Mapping tropical Pacific sea level: Data assimilation via a reduced state space Kalman filter

Mark A. Cane,[1] Alexey Kaplan,[1] Robert N. Miller,[2] Benyang Tang,[3] Eric C. Hackert,[4] and Anthony J. Busalacchi [5]

Abstract. The well-known fact that tropical sea level can be usefully simulated by linear wind driven models recommends it as a realistic test problem for data assimilation schemes. Here we report on an assimilation of monthly data for the period 1975-1992 from 34 tropical Pacific tide gauges into such a model using a Kalman filter. We present an approach to the Kalman filter that uses a reduced state space representation for the required error covariance matrices. This reduction makes the calculation highly feasible. We argue that a more complete representation will be of no value in typical oceanographic practice, that in principle it is unlikely to be helpful, and that it may even be harmful if the data coverage is sparse, the usual case in oceanography. This is in part a consequence of ignorance of the correct error statistics for the data and model, but only in part. The reduced state space is obtained from a truncated set of multivariate empirical orthogonal functions (EOFs) derived from a long model run without assimilation. The reduced state space filter is compared with a full grid point Kalman filter using the same dynamical model for the period 1979-1985, assimilating eight tide gauge stations and using an additional seven for verification [*Miller et al.*, 1995]. Results are not inferior to the full grid point filter, even when the reduced filter retains only nine EOFs. Five sets of reduced space filter assimilations are run with all tide gauge data for the period 1975-1992. In each set a different number of EOFs is retained: 5, 9, 17, 32, and 93, accounting for 60, 70, 80, 90, and 99% of the model variance, respectively. Each set consists of 34 runs, in each of which one station is withheld for verification. Comparing each set to the nonassimilation run, the average rms error at the withheld stations decreases by more than 1 cm. The improvement is generally larger for the stations at lowest latitudes. Increasing the number of EOFs increases agreement with data at locations where data are assimilated; the added structures allow better fits locally. In contrast, results at withheld stations are almost insensitive to the number of EOFs retained. We also compare the Kalman filter theoretical error estimates with the actual errors of the assimilations. Features agree on average, but not in detail, a reminder of the fact that the quality of theoretical estimates is limited by the quality of error models they assume. We briefly discuss the implications of our work for future studies, including the application of the method to full ocean general circulation models and coupled models.

1. Introduction

The product of the work reported here is a sequence of maps of sea level in the tropical Pacific. The intense interest in satellite altimetry is but the most expansive of many testimonies to the importance of sea level as a diagnostic of the ocean state. Beginning with the seminal studies of *Wyrtki* [1973, 1975], sea level measurements have played a crucial role in developing our understanding of the tropical Pacific and El Niño. Its importance for El Niño prediction is

¹Lamont-Doherty Earth Observatory of Columbia University, Palisades, New York.
²College of Oceanic and Atmospheric Sciences, Oregon State University, Corvallis.
³Centre for Earth and Ocean Sciences, University of Victoria, Victoria, British Columbia, Canada.
⁴Hughes STX Corporation, Lanham, Maryland.
⁵NASA Goddard Space Flight Center, Laboratory for Hydrospheric Processes, Greenbelt, Maryland.

also well established [e.g., *Cane*, 1991; *Ji and Smith*, 1995; *Rosati et al.*, 1996].

The maps we produce are based on reports from the Pacific tide gauge network. A dynamical model driven by observed surface wind fields is combined with a data assimilation procedure to map this data. The model plus data assimilation is nothing more (or less) than an elaborate data interpolation scheme, one which imposes some dynamical constraints.

The general objective in data assimilation is to create the best analysis of the system state by combining incomplete and inaccurate measurements with output from an imperfect model. For a broad set of meanings of "best," if certain assumptions about the model and data hold, then it may be rigorously demonstrated that the optimal analysis is generated by the data assimilation procedure known as the Kalman filter. Crucial among these assumptions is the supposition that the model and data errors have a certain form and that we have perfect knowledge of them. For real problems in oceanography or meteorology, we all know this is not true at present. *Dee* [1995] argues cogently that to a significant extent, this is not conditional ignorance but is inevitable.

The usual objection to applying the Kalman filter (KF) to problems in meteorology or oceanography is computational

expense. The computational burden is indeed formidable: at every assimilation time the KF requires the update of the model state error covariance matrix, and since even a modest sized oceanographic or meteorological model has $N = O(10^5)$ variables, this $N{\times}N$ matrix has $O(10$ gigawords) or more.

Limits on available computational power thus compel us to abandon the brute force KF which requires updates of the complete model state error covariance. For those who find it dismaying to give up on the optimal approach for merely logistical reasons, below we will add some theoretical reasons to seek a different path (also see *Dee* [1991]). Briefly, we argue that our imperfect knowledge of error structures makes the full KF superfluous; that the limited sample available to us for data assimilation is unlikely to be in detailed agreement with the long-term error structure, which means that the full KF is not cost effective; and, finally, that the sense in which the full KF analysis would be "best" is not what we want after all.

A number of techniques have been used to reduce the computational burden of the KF. Some involve clever computational procedures for the full KF [cf. *Parrish and Cohn*, 1985; *Cohn and Parrish*, 1991], some involve procedural compromises such as dispensing with the update and instead using the constant asymptotic gain matrix [e.g., *Gourdeau et al.*, 1992; *Fukumori et al.*, 1993; *Fu et al.*, 1993], and others reduce the size of the model until the computation is feasible [e.g., *Miller and Cane*, 1989].

