 3. Truncating the smallest singular value, which Church et al. (1986a) found necessary to stabilize the solution, removes v_4 from the SVD solution, leaving the mode 1 and 2 amplitude time series highly correlated ($r = 0.94$). This correlation is not physically based, but is related to the inability of the instrument array to resolve adequately modes 1 and 2. This conclusion is not changed substantially by the addition of sea-level data.

It should be noted that, generally, no data are necessary to perform this type of analysis, only the instrument locations and theoretical modal structures that go to make up the matrix **G**. However, in this application of the SVD, the eddy mode structure is determined from the data (Church et al. 1986a). If the eddy mode is left out of the analysis, the singular vector structure is similar (see Table 3). Again, the singular vector that contains the most information concerning the separation of mode 1 and mode 2 is associated with the smallest singular value.

It is possible that the loss of instruments from the original designed array has contributed to the difficulty in separating modes 1 and 2. To test this, the SVD of the intended array at line 1 (see Church et al. 1986a for the configuration) was computed (see Table 4). The singular values and vectors have not changed much from those of the recovered array, and we conclude that the loss of instruments did not degrade significantly the ability of the ACE array to resolve coastal-trapped wave modes.

TABLE 3. As for Table 2 but with the eddy mode left out.

Mode	v_1	v_2	v_3
1	0.41	0.55	0.73
2	0.74	0.27	−0.62
3	0.53	−0.79	0.30
s_i/s_1	1.00	0.67	0.20

TABLE 4. As for Table 2 but for the intended ACE array and with the eddy mode left out and no sea-level data. Results are essentially unchanged by the addition of sea-level data.

Mode	v_1	v_2	v_3
1	0.46	0.53	0.71
2	0.78	0.15	−0.61
3	0.43	−0.83	0.35
s_i/s_1	1.00	0.68	0.21

The singular vectors and singular value ratios for our revised method applied to just the velocity data are shown in Table 5. These are determined from the scaled matrix $\hat{\mathbf{G}}$. The main feature is that mode 3 is almost wholly resolved by the fourth singular vector, which is associated with a small singular value. This is because the iterative column scaling technique has determined that mode 3 contains very little energy, which makes it difficult to resolve relative to the more energetic modes. The next feature to note is that the separation between modes 1 and 2 is now mainly determined by v_3, which has a larger singular value (relative to the largest) than the equivalent singular vector from the original SVD analysis (0.19 versus 0.14). Hence our revised method should be able to discriminate more reliably between modes 1 and 2.

Unfortunately, the scaled matrix $\hat{\mathbf{G}}$ depends on the data through the iterative column scaling technique, which in turn depends on GCV to determine the signal-to-noise ratio. Hence, unless good prior estimates of the expected solution size and signal-to-noise ratio were available, it would not be possible to perform this analysis prior to deploying the array. However, there are some features of the original SVD analysis that have not changed much. For example, separation of modes 1 and 2 is clearly going to present some difficulty regardless of scaling. Ideally, one would like to design the instrument array so that the most energetic modes (presumably the first two) are resolved first. If the instrument array could be arranged so that each singular vector resolved just one mode, then column scaling will not alter this structure but merely alter the size of the associated singular value.

The point of this argument is that it is still worth looking at the SVD analysis of a potential instrument

TABLE 5. The $\hat{v}$ singular vectors and singular value ratios of $\hat{\mathbf{G}}$ for our revised analysis of line 1, using only velocity data. Similar results are obtained with the addition of sea-level data.

Mode	$\hat{v}_1$ (16)	$\hat{v}_2$ (17)	$\hat{v}_3$ (18)	$\hat{v}_4$ (19)
1	0.73	0.05	0.66	0.18
2	0.51	0.62	−0.54	−0.26
3	0.01	0.14	−0.29	0.95
eddy	−0.45	0.78	0.44	−0.03
s_i/s_1	1.00	0.56	0.19	0.08

array before deployment, making the best guess possible about column (and row) scales. Techniques for improving the array are beyond the scope of this article; recent work on array design by Barth and Wunsch (1990) gives some indication as to how one might proceed.

4. Comparison of revised and original analyses

The original ACE coastal-trapped wave analysis (Church et al. 1986a,b) has been modified in two ways. The CTW mode-fitting technique has been modified by adopting a weighted least-squares approach, and the original velocity data have been augmented by sea-level data. The initial goal of this analysis is to determine the modal amplitude time series, which depend on both the data used and on a rational choice of weights for the least-squares method. In this section we endeavor to discriminate between changes caused to the original modal decomposition by the use of additional data, and changes due to the analysis method.

We will refer to the original singular value decomposition method by the acronym SVD, while our weighted least-squares method is referred to as WLS. The latter method uses GCV to determine the signal-to-noise ratio, while the column scales are determined by an iterative method. Although it is possible to apply GCV without column scaling and also possible to apply iterative column scaling to the SVD method, these cases would complicate the discussion unnecessarily. Hence we concentrate on the effect on the solution of additional data versus the weighted least-squares analysis method.

In comparing the effect of additional data and the least-squares method on the modal decomposition, the following diagnostics will be used:

1) total CTW mode energy flux, and its partitioning between the fitted modes
2) correlation between the time series of mode 1 and mode 2 amplitudes
3) percentage of total variance of the data not explained by the fitted modes
4) coherence of the CTW mode amplitude time series between lines 1 and 2
5) ability of a numerical wind-forced CTW model (Church et al. 1986a,b) to predict both alongshore currents at line 2 and sea-level data along the coast when the analyzed modes at line 1 are used as boundary conditions to the model.

The last diagnostic is particularly important because the increased ability to predict sea-level and currents along the coast is the major practical benefit of an experiment such as ACE.

a. Effect of mode-fitting method

Consider first the difference between the CTW mode energy fluxes obtained using the SVD and WLS mode-fitting methods. Each method has been used to analyze both the original dataset (consisting of velocity measurements alone), and a larger dataset created by the addition of sea-level data. The energy fluxes for the two methods applied to both datasets are given in Table 6 for line 1 and Table 7 for line 2. In all cases the WLS method acts to increase the energy in mode 1 by at least a factor of 2 over the SVD method. In contrast, the energy in mode 2 remains relatively unchanged. The net effect is to change the conclusion of the original ACE study that mode 1 contained less energy than mode 2. The results of the WLS analysis suggest that mode 1 contains about 70% of the total CTW energy at line 1, and about 60% of the total CTW energy at line 2.