These approaches do not adequately address the other concerns mentioned above. We opt for a procedure that reduces the size of the state space for the KF without necessarily reducing the size of the model being integrated. This is a common strategy for adaptive estimation procedures where it is essential to reduce the number of unknowns. *Dee et al.* [1985], for example, parameterize the error covariance in terms of a few parameters that, multiply fixed, specified error structures. *Dee* [1991] reduces the state space by a factor of 3 by assuming a geostrophic relation between velocities and pressure. The long wave approximation in the tropical model we use [*Cane and Patton*, 1984] would allow a similar strategy, but a more general approach is taken here. Our method employs a relatively small set of basis functions assumed able to capture the significant structure. In principle, this set is arbitrary, but here we use multivariate empirical orthogonal functions (EOFs). A similar approach was taken by *Hernandez and Calderon* [1991], who used a severely truncated set of spherical harmonics in their application of the KF to an atmospheric model. The oceanographic work most similar to what is done here is that of *Fukumori and Malanotte-Rizzoli* [1995] (FMR hereafter), who reduced the KF state space by using a grid coarser than that of the primary model for the KF. They used EOFs to guide the choice of grid but did not use the EOFs as a basis for the KF updates as is done here. They also differ in using the asymptotic filter, as given by *Fukumori et al.* [1993]. They justify the reduction by computational necessity, whereas we believe that the full KF would not be worthwhile even if it were affordable. Nonetheless, the parallels between their work and ours are strong, and the success of the reduced state space approach in two very different oceanographic contexts (ours is a linear, large-scale tropical assimilation of real tide gauge data; FMR's is a nonlinear eddy rich simulation of an idealized midlatitude jet) builds confidence in the soundness of the general approach. Furthermore, our methods generalize to more complex models,

including incorporation of diabatic physics in the error estimation procedure.

The plan of the paper is as follows. In the next section we quickly review the KF and present our case that the full KF would be undesirable even if feasible. Section 3 explains our general methods and section 4 their implementation for the tropical Pacific sea level problem: our procedures for reducing the state space, for performing the covariance updates and combining model forecast and data, and for generating an estimate of the system noise. Section 5 introduces the data we use and compares our approach to a full grid point KF. The main body of results is given in section 6. We conclude by discussing some of the implications of our work, including possible extensions to more complex situations.

2. A Brief Review of the Kalman Filter

To establish notation and to attempt to justify our approach, it is first necessary to review some aspects of the KF. Our thinking on this subject has been strongly conditioned by the pioneering work adapting the KF to geophysical fluids by M. Ghil and collaborators [e.g., *Ghil et al.*, 1981; *Ghil*, 1989; *Dee*, 1991; *Ghil and Malanotte-Rizzoli*, 1991; *Jiang and Ghil*, 1993]. Since our exposition follows their approach, especially that of *Ghil and Malanotte-Rizzoli* [1991], it can be rather abbreviated. The reader in need of more detail should consult that reference (or *Miller and Cane* [1989], or *Jiang and Ghil*, [1993]).

As noted in the introduction, our object is to obtain the best analysis field w^a, i.e. the one closest to the true state w^{true} by combining the (forecast) model state vector w^f with the vector of observations w^o. We assume that both the model and observations are unbiased estimates of the true state, in which case any linear estimate combining the two may be written in the form

$$w^a(t) = w^f(t) + K\,[w^o(t) - Hw^f(t)]. \qquad (1)$$

The matrix H allows for a general linear relation between the model state and the observations. The observations need not be at the model grid points or even be of the same variables as in the model. (In the former case, H would be an interpolation formula.) Typically (and especially in oceanography), the dimension of w^o is much less than the dimension N of w^f, which includes all model variables at all grid points. If the observing locations vary with time (e.g., if there are data dropouts), then H varies with time.

The goal is to find the $K(t)$ that minimizes the expected error in $w^a(t)$; the least squares answer is

$$K = P^f\,H^T\,(H\,P^f H^T + R)^{-1} \qquad (2)$$

where $P^f(t)$ is the forecast error covariance, $R(t)$ is the observational error covariance, and superscript T denotes transpose. The essence of the formulas (1) and (2) is to weight the model and observed variables inversely as their expected errors.

This K gives the best analysis in the sense of minimizing

$$<(w^a - w^{true})^T (w^a - w^{true})>, \qquad (3)$$

where the brackets denote the expected value of the ensemble. Minimizing (3) yields a "best" analysis that may violate desirable dynamical or smoothness constraints. If the latter are desired, they must be imposed, as is often done in variational data assimilations. We return to this point below.

Calculating $\mathbf{K}$ requires knowledge of the observational error covariance $\mathbf{R}$ and the forecast error covariance $\mathbf{P}^f$. The former is determined by the characteristics of the instruments in the observing array. It is also influenced by sampling considerations. For example, in situ instruments typically sample only a single point, while the model values are meant to be representative of averages over a grid box ($10^4 \, km^2$ in our case). It is usually appropriate to take the errors as spatially uncorrelated (at least for in situ data). How to find $\mathbf{P}^f$ is less obvious, and it is the special genius of the KF to use the model dynamics to do so.