The original SVD analysis of velocity data alone found that mode 3 contained 20% of the total CTW energy at line 1, and 31% at line 2. However, the WLS analysis finds that there is negligible energy in mode 3. We attribute this result to the use of column scaling, that is, weighting the modes in the least-squares procedure by the reciprocal square root of the expected energy in the modes.

Overall, the WLS method results in an increase in the total CTW energy relative to the SVD method of 60% at line 1 for the analysis of velocity data alone, or 21% if sea-level data are included. At line 2, the increases are 28% and 35%, respectively, for the two datasets.

When considering the eddy mode, we note that the squared amplitude of this mode is proportional to the energy in this mode but may not be interpreted directly as energy, as may be done with the CTW modes. The squared size of the eddy mode is essentially independent of the analysis method used. This result is probably due to the fact that the eddy mode is well resolved by the array. However, the fact that the energy in mode 2 also is essentially independent of the analysis method

TABLE 6. Line 1 time-averaged energy fluxes (in units of 10^8 W), velocity (v) and sea-level (η) residuals as a percentage of total variance in the observed data, and correlation (r) between modes 1 and 2 for different models. SVD refers to the original truncated singular value decomposition analysis of Church et al. (1986a,b), and WLS refers to our weighted least-squares method. Subscript values are percentage of total CTW energy in a particular mode. The 95% significance level for correlations is about 0.30. Analysis period is from 21 September 1983 to 20 March 1984.

Mode	SVD v data	WLS v data	SVD $v + \eta$ data	WLS $v + \eta$ data
1	0.58_{31}	2.18_{71}	0.79_{40}	1.66_{69}
2	0.94_{49}	0.85_{28}	0.86_{43}	0.71_{30}
3	0.39_{20}	0.03_{1}	0.34_{17}	0.04_{2}
total CTW	1.91	3.05	2.00	2.41
eddy	0.78	0.68	0.89	0.85
v residual (%)	13	13	16	17
η residual (%)	49	59	25	20
r mode 1 vs mode 2	0.94	0.04	0.95	0.28

TABLE 7. Same as Table 6 but for line 2. Analysis period is from 10 September 1983 to 8 March 1984.

Mode	SVD v data	WLS v data	SVD $v + \eta$ data	WLS $v + \eta$ data
1	0.65_{28}	1.90_{63}	0.71_{32}	1.64_{56}
2	0.98_{42}	1.05_{35}	0.71_{33}	1.28_{43}
3	0.72_{31}	0.05_{2}	0.77_{35}	0.03_{1}
total CTW	2.35	3.00	2.19	2.95
eddy	1.61	1.66	1.81	2.17
v residual (%)	19	18	21	20
η residual (%)	101	96	55	32
r mode 1 vs mode 2	0.81	−0.23	0.82	−0.37

cannot be attributed to the resolution of the array; it was noted in section 3 that the array was not well suited to distinguishing mode 1 from mode 2.

Another point made in section 3 was that the SVD analysis virtually guaranteed that the mode 1 and mode 2 amplitude time series would be highly correlated. This is because the information that would allow the separation of these modes is discarded due to the necessity of stabilizing the SVD solution by truncating the SVD. Hence the analyzed amplitude time series of mode 1 and mode 2 will be highly correlated whether these modes are actually correlated or not. The original analysis found that the correlation between mode 1 and mode 2 was 0.94 at line 1 and 0.81 at line 2. One of the advantages of the weighted least-squares method is that no information is discarded, although it may be down-weighted. Consequently, we find that for the WLS method the correlation between mode 1 and mode 2 is considerably reduced relative to the SVD values. At line 1 the correlation coefficient is 0.04 if velocity data alone are used and 0.28 if velocity and sea-level data are used. At line 2, the correlation is small and negative, being −0.23 for velocity data, and −0.37 for velocity and sea-level data. The significance level for all correlations is about 0.30 so that the WLS method suggests that mode 1 and mode 2 are not significantly correlated. We have tested the WLS method on an artificial dataset in which mode 1 and mode 2 were highly correlated and demonstrated that the WLS method is capable of producing a high correlation between mode 1 and mode 2.

We find that the SVD and WLS techniques give solutions that explain a similar fraction of the variance in the data. At line 1, the residual velocity variance is between 13% and 17% of the total velocity variance, depending on the dataset used. At line 2 the unexplained velocity variance is slightly higher at between 18% and 21% of the total. The residual sea-level variance depends on whether sea-level data are included explicitly in the fit or not. The WLS method generally leads to a slightly better fit, except at line 1 in the case where velocity data only are used (see Tables 6 and 7).

Looking at the residual variance at individual current meters reveals a similar pattern for both methods. The unexplained variance is roughly uniform in absolute terms, which is not surprising considering that each instrument has been given the same weight by choosing uniform row (equation) weights. Recall that the ACE current meter arrays extended from the 135-m isobath to the 2000-m isobath. The signal variance is generally higher for instruments on the shelf than off the shelf, and hence the unexplained variance as a percentage of the individual instrument variance increases for instruments off the shelf. It is possible to encourage uniform fractional residual variances at all instruments by choosing row weights proportional to individual instrument variances. However, such a choice does not allow for the fact that the dominant error at most instruments off the shelf will be due to eddy noise not accounted for by the time-independent eddy mode. Hence the choice of uniform row weights is a realistic first estimate, and the resultant uniform residual variances at all instruments is a reflection of the expected error in fitting the data.

We now examine the coherence between the mode amplitudes analyzed independently at line 1 and line 2. Of particular interest is the phase of the coherence between similar modes at the two lines. The phase is inversely proportional to the phase speed at a given frequency. The phase speed can be compared with the theoretical estimate of the phase speed; a favorable comparison is evidence of the existence of coherent coastal-trapped waves in the region. At this stage, we consider only the original velocity dataset. The effect of including sea-level data is considered in the following subsection.

Results for the original SVD analysis of velocity data alone (Church et al. 1986a) indicated that both mode 1 and mode 2 were highly coherent between line 1 and line 2 (Fig. 2a). The mode 1 phase speed was slightly lower than the theoretical value of 3.5 m s^{-1} at all frequencies for which there was both significant energy and significant coherence. For mode 2, the phase speed was close to the theoretical value of 1.9 m s^{-1}, but tended to increase with frequency.

The corresponding results for our WLS analysis of velocity data alone reveal a marked decrease in the alongshore coherence relative to the SVD results, especially in the case of mode 1 (Fig. 2b). The mode 1 phase speed is further from the theoretical value than was the SVD analysis, and the coherence levels are below the 95% significance level at the energetic frequencies. However, the phase speed for mode 2 is close to the theoretical value, and is now almost constant over the lower frequencies where the majority of the modal energy is contained.