Let Ψ be the dynamical model that advances the system state from one time to the next, i.e.,

$$\mathbf{w}^f(t+1) = \Psi \, \mathbf{w}^a(t) + \Phi \, \tau(t). \qquad (4)$$

We take the dynamics to be linear and include a possible external forcing τ mapped onto the model response by a mapping Φ. We allow the possibility that the dynamics vary in time: $\Psi = \Psi(t)$. In the application to tropical sea level under study here a constant linear model is adequate [e.g., *Busalacchi and O'Brien*, 1981; *Cane*, 1984; *Busalacchi and Cane*, 1985] and the forcing τ is the surface wind stress. Given (4), it follows that

$$\mathbf{P}^f(t+1) = \Psi \, \mathbf{P}^a(t)\Psi^T + \mathbf{Q}(t) \qquad (5)$$

where $\mathbf{P}^a$ is the analysis error covariance, i.e.,

$$\mathbf{P}^a = <(\mathbf{w}^a - \mathbf{w}^{true}) \, (\mathbf{w}^a - \mathbf{w}^{true})^T>$$

which may be found from the formula

$$\mathbf{P}^a = (\mathbf{I} - \mathbf{KH}) \, \mathbf{P}^f. \qquad (6)$$

$\mathbf{Q}$ is the single step "system noise" covariance, the recognition that the model is not perfect:

$$\mathbf{Q}(t) = <\mathbf{q}(t) \, \mathbf{q}(t)^T> \qquad (7)$$

where $\mathbf{q}(t)$ is the error the model makes in the transition from time t to time $t+1$:

$$\mathbf{w}^{true}(t+1) = \Psi\mathbf{w}^{true}(t) + \Phi \, \tau \, (t) + \mathbf{q}(t). \qquad (8)$$

Note that $\mathbf{q}$ includes the influence of wind errors.

We further assume that the system noise is uncorrelated in time [cf. *Dee*, 1995;] (FMR, section 3.2):

$$<\mathbf{q}(t) \, \mathbf{q}(t')^T> = 0 \quad \text{if } t \neq t'.$$

It remains to specify the system noise spatial covariance $\mathbf{Q}$. The KF might be said to trade the problem of specifying $\mathbf{P}^f$, which is essentially what is done in optimal interpolation (OI) schemes, for that of specifying $\mathbf{Q}$. This trade is quite favorable. $\mathbf{P}^f$ evolves as the system evolves and changes according to the quantity and quality of the observational data assimilated as well as the model dynamics (see (5) and (6)). $\mathbf{Q}$ is a property of the system independent of the data and the assimilation process, which makes it easier to determine. The method we use below to specify it clearly relies on this independence.

Nonetheless, a detailed specification of $\mathbf{Q}$ remains a formidable task. In principle, it requires us to specify $N\times(N+1)/2$ independent numbers, where N is $O(10^5)$ or more. It is questionable whether we know that many meaningful numbers about the ocean, let alone its differences from an ocean model. Any meaningful $\mathbf{Q}$ we could write down would have to be specified by a far smaller number of parameters.

For the moment, let us suppose that we somehow possess a complete and precise $\mathbf{Q}$. We then face the famous problem of the KF's insupportable expense: calculating the $\mathbf{P}$ from (5) and (6) requires several times N^3 multiplications per time step. Let us go a step further and assume that the calculation is feasible, though expensive. We now argue that carrying out the detailed calculation is unlikely to be much help.

The crux of the argument is the mismatch between the very short duration of oceanographic time series and the very slow convergence of the covariance matrices with sample size. (Note that this is not the issue of the convergence of the KF equations (2), (5), and (6), which is fast.) There are very few oceanographic time series spanning more than a few decades. For the problem addressed in the present study only a handful of tropical tide gauge records exists before the 1980s. Since only monthly data are meaningful for our purpose of mapping the climatically important variations, the sample over which we will be assimilating data consists of fewer than 300 time points. We plan to work with $\mathbf{Q}$, the expected system noise covariance for the ensemble, whereas we would do better to use $\mathbf{Q}_s$, the expected noise statistics for the sample we are working on. (Of course, it would be best to know the precise sequence of errors $\mathbf{q}(t)$, but then all this machinery would be unnecessary.) Since we don't know $\mathbf{Q}_s$, we use the best estimate we have for it, which is $\mathbf{Q}$.

Unfortunately, since the convergence of a sample covariance to the true ensemble covariance is quite slow, like the square root of the sample length, it is unlikely that $\mathbf{Q}$ is very close to $\mathbf{Q}_s$ for samples so small compared to N. A nice illustration of this point relevant to the tropical sea level problem appears in Figure 1 of *Miller* [1990]. The major structures of $\mathbf{Q}$ and $\mathbf{Q}_s$ (e.g., the largest-scale eigenfunctions) are likely to be close, but it is statistically unlikely that the two will agree in detail. We conclude that the expensive calculation required to compute the $\mathbf{P}$ in great detail is unlikely to pay off in a great improvement in the average analysis over the necessarily short duration of oceanographic assimilations.

Thus far we have argued that we can't fill the system noise covariance fully with meaningful numbers and that even if we could the fine details are statistically unlikely to help over the relatively short duration of our assimilations. We are thus encouraged to reduce the number of degrees of freedom in the covariance calculation to reflect what error information we can know and to avoid expending computing resources on details unlikely to be useful.

We further argue that the analysis will actually be better if error covariance detail is omitted: because oceanographic data are too sparse to support it, the detail creates mischief. The sketch in Figure 1 illustrates this. In the usual circumstance where the distance between data points is large compared to error decorrelation scales, the minimization (3) overfits the data locally, resulting in a solution that is not smooth and generally not consistent with expected dynamical constraints such as geostrophy; the "bump" in Figure 1 is a unphysical feature, and most of us would be unhappy with it. A smoother analysis would be preferable, even though it would have a larger error at the data points, and thus not be "optimal" in the usual KF sense of minimizing the least squares error (3).

This is essentially the issue of regularity discussed by *Bennett and Budgell* [1987]. They showed that difficulties arise if $\mathbf{Q}$ is insufficiently red. In the common case where the model does poorly at small scales, the true system noise $\mathbf{Q}$

Figure 1. Sketch illustrating the possibility of data assimilation procedures overfitting the data. Heavy line depicts the true system state, the light line the forecast state, and the dashed line the analysis after data assimilation. The dot represents the observation.

will have appreciable power at small scales. If this is not altered, the analysis puts no credence in the model at small scales, so it "draws to the data" as in Figure 1. The remedy is to insist that Q be red or to impose a smoothness constraint in some other way. *Dee* [1991] and *Jiang and Ghil* [1993] effectively reddened Q by insisting that the system noise obey a geostrophic relation between pressure and velocity errors. Imposing smoothness constraints is a long-standing strategy in variational assimilation procedures [e.g., *Sasaki*, 1970]. *Provost and Salmon* [1986] dealt with the sparsity of data by restricting the analysis to a limited set of (trigonometric) basis functions. Removing small-scale variability from the basis set obviously makes local overfitting impossible. With any of these strategies, Q is being forced to have fewer significant degrees of freedom than its full size would allow.

The arguments presented above prompt the strategy described in the next section. For purposes of specifying Q and calculating the P we reduce the state space to the small number of degrees of freedom adequate to carry what little we know about the system noise, while maintaining enough of the KF's ability to use the dynamics to propagate error information. Inter alia, this reduction ensures a satisfactory degree of smoothness. The procedure may also be viewed as a parameterization of the large matrices Q and P in terms of a relatively small number of parameters. Our version of (5) propagates these parameters in time rather than the full model error covariance. Full covariance matrices can be reconstituted from the parameters.

3. General Method

We seek representations of the error covariance matrices with fewer degrees of freedom than implied by the dimension of the model state space. We begin by finding a reduced representation of the model state space and a reduced model to accomplish the transitions from one update time to the next. Similar developments are given by FMR, *Xue et al.* [1994], and Y. Xue et al. (Predictability of a coupled model of ENSO

using singular vector analysis. Part 1: Optimal growth in seasonal background and ENSO cycles, submitted to *Monthly Weather Review*, 1996; hereinafter referred to Xue et al., submitted manuscript, 1996). Let w be the vector of all state space variables. For example, if the model variables are u,v,p defined at the set of space points x_u, x_v, x_p, respectively, then

$$\mathbf{w}(x,t) = (\mathbf{u}(\mathbf{x}_u, t), \mathbf{v}(\mathbf{x}_v, t), \mathbf{p}(\mathbf{x}_p, t))^T.$$

There is no need for the different model variables to be defined on the same grid or with the same number of points. Now write the state space vector in terms of factors in time and space:

$$\mathbf{w}(x,t) = \mathbf{E}(x)\, \mathbf{u}(t). \tag{9}$$

The equality in (9) means that the columns of the $N \times N$ matrix E are a complete set of basis functions for the model state space. The vector $\mathbf{u}(t)$ holds the amplitudes at time t of these basis functions. There are many possible choices for the columns of E, e.g., Fourier components for each model variable. The important requirement for us is that the choice is efficient in that it lets us truncate the number of columns in E (and hence the dimension of u) to $M \ll N$ without sacrificing anything essential. We choose multivariate EOFs (MEOFs) for the columns of E, in which case the elements of u are the principal components (PCs). It is convenient to take the basis set to be orthonormal; $E^T E = I$. Because of the truncation it is not also true that $EE^T = I$.

Then since

$$\mathbf{u} = \mathbf{E}^T \mathbf{w}, \tag{10}$$

left multiplying the model evolution equation (4) by E^T yields

$$\mathbf{u}^f(t+1) = \mathbf{A}\, \mathbf{u}^a(t) + \mathbf{E}^T\Phi\, \tau(t) + \mathbf{u}' \tag{11}$$

where

$$\mathbf{A} = \mathbf{E}^T\Psi\, \mathbf{E} \tag{12}$$

and u' accounts for the influence of the discarded modes at time t on the retained ones at time $t+1$. As in FMR, we assume that u' is negligible. For our sea level simulations it is easy to make this true by retaining enough MEOFs. As appears below, it is more problematic for error propagation.

There is a straightforward way to calculating the transition matrix A (cf. FMR or Xue et al. (submitted manuscript, 1996) for details). With A in hand we find the error covariances P in the reduced representation from reduced state space versions of (5), (6), and (2):

$$\mathbf{P}^f(t+1) = \mathbf{A}\, \mathbf{P}^a(t)\, \mathbf{A}^T + \mathbf{Q}(t); \tag{13}$$

$$\mathbf{P}^a = (\mathbf{I} - \mathbf{KH})\, \mathbf{P}^f; \tag{14}$$

$$\mathbf{K} = \mathbf{P}^f \mathbf{H}^T\, (\mathbf{H}\, \mathbf{P}^f \mathbf{H}^T + \mathbf{R})^{-1} \tag{15}$$

While these equations are formally identical to the earlier ones (apart from the change from Ψ to A), the meanings of the symbols have changed. Here

$$\mathbf{P}^{a,f} = \langle (\mathbf{u}^{a,f} - \mathbf{u}^{\text{true}})\, (\mathbf{u}^{a,f} - \mathbf{u}^{\text{true}})^T\rangle \tag{16}$$

and Q is the system noise appropriate to the reduced system. There is a relation between the new reduced mapping matrix H and the original full state space one $H^\dagger$

$$\mathbf{H} = \mathbf{H}^\dagger \mathbf{E}, \tag{17}$$

which shows that the observational data are now approximated

in terms of the M retained structures instead of the original N variables. The analogs of (17) for the statistical matrices, i.e.,

$$Q = E^T Q^\dagger E; \quad P = E^T P^\dagger E; \quad K = E^T K^\dagger; \qquad (18)$$

need not hold if the influence of the discarded modes on the retained ones is significant. The full system allows errors in modes omitted in the reduced system to propagate "upscale" into the retained modes. If this happens, then Q, the effective system noise in the reduced space, is not just a truncated version of the full $Q^\dagger$ but should be modified to account for it. The second equality shown in (18) holds only if this accounting is precise. We return to this issue below when we consider the construction of the system noise covariance.