In agreement with Church et al. (1986a), neither the SVD nor WLS analysis shows any alongshore coherence for either mode 3 or the eddy mode.

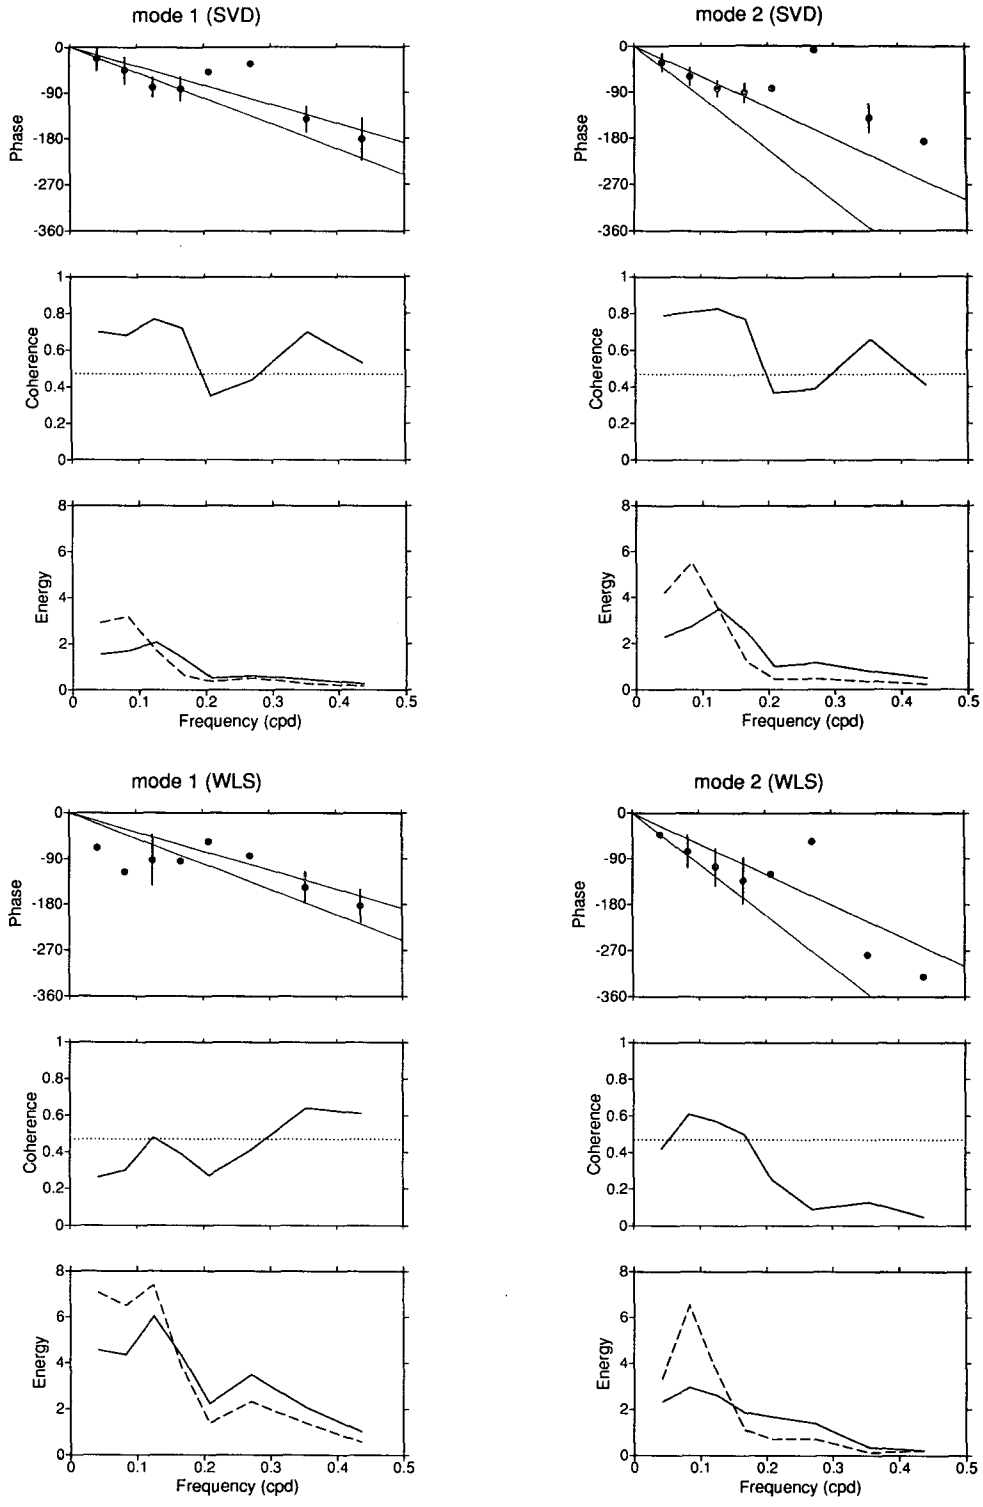

FIG. 2. Coherence, phase, and spectral energy of mode 1 and mode 2 amplitudes between line 1 and line 2; analysis period 336 12 h values from 21 September 1983. (a) Original SVD analysis; (b) WLS method using velocity data alone; (c) SVD method using velocity and sea-level data; (d) WLS method using velocity and sea-level data. Error bars indicate the 95% confidence limits for phase only at frequencies for which there is significant coherence. The solid lines on the phase plots indicate phase speeds of 3.0 and 4.0 m s^{-1} for mode 1 and 1.5 and 2.5 m s^{-1} for mode 2. The theoretical phase speeds for mode 1 (3.5 m s^{-1}) and mode 2 (1.9 m s^{-1}) lie approximately halfway between these lines. A negative phase implies that line 1 leads line 2. The 95% significance level for coherence is indicated by a dotted line. The solid and dashed energy curves correspond to line 1 and line 2, respectively; the energy units are arbitrary.