Even if the first two equalities in (18) hold, the relation between the Kalman gain matrices in (18) need not be true. To see this, first write

$$P^\dagger = EPE^T + P'. \qquad (19)$$

P' contains all the entries in $P^\dagger$ related to the truncated modes. We will assume that P is correct so that the elements of P' representing the covariance of retained modes are all zero. Now define

$$R' = H^\dagger P' H^{\dagger T}, \qquad (20)$$

which allows us to write

$$
\begin{aligned}
K^\dagger &= EPH^T (HPH^T + R' + R)^{-1} + P'H^{\dagger T}(HPH^T + R' + R)^{-1} \\
&= \quad K'' \quad + \quad K'. \qquad (21)
\end{aligned}
$$

The second gain matrix K' does not project onto the retained modes at all $(EK' = EP' = 0)$; it is the part of the full gain matrix that puts the innovations into the discarded scales. As discussed in section 2, we consider it desirable to eliminate it to avoid overfitting to the data.

K'' differs from the K defined in (15) by the inclusion of the extra noise term R' in the denominator. As derived here, this term, which is reckoned at the observation points, arises from the model error covariance in discarded modes. It reduces the impact K'' gives the innovations (the observations) by "reminding" the filter that some of the innovation belongs in the discarded modes.

This suggests an alternate interpretation of R' as additional sampling error associated with the observations, appropriate if we regard the reduced KF procedure as defining the analysis at the retained scales only. Regarding $R'+R$ in (21) as total observational error R makes K'' identical to the K of (15). Discarding K' is an immediate consequence of this interpretation. Though they phrase it somewhat differently, this second interpretation is the one taken by FMR. (See their discussion following equations (14) and (17). Since Kalman filter theory assumes that H applied to the true state gives the true observation, it is clear that their n' is being treated as observational noise. Consequently, no term like K' appears in their derivation of an approximate filter.)

4. Implementation for Tropical Pacific Sea Level

4.1. Models and Wind Forcing

The numerical model used in our sea level calculation is a two vertical mode version of the *Cane and Patton* [1984] algorithm for solving the linear long wave approximation to

the shallow water equations on an equatorial beta plane. The implementation is quite close to that given by *Cane* [1984] and *Busalacchi and Cane* [1985]; additional details are given in those references. For present purposes the most serious shortcoming is probably the approximation that the stratification is horizontally uniform. As given by the cited references and *Miller and Cane* [1989], we use values typical of the mid Pacific: wave speeds of 2.86 and 1.85 m s^{-1} and length scales of 354 and 285 km for the first and second baroclinic modes, respectively. Note that sea level is not a primary model variable but is derived as a linear combination of the displacements of the two modes [e.g., *Cane*, 1984].

The standard model configuration used here is a domain that covers the tropical Pacific from 28.75°S to 28.75°N and 124°E to 80°W; the exact model domain is evident in Figure 2. The model time step is 0.25 months. Grid spacing is 2° in longitude and 0.5° in latitude, requiring approximately 9000 grid points in each of the two vertical modes. For the two-mode *Cane and Patton* [1984] model, six different variables are needed to define a complete state space: for each vertical mode $(k = 1,2)$ these are the two-dimensional fields of Rossby mode zonal velocity $u_R(x,y,t,k)$ and displacement $h_R(x,y,t,k)$ and the zonal array of Kelvin wave amplitudes $a_K(x,t,k)$. (There are a few additional numbers related to the peculiar way the algorithm handles boundary conditions in a domain that is not a simple rectangle [cf. *Cane and Patton*, 1984].)

The model has no thermodynamics or salinity; it is forced by surface wind stress. The wind anomaly fields we use derive from the Florida State University (FSU) pseudostress analysis [*Goldenberg and O'Brien*, 1981] smoothed and detrended in the manner described by *Cane et al.* [1986].

4.2. Reduced State Space

Our first task is to choose a suitable reduced basis E. We will use multivariate EOFs. First, the standard model is run for the period 1964-1991. After a 3 year spin-up period, all model variables are saved midmonth from January 1967 to September 1991, yielding time series of length $N_t = 297$. From these time series, MEOFs are calculated by the procedure described given by *Xue et al.* [1994] or Xue et al. (submitted manuscript, 1996).

The model state is thus transformed into the form (9) in terms of multivariate PCs $u(t)$ and EOFs $E(x)$. This set of multivariate EOFs (MEOFs) is the most efficient representation of the "total model variance" (with each variable weighted the same [cf. *Xue et al.*, 1994]). For our purposes, it is the "data compression" efficiency of EOFs that recommends this basis set (cf. the discussion given by *Lorenz*, [1956]). The least severe truncation we will use in what follows retains 93 MEOFs, which keeps 99% of the variance. In section 6 we will explore the consequences of more drastic truncations. Figure 2 illustrates how rapidly the series of MEOFs converges: it takes very few patterns to reproduce the large-scale features of the sea level field, and one might well question the realism of the details that emerge as more variance is retained.

While the MEOFs are the most efficient basis for the model states occurring in the simulation history, they are to be put to a different use, so certain potential problems must be considered. It may be that the truncated set is inadequate to represent the data (see (17)). Missing some structure in the data error is not a problem but a virtue: the data noise will be

Figure 2. Sea level maps for December 1990 based on the forced model run with no data assimilation. The different panels show the results of retaining 9, 17, 32, 51, and 93 multivariate empirical orthogonal functions (MEOFs) which keep 70%, 80%, 90%, 95%, and 99%, respectively, of the total model variance. The final panel shows the full field from the forced model run.

Figure 3. Locations of the tide gauge stations used in the assimilations and verifications.

automatically filtered out of the assimilation. The unhappy case is when the MEOFs, which derive from a limited simulation with an imperfect model, miss the true signal. Fortunately, as will be seen shortly, the truncated MEOF set we obtained appears to be adequate for the work reported here.

4.3. Tide Gauge Data and Error Structure

The data to be assimilated are the monthly mean sea level at the 34 tide gauge stations in the tropical Pacific shown in Figure 3. The data are obtained from the Sea Level Data Center (SLDC) at the University of Hawaii as monthly means with the tides removed. Additional processing was needed to convert to monthly anomalies. Following the procedure used by the SLDC, we choose the period 1975 to 1986 as a basis for calculating climatological monthly means. The record for this period was too sparse for this purpose at two stations, Baltra and Nuku Hiva. At these stations the average annual cycle was first calculated for the period 1987-1993. To this we then added the difference between the annual cycle from 1975 to 1986 and that from 1987 to 1993 at their nearest neighbors (Santa Cruz for Baltra and Penrhyn for Nuku Hiva). We excluded eight additional stations within the model domain: five (Naha, Chichijima, Midway, Bundaberg, Easter) are too close to the model's artificial boundaries; Quepos is too contaminated by tectonic movement; and for unknown reasons, Rikitea and Suva correlate poorly (< 0.2) with both the model run and nearby expendable bathythermographs (XBTs).

The error covariance for the tide gauge data is as discussed by *Miller and Cane* [1989]. Each station is taken to be accurate to 3 cm, and the errors at different stations are assumed to be uncorrelated: $R = (3 \text{ cm})^2 I$.

4.4. Modeling the System Noise

Our starting point for constructing a model system noise covariance matrix is the same as given by *Miller and Cane* [1989] and *Miller* [1990]. The basic assumptions are (1) the dominant source of model error is in the wind stress and (2) the wind errors are statistically homogeneous. As discussed in

those papers, not only is the first assumption quite reasonable, but those unhappy with it are free to interpret it merely as a device for generating a structure for errors due to other sources. The second assumption ignores the inhomogeneities in the variability of the wind field and in the distribution of ship tracks and so cannot be strictly correct. Nonetheless, it is a reasonable starting point, and since an important purpose here is to compare with those earlier KF studies and with *Miller et al.* [1995], we continue to use it.

Thus, denoting the wind error at (x,y) for month t by $e(x,y,t)$, we assume that

$$<e(x,y,t)\ e(x',y',t')>$$

$$= W \, \delta(t\text{-}t') \exp[-(x\text{-}x')^2/L_x^2 - (y\text{-}y')^2/L_y^2]. \quad (22)$$

Errors are uncorrelated month to month, and the spatial covariance is determined by three parameters: an amplitude (wind stress squared) and two length scales. *Miller et al.* [1995] and *Miller and Cane* [1989] discuss this form and the parameter values. We will use

$$W= (\rho_w/\rho_a C_D)^2 \ (164 \text{ m}^2/\text{s}^2), \quad L_x = 10°, \quad L_y = 4° \quad (23)$$

The values of the length scales are identical to those in *Miller et al.* [1995]. W is obtained by fitting to the differences from observations of a model run without assimilation. This value is approximately one third of the value of *Miller et al.* [1995] due to procedural differences, the most important of which is that they generate new random wind field three times per month while we do it once per month.

We will not reprise the discussion in these earlier papers but concentrate on the new issue of determining the system noise covariance in the reduced space, taking (22) and (23) as given. In doing so we make use of two Monte Carlo runs of the numerical model. In the Monte Carlo P run the model is forced by a random zonal wind stress generated from a covariance of the form (22) for 2001 months. The Monte Carlo Q run differs from the P run in one important feature: the model state is reset to zero at the midpoint of each month. The matrix of covariances among the model state variables for the Q run is a sample estimate of Q. The covariance matrix for the P run is a

sample estimate of P^f for the case of no data assimilation. The time mean is taken in place of the ensemble mean, and statistics are calculated from the 1965 months that remain after the first 36 months of spin up are discarded.

We do not calculate the huge covariance matrices for the full state space but only for the M (=93) retained MEOFs. At each month t the MEOFs are projected onto the full model fields w and we save the PC values, p(t) for the P run and q(t) for the Q run (as in (10) p or q = E^Tw). The covariances are then calculated from the resulting time series (compare (17)):

$$P^{(i)} = <p\ p^T>_t;\quad Q^{(i)} = <q\ q^T>_t.\tag{24}$$

We then ask whether p and $P^{(i)}$ satisfy reduced space equations like (11) and the steady form of (13):

$$p(t+1) = A\ p(t)\ +\ q(t);\tag{25}$$

$$P = A\ P\ A^T + Q.\tag{26}$$

Agreement with (26) was found to be unsatisfactory. We checked that the sample used in (25) was long enough not to be the problem. Hence we concluded that the difficulty must be the influence of the truncated modes on the retained ones, an influence which is omitted in (25) and (26). Trouble could be anticipated when working with the error scales (24) and only $O(100)$ modes. Crudely estimating that these modes divide the model domain into 10 x 10 regions gives a resolution of about 15° x 6°, leaving too much of the error energy at truncated scales. According to the arguments of section 2, the best diagnosis of the difficulty is that the full space system noise Q following from (23) and (24) is not red enough and should be changed. A possible fix is to change from the full Q to the truncated $Q^{(i)}$, which is a filtered version of the full space noise. However, in order to compare with grid point models that used the full space system noise, we pursued other modifications.