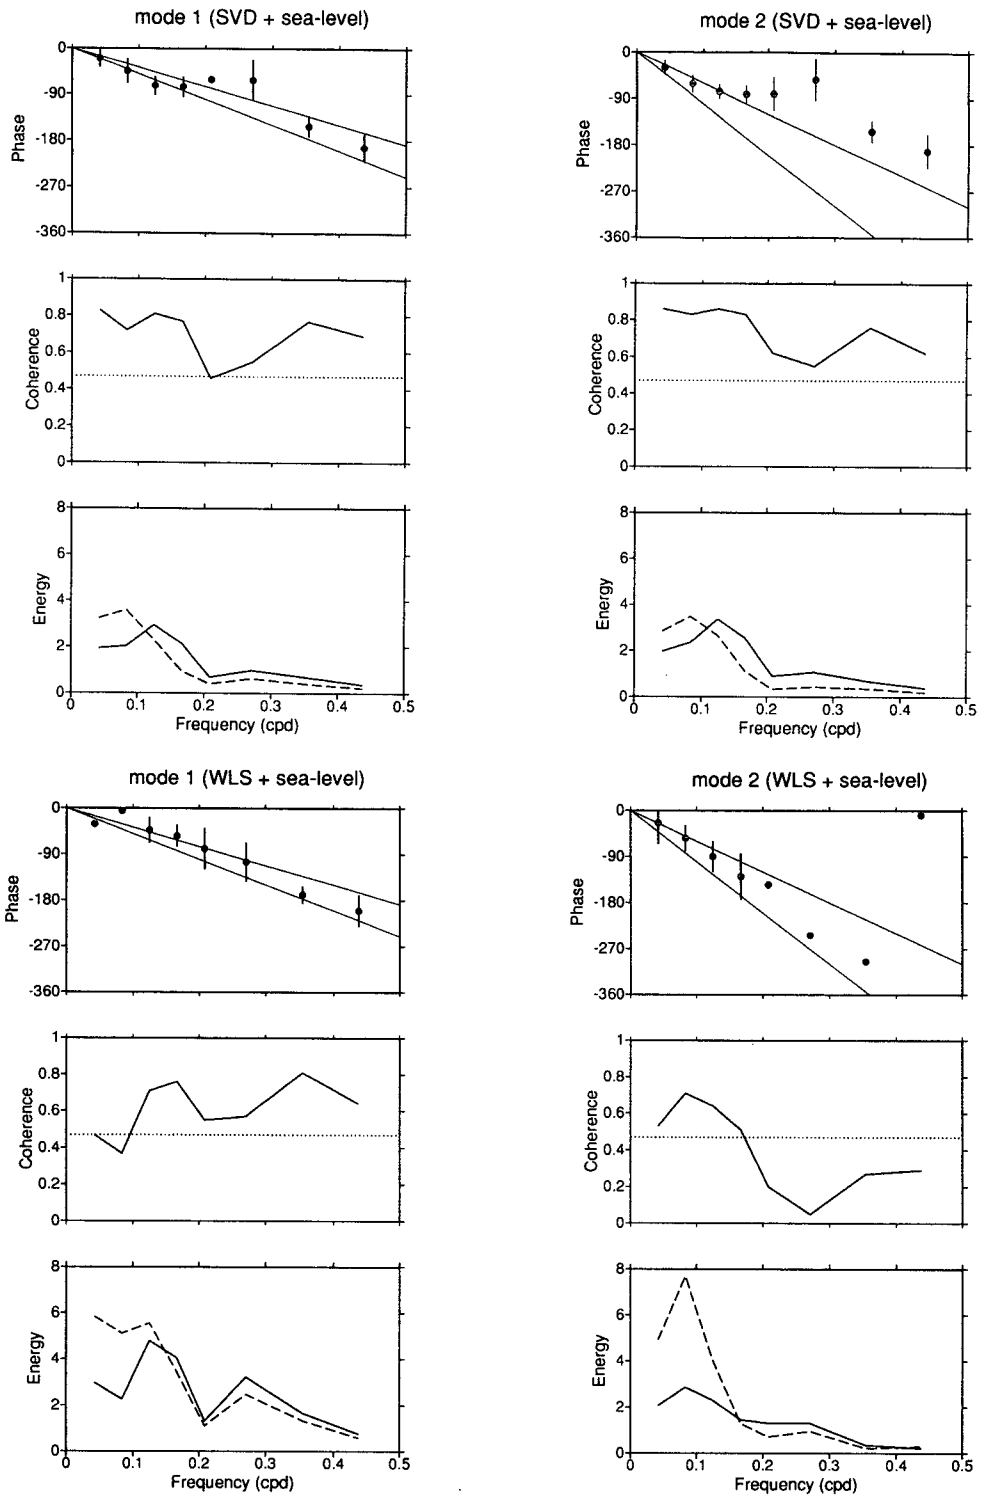

FIG. 2. (*Continued*)

Our final comparison involves the use of the analyzed CTW modes at line 1 as upstream boundary conditions for a wind-forced prognostic model (Church et al. 1986a,b). Observations at line 2 are compared with predictions from the model. The model propagates the input CTW modes along the coast, and incorporates

enhancement of the modes from observed winds, decay due to bottom friction, and exchange of energy between the modes. The eddy mode cannot be included in this procedure; to minimize the effect of eddies on our comparison, we consider only a 90-day subset of the ACE period that was relatively free of the effects of the East Australian Current. This period is the same as that used by Church et al. (1986b). Again, at this stage, we consider only the modes obtained from analyzing the original velocity dataset.

When the three CTW modes obtained at line 1 using the SVD technique are used as boundary conditions for the prognostic model, the observed and predicted alongshore currents at line 2 are most highly correlated at the nearshore sites. All these correlations are significant at the 95% level, as are the predictions of coastal sea level at line 2 and line 3. At each of the two deep-water offshore moorings (depths of 1200 m and 2000 m), only the prediction at the instrument closest to the surface had a correlation with the data that was significant at the 95% level (see Table 8). Similar results are obtained using the CTW modes obtained from the WLS analysis (not shown in the Table), although there is no significant correlation at the shallowest 1200-m instrument.

The two methods also give similar results if we use as our measure of fit the linear regression coefficient, β, between the observed and predicted data rather than the correlation coefficient. (Here β is defined as that number that when multiplied by the observed time series gives the best fit to the predicted time series in the least-squares sense; see Draper and Smith 1981.) However, it is noteworthy that at all instruments on the two inshore moorings, the WLS method provides a slight improvement in the linear regression coefficient. Each increase, taken individually, is not significant, but the consistent increase may be significant. We will return to this point in a later section, where a simple statistical test is used to show that a consistent but small increase in linear regression coefficients may indicate that one analysis method and/or dataset has significantly increased predictive skill over others.

A note here on statistical matters. The standard error of the regression coefficients is calculated using the method of Allen and Kundu (1978), while the equivalent degrees of freedom required for its computation (and also required for determining the statistical significance of the correlation coefficients) were obtained using the method of Davis (1976). A conservative estimate of the integral time scale is 5 days (Freeland et al. 1986; Church et al. 1986b). Confidence limits for correlation coefficients were computed as in Draper and Smith (1981).

If we consider the root mean-square (rms) of the signal at individual instrument sites, we find that the SVD method leads to a consistent underprediction of the measured rms; the two inshore instruments are the only exceptions. Although the WLS method also leads to underprediction of the measured rms, there is a general improvement at most sites. The exceptions are at the two instruments on the inshore mooring, where the WLS method overpredicts the rms.

One tentative conclusion we can draw from these results is that the WLS method gives a more energetic solution than the SVD method, but with increased noise. This is evidenced by the increase in predicted rms values without a corresponding increase in the correlation or regression coefficients.

TABLE 8. Predictions of velocity at line 2 and sea level (η) at lines 2 and 3 from the prognostic CTW model with modes analyzed at line 1 as upstream boundary conditions. Only two combinations of model and data are shown: the original SVD method using only velocity data, and the WLS method with the addition of sea-level data. Analyses are for a 90-day "eddy-free" period, 20 October 1983 to 18 January 1984, except for those instruments marked † and ‡, where only 67 and 28 days, respectively, of observed data were available. Root-mean-square (rms) units are centimeter per second for velocity and centimeter for sea level. The correlation coefficient, r, between the observed and predicted time series has a 95% significance level of about 0.30; double asterisks denote an insignificant correlation. The standard error for the linear regression coefficient β is about 0.10.

Instrument depth/ water depth	Obs. rms	Predicted rms		Error rms		r		β	
		SVD	WLS + η	SVD	WLS + η	SVD	WLS + η	SVD	WLS + η
75/135	12.4	12.3	13.8	8.7	8.6	0.75	0.79	0.74	0.88
125/135	12.4	13.2	14.7	8.7	8.4	0.77	0.82	0.82	0.97
75/200	12.4	8.7	10.3	10.5	10.2	0.55	0.61	0.38	0.51
125/200	12.3	8.7	10.4	10.2	9.9	0.57	0.63	0.40	0.53
190/200	10.4	6.8	8.8	8.2	8.3	0.62	0.64	0.40	0.54
190/1200†	13.1	2.2	3.6	12.3	12.7	0.42	**	0.07	0.06
450/1200	9.8	1.9	2.5	9.5	9.8	**	**	0.05	0.04
650/1200	9.7	2.4	2.5	9.9	9.9	**	**	0.01	0.01
1000/1200	7.6	2.3	2.4	7.2	7.6	0.31	**	0.10	0.05
450/2000‡	3.6	1.2	2.5	3.2	2.9	0.48	0.62	0.16	0.42
1000/2000	7.1	1.6	1.8	7.0	7.0	**	**	0.04	0.04
1900/2000	1.7	2.6	1.7	2.9	2.2	**	**	0.25	0.20
η (line 2)	5.5	5.1	6.1	3.7	3.5	0.76	0.82	0.70	0.91
η (line 3)	5.5	4.3	5.6	3.0	3.1	0.83	0.84	0.65	0.86

To summarize, we see that the weighted least-squares technique changes the composition of the CTW signal at both lines 1 and 2, with the first mode containing more energy than the second mode. The energy in the third mode is considerably reduced relative to the SVD analysis. The unexplained variance is the same for both methods, and there is little to choose between the predictions from the prognostic model. Both methods indicate the existence of a coherent mode 2 coastal-trapped wave propagating from line 1 to line 2 at approximately the theoretical phase speed. However, only the SVD method indicates a coherent mode 1 coastal-trapped wave.

b. Effect of adding sea-level data

We examine now the effect on the CTW mode composition of using additional data in the analysis. At line 1, the original nine current measurements are augmented by three sea-level time series, while at line 2, coastal sea-level data are added to the original 12 current measurements. The modal energy decomposition, total CTW energy, residual variance and correlation between mode 1 and mode 2 are summarized in Table 6 for line 1 and Table 7 for line 2.

The effect of additional data on the energy in the CTW modes depends on the analysis method used. At both line 1 and line 2, the SVD method shows an increase in the energy in mode 1 (36% and 9%, respectively), and a decrease in the energy in mode 2 (9% and 28%, respectively), to the extent that these two modes now contain about the same energy. The energy in mode 3 decreases by 13% at line 1 and increases by 7% at line 2, while the squared amplitude of the eddy mode increases by 12%–14%.

In contrast, the addition of sea-level data to the WLS analysis gives a decrease in mode 1 energy of 24% at line 1, and a decrease of 14% at line 2. The energy in mode 2 decreases by 16% at line 1, but increases by 22% at line 2. The net effect is to bring the ratio of mode 1 to mode 2 energy for the WLS method a little closer to the ratio obtained using the SVD method. However, the difference between the methods is still marked; the SVD method estimates that mode 1 and mode 2 contain approximately equal energy at both lines, while the WLS analysis suggests that mode 1 contains more than twice the energy of mode 2 at line 1, and 28% more energy at line 2.

The WLS analysis of mode 3 is unchanged by additional data; there is very little energy in this mode. The squared amplitude of the eddy mode is increased by 25% at line 1 and 31% at line 2.

The only change of any note in the total CTW energy caused by adding sea-level data is at line 1 for the WLS method, where the total energy drops by 21%. It still exceeds the total CTW energy obtained from the SVD method by 21%.

The residual velocity variance at both lines and for both techniques is increased slightly by the addition of sea-level data. However, there is a considerable decrease in the associated residual sea-level variance (see Tables 6 and 7). To fit the sea-level data, it is necessary to relax the fit to the velocity data. The advantage is that there is now a relatively uniform residual variance for all available measurements, although in all cases the sea-level residual exceeds the velocity residual. This might indicate that we have not chosen our row (equation) weights exactly right.

For all combinations of methods and datasets, the velocity and sea-level residuals are somewhat larger at line 2 than at line 1. We attribute this to the increased eddy activity at line 2, as evidenced by the fact that the squared size of the eddy mode is about twice as large at line 2 as at line 1. The time-independent eddy mode will not be able to represent fully the eddy activity at either line. Hence the portion of the eddy signal that cannot be represented as either the eddy mode or one of the CTW modes will appear as residual variance.

The inclusion of sea-level data leads to an increase in the alongshore coherence of modes 1 and 2 for both methods, while mode 3 remains incoherent. For the SVD method, there is a slight increase in coherence at all frequencies of interest (0.05–0.50 cycles per day), although changes to the corresponding phase speeds are not significant (Fig. 2c). However, for the WLS technique, the increase in coherence for mode 1 is quite marked, with significant coherence now found at all frequencies higher than 0.1 cpd (cycles per day) (Fig. 2d). Mode 2 is coherent at the 95% level only at frequencies less than 0.2 cpd, but it is at these frequencies that the majority of energy resides. The mode 1 and mode 2 phase speeds estimated from all data using the WLS method are close to being independent of frequency, and the values are closer to the theoretical values than are those of the SVD analysis.