The noise model we use is the following. Using the time series p(t) we find A', the model that is the best fit to (26) in a least squares sense:

$$A' = <p(t+1)\ p(t)^T>_t <p(t)\ p(t)^T>_t^{-1}.\tag{27}$$

(This is a standard procedure for constructing a multivariate AR(1) model; for a related example see *Blumenthal* [1991]). We then calculate the system noise Q' from the appropriate version of (27):

$$Q' = P^{(i)} - A'\ P^{(i)}\ A'^T.\tag{28}$$

We have now effectively parameterized the "sub grid scale" (truncated mode) influence in two ways. First, we added A'-A to the model in an attempt to capture influences going from a retained mode to the set of discarded modes and then back to retained modes. Second, the noise covariance Q was changed to Q' to try to capture the direct influence of the discarded modes.

To best maintain the form (23) for the system noise in the reduced space, we were led to modify the transition model A as well as the noise estimates. FMR left A untouched but made ad hoc adjustments to the observational (R) and system noise (Q) estimates. (We recommend the discussion in their section 3.2.) As we argued earlier, these errors are poorly known and we see no clear reason a priori to prefer one approach to the other.

It remains to estimate R' from (19) and (20):

$$R' = H^\dagger P' H^{\dagger T} = H^\dagger P^\dagger H^{\dagger T} - HPH^T.\tag{29}$$

The first term on the right is the error covariance matrix $P° = <h(t)h(t)^T>$ of the model produced sea level at the observation points (the locations of the tide gauges), estimated from the Monte Carlo P run. We will further assume that all the correlated structure in $P°$ is captured by the P term; that is, by the retained modes. Then R' is diagonal.

Our estimate of R' is based on the model error without assimilation and is likely to be somewhat high. Recall that while we use the reduced state space for the error estimation component of the KF assimilation procedure, the full model is retained for the simulation itself. Thus R' is likely to become smaller as the assimilation reduces the error in the retained modes, and the model dynamics propagate the improvement in this part of the state into the remainder.

5. Comparison of the Reduced State Space KF With a Grid Point KF

The purpose of this section is to demonstrate that the performance of the reduced state space KF is comparable to that of a full state space KF, i.e., one which retains the full error covariance matrix with a row and column for every variable at every model grid point. We compare to the results of *Miller et al.* [1995], who used the same *Cane and Patton* [1984] numerical model employed here with the same parameter settings. The principal differences are that Miller et al. had to use a coarser grid (2° latitude x 5° longitude; ≈ 10^3 grid points) in order to reduce the state space to a computationally feasible size and that they add dissipation in order to suppress grid scale noise excited by the assimilation process.

Miller et al. [1995] used a total of 15 tide gauge stations, assimilating data from eight of them and withholding the other seven for verification. The period of the assimilation is the 7 years 1979-1985. We assimilated the same stations for the same period; results are presented in Table 1 in terms of rms errors. Comparing first the 93 EOF run with the grid point filter, we reach the general conclusion that the results are comparable. The rms errors rarely differ by even as much as 1 cm, and at only one point (Kanton) do the correlation coefficients (not shown) differ by as much as 0.1.

The only systematic difference we discern is a tendency for the grid point filter to be closer to the data at the assimilated stations and farther from it at the withheld stations. It would be consistent with the discussion of section 2 to suggest that the "higher-resolution" grid point filter tends to overfit the data locally, perhaps causing it to underweight more remote connections. Figure 4 illustrates results at a point where data are assimilated and one where they are withheld. These systematic differences are slight. More important, both assimilations improve markedly on the unfiltered simulation. At the withheld point, Yap, the two filtered runs do tend to be closer to each other than either is to the observations. However, they are about as far from the unfiltered run as they are from the data. The only important difference between the two filtered runs is that the reduced space version better captures the low at the end of 1982.

Table 1 also gives results from runs with fewer EOFs retained in the KF procedure. The comparisons at the assimilation points worsen noticeably as the number of EOFs is reduced. This is to be expected since the fewer the spectral (EOF) components, the less complete the fit to the data. But there is surprisingly little falloff in skill at the withheld

Table 1. Rms Differences with Observations of Grid Point and Reduced Kalman Filter (KF) Experiments for the Period 1979-1985