Finally, we consider predictions of velocity and coastal sea level from the prognostic CTW propagation model described in the previous section. Including sea-level data at line 1 in either the SVD or the WLS analysis increases the amplitude of the predictions at line 2 by an average of about 7%. This increase is desirable because of the tendency of both methods to underpredict the amplitude at most instruments (see the previous section). There is also a general improvement in the correlation between the predicted and observed data; the extent of this increase is largest for the sea-level time series. However, in all cases the increase in the value of the correlation coefficient is not significant at the 95% confidence level. The addition of sea-level data to both methods also has the effect of increasing the linear regression coefficients at most instruments. However, in most cases these increases were below the standard error in these coefficients and so individually could not be considered significant. In the next section it will be shown that the fact that nearly all regression coefficients increased is significant, thus allowing the selection of a preferred method and dataset.

To summarize this section, the addition of sea-level data has a different effect on the results of the SVD and WLS methods. The mode 1 to mode 2 energy ratio increases for the SVD method, and decreases for the WLS method, thus narrowing the gap between the results from these methods. However, the WLS method still results in mode 1 containing at least twice the energy of mode 2, while the SVD results indicate these modes contain about the same energy. Estimates of the energy in mode 3 are not changed substantially; the WLS method results in very little energy in this mode, while the SVD method suggests mode 3 contains about one-fifth of the total CTW energy. For both methods, the sea-level residual is reduced considerably at the expense of a slightly higher velocity residual. The additional data improves the coherence of both mode 1 and mode 2 between line 1 and line 2; the coherence for the WLS method in particular is improved, with the phase speed close to being independent of frequency over the range of energetic frequencies. Predictions of line 2 velocity and coastal sea-level by the prognostic model are generally improved, although none of the improvements by themselves could be considered statistically significant.

c. *Are the modifications significant?*

We observed in the previous section that the WLS method led to a slight improvement in the prediction of velocity and sea-level data along the coast. We also observed that the addition of sea-level data to the analysis led to a similar improvement for both methods. These improvements occurred at most instruments, whether measured by the correlation coefficient or by the linear regression coefficient. The fact that the improvement was consistent across most instruments suggests that these results are not random. What is needed is a method for assessing the likelihood that such a consistent result is due to chance.

We will concentrate on the linear regression coefficient between the predicted and observed data. The fact that this value is normally distributed if the data are normally distributed (Draper and Smith 1981) makes the analysis easier. To remove the effect of instrument location, we use a paired test in which the variables are the differences between the linear regression coefficients at identical sites for the two models being tested. These differences are scaled by the standard error of the difference; this standard error is the square root of the sum of squared standard errors of the two individual regression coefficients. The scaled differences are normally distributed with mean zero and standard deviation of unity. We can now test the mean of these differences to see how likely it is that such a value occurred by chance.

The mean scaled difference for each combination of experiments is shown in Table 9, together with the probability that this value could occur by chance. Only those instruments that have a statistically significant regression coefficient for all experiments are included in this test; these are the 5 shallow current meters and the sea-level measurements at line 2 and line 3. However, the conclusions would be unchanged if all instruments were included. If we reject the null hypothesis that two experiments give the same linear regression coefficients at the 95% level, then any two experiments for which the probability is less than or equal to 0.05 can be considered as having significantly different predictive skill. If the mean scaled difference is negative, then experiment 2 has greater skill than experiment 1. The figures in Table 9 show that the WLS method used to analyze both velocity and sea-level data has significantly greater predictive skill than any other combination of method or dataset. This is the only conclusion that may be drawn from these figures at the 95% level.

TABLE 9. Statistical measure of whether one combination of method and dataset is significantly better or worse than another. The third column consists of the difference between linear regression coefficients at line 2, scaled by the standard error and averaged. A negative value means that experiment 2 has larger regression coefficients than experiment 1. The fourth column is the probability that the mean scaled difference could occur by chance. See the text for further explanation.

Experiment 1	Experiment 2	Mean scaled difference	Probability
SVD	SVD + η	−1.25	0.211
SVD	WLS	−1.12	0.263
SVD	WLS + η	−3.19	0.001
SVD + η	WLS	0.08	0.935
SVD + η	WLS + η	−1.96	0.050
WLS	WLS + η	−1.96	0.050

d. *Presence of mode 3*

To establish whether a particular mode has a significantly nonzero amplitude relative to the noise level in the analysis, we conduct a Student's t-test of the modal amplitudes relative to their estimated errors. The estimated errors for the WLS method are calculated from (15), while the estimated errors for the SVD method are calculated from a similar expression (see Lawson and Hanson 1974). The test is applied at all times (every 12 h), and the percentage of times that a mode is significantly different from zero at the 95% level recorded. Because the modal amplitudes fluctuate about zero, a mode will be considered insignificant by this test near a zero crossing. However, if the mode is significantly nonzero and of opposite sign either side of the zero crossing, then the zero crossing must be considered significant. We have not found a way to include this effect in the test. The results are clear nonetheless.

The only experiments in which mode 3 is significant more than 2% of the time are those involving the SVD method at line 1. The other modes are typically sig-

nificant 30%–50% of the time. We note that the SVD method generally gives somewhat larger values for the percentage of time a mode is significant and that adding sea-level data generally increases this percentage.

This analysis is likely to be quite strongly dependent on the weight matrices used in the analysis. Not only does the solution depend on the weight matrices, but so does the standard error. Hence the t values (ratio of amplitude to standard error) might vary considerably as the assumptions about weighting change. Since we believe that the WLS method uses more realistic weights than the SVD method, we conclude that mode 3 is not significant at either line 1 or line 2.

e. *Significance of energy estimates*

If it is assumed that the modal amplitudes have a normal distribution, then the modal energy estimates have a chi-square distribution. The number of degrees of freedom is given by the length of the experiment divided by the integral time scale. The latter was estimated by both Freeland et al. (1986) and Church et al. (1986b) to be between 3 and 5 days. Hence a conservative estimate for the number of degrees of freedom for estimating modal energy is $180/5 = 36$. With this estimate, the 95% confidence interval for the modal energy estimates, given as a percentage of the estimate, is +69% and −34%. If the shorter integral time scale of 3 days is used, the 95% confidence limits are +48% and −28%.

5. Discussion

The results we obtain using the weighted least-squares method raise a number of interesting issues regarding the oceanography of the ACE region, and CTW measurement experiments in general. The first of these is the lack of correlation between the mode 1 and mode 2 amplitude time series at both line 1 and line 2. This indicates that the generation mechanism for the two modes may not be as strongly linked as the analytical studies of Buchwald and Kachoyan (1987) and Middleton (1988) might suggest.

Theoretical studies by Clarke and Van Gorder (1986) and Lopez and Clarke (1989) suggest that seven or more modes may be needed to describe adequately the alongshore velocity field associated with CTWs (although fewer are required to describe coastal sea level). However, the work of Chapman (1987) in the CODE region, and Church et al. (1986b) in the ACE region, indicate that fewer modes are necessary to explain adequately the observed signal. In addition, both works point to the impracticality of extracting more than a few CTW modes from most instrument arrays. Our finding that the amplitude of mode 3 is insignificant relative to the other modes is consistent with these conclusions, although we cannot discount the existence of modes higher than the third. (A referee has pointed out that Figs. 2 and 3 of Lopez and Clarke 1989 indicate that many modes are necessary only on the inner shelf, and that elsewhere fewer modes are needed to describe the flow. This comment helps reconcile theory and observation.)

Our best estimate of the total CTW energy at line 1 during the ACE period, using the weighted least-squares method and sea-level data, is 2.4×10^8 W. This is an increase of 26% over the original ACE results. Three separate energy sources have been identified in the literature. Estimates of the energy flux due to local wind forcing in Bass Strait range from 1.1 to 2.0 ($\times 10^8$ W) (Clarke 1987; Morrow et al. 1990). Local wind forcing between the eastern end of Bass Strait and line 1 is estimated to add about 0.3×10^8 W (Freeland et al. 1986). Finally, CTWs incident at the western end of Bass Strait are thought to transfer energy through the Strait (Church and Freeland 1987; Middleton and Viera 1991; Baines et al. 1991). The last two studies showed incident CTWs to be at least as important as local wind forcing in driving an eastward energy flux in Bass Strait, though it should be noted that these conclusions were based on data from the post-ACE period April to June 1984. Adding the energy fluxes from the three sources gives a tentative lower bound of 2.5×10^8 W for the energy flux incident at line 1. The fact that our estimate of total CTW energy is slightly lower than this lower bound could be attributed to the penalty function minimized to obtain a stable solution. One term in the penalty function involves the sum of energy fluxes in the modes, so it is likely that our estimate of energy flux is also a lower bound. It is difficult to draw any stronger conclusions given the large confidence limits on our calculation of CTW energy, and the inherent uncertainties in estimating energy sources.

An unexpected result of the original ACE study was the finding that mode 2 contained more energy than mode 1 at both line 1 and line 2 (Church et al. 1986a). Our results show that both the addition of sea-level data and the use of the weighted least-squares method act to increase the energy in mode 1 at both lines, marginally decrease the energy in mode 2 at line 1, and generally increase the energy in mode 2 at line 2. Overall we find that at line 1, mode 1 contains 2.3 times as much energy as mode 2, while at line 2 the ratio is 1.3 (see Tables 6 and 7). Our modal energy partitioning agrees with the "conventional" view of a CTW signal dominated by the first mode, but it should be noted that arguments have been proposed that lend weight to the opposing view. The analytical studies of Buchwald and Kachoyan (1987) and Middleton (1988), for example, both indicate that the dimensions of Bass Strait are such that an oscillatory flux through the strait may act to generate mode 2 in preference to mode 1 on the East Australian shelf. However, Middleton (1988) also points out that the ratio of mode 2 to mode 1 may change if allowance is made for the mismatch

in the strait and shelf water depths. Griffin and Middleton (1991) show that the energy flux in the first CTW mode in the vicinity of Sydney (50 km north of line 2) due to free propagation from Bass Strait was three to five times larger than that of the second mode. However, they concede that this figure may be an overestimate. In addition, their data comes from the 1984/85 summer period, a year after ACE.

An unusual aspect of our reanalysis is that, although the energy in mode 1 remains essentially unchanged from line 1 to line 2, the energy in mode 2 increases by a factor of 1.8. Wind forcing in the intervening region is unlikely to account for such an increase (Schahinger and Church 1994). The most obvious physical explanation is linked to the increased activity of the East Australian Current at line 2 relative to line 1. It is possible that CTWs may be generated by eddies hitting the continental slope and shelf. Louis (1989) has suggested that this mechanism may have been responsible for the eventlike bursts of wave activity evident in many of the ACE current meter records, particularly at line 2 during January 1984. Such a mechanism may preferentially generate higher-order CTW modes.

In all the discussions above involving energy estimates, it must be remembered that the confidence limits on such estimates are quite large due to the small number of degrees of freedom. Although the ACE experiment ran for about half a year, the signals it was designed to measure had integral time scales of between 3 and 5 days, giving an effective number of degrees of freedom of only 36 to 60.

The modal decomposition at line 1 differs considerably between the original SVD analysis and our weighted least-squares analysis with additional sea-level data. The original ratio of modes 1, 2, and 3 was about 3:5:2, whereas our analysis suggests the ratio is more like 7:3:0. It was shown that the ability of our analysis to predict sea level and velocity along the coast was statistically significantly enhanced over the original method. However, we were surprised to find that the amplitude of predicted sea level and velocity increased by only 10%–20%, despite the substantial change in energy partitioning. The square root of the total CTW energy at line 1 in our reanalysis increased by 12% over the SVD analysis, and it is this increase that is reflected in the predictions. We conclude that it is the total CTW energy input at line 1 that is more important than the distribution of that energy into CTW modes. However, this conclusion is almost certainly dependent on the residual fit at line 1 being acceptable.

6. Conclusions

We have applied a weighted least-squares method to reanalyzing data from the Australian Coastal Experiment. The original time-domain analysis used a singular-value decomposition. Novel features of our method are the use of generalized cross validation to estimate the signal-to-noise ratio of the dataset, and the use of an iterative method to estimate the matrix column scales. Sea-level data not used in the original analysis were also incorporated. A comprehensive study of the effects on the solution of the analysis method versus the use of additional data revealed a complicated interaction. Adding sea-level data had a different effect on the two methods, although it did improve alongshore coherence between modes in both cases. The additional data tended to bring the two solutions into closer agreement, indicating that with more data the solution is less determined by the method used. Applying the original analysis method to the larger dataset led to mode 1 containing the same energy as mode 2 at both lines. Our weighted least-squares method applied to the entire dataset increased the energy in mode 1 substantially, left mode 2 essentially unchanged, and indicated that there was almost no energy in mode 3.

Furthermore, our reanalysis indicates that there is little correlation between the mode 1 and mode 2 amplitude time series at line 1. The original analysis method enforced such a correlation. The implication is that there may not be a common energy source for these two modes. This, together with our finding that mode 1 has more than twice the energy of mode 2 at line 1, suggests that further work is needed on the energy source of coastal-trapped waves on the southeastern Australian shelf.

We have demonstrated that our method together with the additional sea-level data gives significantly greater skill at predicting alongshore currents and sea level along the coast. However, we have noted deficiencies in the ability of the ACE instrument array to separate completely mode 1, mode 2, and the eddy mode. In addition, the duration of the ACE experiment leads to between 36 and 60 degrees of freedom, which in turn leads to disappointingly large confidence limits on estimates of modal energies. Results in this paper should be tempered with that knowledge.

Acknowledgments. This work was initiated and encouraged by John Church, and discussions with him have been most helpful. We thank Kathy Haskard for help with a number of statistical issues. We also thank Stephen Walker and John Wilkin for comments on two versions of this manuscript. PCM and RBS were funded by Australian Research Council grants.

REFERENCES

Allen, J. S., and P. K. Kundu, 1978: On the momentum, vorticity and mass balance on the Oregon shelf. *J. Phys. Oceanogr.*, **8**, 13–27.

Baines, P. G., G. Hubbert, and S. Power, 1991: Fluid transport through Bass Strait. *Contin. Shelf Res.*, **11**, 269–293.

Barth, N., and C. Wunsch, 1990: Oceanographic experiment design by simulated annealing. *J. Phys. Oceanogr.*, **20**, 1249–1263.

Buchwald, V. T., and B. J. Kachoyan, 1987: Shelf waves generated by a coastal flux. *Aust. J. Mar. Freshwater Res.*, **38**, 429–437.

Chapman, D. C., 1987: Application of wind-forced, long, coastal-trapped wave theory along the California coast. *J. Geophys. Res.*, **92,** 1798–1816.

Church, J. A., and H. J. Freeland, 1987: The energy source for the coastal-trapped waves in the Australian Coastal Experiment region. *J. Phys. Oceanogr.*, **17,** 289–300.

——, ——, and R. L. Smith, 1986a: Coastal-trapped waves on the east Australian continental shelf. Part I: Propagation of modes. *J. Phys. Oceanogr.*, **16,** 1929–1943.

——, N. J. White, A. J. Clarke, H. J. Freeland, and R. L. Smith, 1986b: Coastal-trapped waves on the east Australian continental shelf. Part II: Model verification. *J. Phys. Oceanogr.*, **16,** 1945–1957.

Clarke, A. J., 1977: Observations and numerical evidence for wind-forced coastal trapped long waves. *J. Phys. Oceanogr.*, **7,** 231–247.

——, 1987: Origin of the coastally trapped waves observed during the Australian Coastal Experiment. *J. Phys. Oceanogr.*, **17,** 1847–1859.

——, and S. Van Gorder, 1986: A method for estimating wind-driven frictional, time-dependent, stratified shelf and slope water flow. *J. Phys. Oceanogr.*, **16,** 1011–1026.

Davis, R. E., 1976: Predictability of sea surface temperature and sea-level pressure anomalies over the North Pacific Ocean. *J. Phys. Oceanogr.*, **6,** 249–266.

Draper, N. R., and H. Smith, 1981: *Applied Regression Analysis.* 2d ed. Wiley, 709 pp.

Forbes, A. M. G., 1985: Sea-level data from the Australian Coastal Experiment—a data report. CSIRO Marine Lab. Rep. No. 171, 16 pp.

Freeland, H. J., F. M. Boland, J. A. Church, A. J. Clarke, A. M. G. Forbes, A. Huyer, R. L. Smith, R. O. R. Y. Thompson, and N. J. White, 1986: The Australian Coastal Experiment: A search for coastal trapped waves. *J. Phys. Oceanogr.*, **16,** 1230–1249.

Golub, G., M. Heath, and G. Wahba, 1979: Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics,* **21,** 215–224.

Griffin, D. A., and J. H. Middleton, 1991: Local and remote wind forcing of New South Wales inner shelf currents and sea-level. *J. Phys. Oceanogr.*, **21,** 304–322.

Huyer, A., R. L. Smith, P. J. Stabeno, J. A. Church, and N. J. White, 1988: Currents off south-eastern Australia: results from the Australian Coastal Experiment. *Aust. J. Mar. Freshwater Res.*, **39,** 245–288.

Lawson, C. L., and R. J. Hanson, 1974: *Solving Least Squares Problems.* Prentice-Hall, 340 pp.

Lopez, M., and A. J. Clarke, 1989: The wind-driven shelf and slope water flow in terms of a local and a remote response. *J. Phys. Oceanogr.*, **19,** 1091–1101.

Louis, J. P., 1989: Current-shelf interactions during the Australian Coastal Experiment. *Aust. J. Mar. Freshwater Res.*, **40,** 571–585.

McIntosh, P. C., and G. Veronis, 1993: Solving underdetermined tracer inverse problems by spatial smoothing and cross validation. *J. Phys. Oceanogr.*, **23,** 716–730.

Menke, W., 1984: *Geophysical Data Analysis: Discrete Inverse Theory.* Academic Press, 260 pp.

Middleton, J. F., 1988: Longshelf waves generated by a coastal flux. *J. Geophys. Res.*, **93,** 10 724–10 730.

——, and F. Viera, 1991: The forcing of low frequency motions within Bass Strait. *J. Phys. Oceanogr.*, **21,** 695–708.

Morrow, R. A., I. S. F. Jones, R. L. Smith, and P. J. Stabeno, 1990: Bass Strait forcing of coastal trapped waves: ACE revisited. *J. Phys. Oceanogr.*, **20,** 1528–1538.

Schahinger, R. B., and J. A. Church, 1994: The prediction of wind-forced currents and sea-level on the southeast Australian continental shelf. *J. Phys. Oceanogr.*, submitted.

Tarantola, A., 1987: *Inverse Problem Theory.* Elsevier, 613 pp.

Wahba, G., 1990: *Spline Models for Observational Data.* SIAM, 169 pp.

——, and J. Wendelberger, 1980: Some new mathematical methods for variational objective analysis using splines and cross-validation. *Mon. Wea. Rev.*, **108,** 1122–1145.

Wang, D.-P., and C. N. K. Mooers, 1976: Coastal-trapped waves in a continuously stratified ocean. *J. Phys. Oceanogr.*, **6,** 853–863.

Wunsch, C., 1978: The North Atlantic general circulation west of 50°W determined by inverse methods. *Rev. Geophys. Space Phys.*, **16,** 583–620.