No. of Station	Name of Station	Location		Grid Point Runs		Reduced KF Runs for Different Dimensions of a Reduced Space					
		Lat	Lon	Unf	Filter	Unf	93 EOFS	32 EOFS	17 EOFS	9 EOFS	5 EOFS
1	RABAUL	4S	152E	4.48	0.86	5.49	1.47	1.94	2.43	2.48	3.28
2	JARVIS	0	160W	4.81	1.65	4.40	2.04	2.23	2.01	2.23	2.82
3	CHRISTMAS	2N	157W	6.28	2.41	6.07	3.26	3.70	3.83	4.11	4.87
4	SANTA CRUZ	1S	90W	6.41	1.87	6.61	2.13	2.28	2.43	2.57	3.01
5	CALLAO	12S	80W	6.40	1.42	6.74	2.59	2.90	3.03	3.22	3.37
6	KWAJALEIN	9N	168E	5.93	1.12	6.64	2.46	3.00	3.58	4.20	4.48
7	PENRHYN	9S	158W	6.01	1.41	6.65	3.61	4 19	4.29	4.07	4.09
8	TARAWA	1N	173E	5.66	1.66	5.92	2.44	2.78	3.16	4.06	4.59
9	KAPINGAMAR	1N	155E	5.26	3.62	5.71	3.58	3.69	3.74	3.61	3.76
10	KANTON	3S	172W	5.73	4.04	6.57	3.67	3.69	3.73	4.07	4.54
11	HONIARA	9S	160E	7.30	6.40	6.93	6.62	6.65	6.86	6.32	6.42
12	YAP	10N	38E	7.14	6.25	6.66	5.74	6.19	6.03	5.67	6.96
13	TRUK	7N	152E	6.56	5.41	6.68	5.16	5.42	5.08	4.75	4.69
14	NAURU	1S	167E	7.70	5.57	7.72	5.28	5.41	5 64	5.76	5.91
15	FANNING	4N	159W	7.76	5.30	7.96	5.67	5 46	5.86	6.70	7.27

	93 EOFS	32 EOFS	17 EOFS	9 EOFS	5 EOFS
Better than unf by 1 cm on number of withheld stations	5	5	5	5	4
Worse than unf by 1 cm on number of withheld stations	0	0	0	0	0
Better than grid KF on number of withheld stations	5	3	3	4	1
Worse than grid KF on number of withheld stations	2	3	4	3	6

The grid point assimilation uses a coarse grid (5° x 2°), and the reduced KF assimilations use a fine grid (2° x 0.5°). The latter were run with 93, 32, 17, 9, and 5 empirical orthogonal functions (EOFs) retained in the representation of the error covariance matrices. "Unf" indicates the unfiltered (no data assimilation) results. Data from Stations 9-15 were withheld in all assimilations

stations. We will pursue this issue in the next section, where more tide gauge stations are used in the assimilation.

The reduced KF runs use a finer grid for the dynamical model. Presumably, this is advantageous (although it is not evident in the two no assimilation (Unf) columns of Table 1). We think it is a fair advantage, since our underlying philosophy is to simplify the filter without sacrificing the complexity of the model, which would lose numerical accuracy and perhaps physical verisimilitude. Nonetheless, we reran the reduced KF experiment with the same coarse grid model used by *Miller et al.* [1995] for the grid point KF. Results were quite similar to those for the fine grid.

6. Principal Results

6.1. Comparison With Tide Gauge Data

In this section we report the results of assimilation experiments for the period 1975 to 1992 using 34 of the 36 tide gauge stations shown in Figure 3. Fanning and Jarvis, two of the stations used by *Miller et al.* [1995], are not

Figure 4. Comparison between results of Kalman filtering (KF) in the full grid point (dashed line) and the reduced (solid line) spaces. Ninety-three EOFs are retained in the reduced KF. Shown are sea level height at (top) Santa Cruz, a station where data are assimilated, and (bottom) Yap, a station where the data are withheld. Unfiltered model output (dotted line) and observations (stars) are also shown.

Table 2. Rms Differences with Observations of Reduced State Space KF Experiments for the Period 1975-1982

No. of Station	Name of Station	Location			Results of KF Runs (EOFs Retained / Variance Explained)				
		Lat	Lon	Unf	93 /99%	32 /90%	17 /80%	9 /70%	5 /60%
1	SANTA CRUZ	1S	90W	5.74	2.75	2.75	2.94	3.04	3.09
2	BALTRA	0	90W	5.45	2.41	2.38	2.43	2.62	2.74
3	NOUMEA	22S	166E	6 63	6.60	6.82	6.35	6.49	6.32
4	SAIPAN	15N	146E	8.78	6.48	7.19	6 55	6.40	5.80
5	YAP	10N	38E	6.89	7.83	8.10	7.35	6.61	5.88
6	MALAKAL	7N	134E	8.02	6.28	6.21	6.22	6.59	6.82
7	POHNPEI	7N	158E	6.73	3.62	4.15	4.53	4.28	4.70
8	KAPINGAMAR	1N	155E	6.44	3.82	4.08	4.16	4.11	4.13
9	NAURU	1S	167E	7 48	4.89	5.33	5.64	5.81	5.68
10	TARAWA	1N	173E	7.45	4.42	4.56	4.82	5.17	5.39
11	MAJURO	7N	171E	6.52	5.17	4.97	4.95	5.62	5.81
12	FRENCH FRI	24N	166W	8.57	8.73	8.51	8.36	8.52	8.61
13	HONOLULU	21N	158W	5.60	4.78	4.94	4.81	4 80	5.04
14	RAROTONGA	21S	160W	7.43	8.89	7.80	7.89	7.67	7.76
15	KANTON	3S	172W	6 93	4.54	4.81	5.10	5.05	5.36
16	PENRHYN	9S	158W	6.99	5.90	5 65	5.20	5.28	5.42
17	CHRISTMAS	2N	157W	6.79	5 42	5.84	5.64	5.43	5.87
18	PAPEETE	18S	150W	5 02	5.01	4.82	4.33	4.48	4.44
19	FUNAFUTI	9S	179E	6.26	5.34	5.20	5.32	5.88	6.16
20	HONIARA	9S	160E	7.39	6.06	6.77	6.68	7.27	8.01
21	RABAUL	4S	152E	7.25	4.95	4.54	4.61	5.04	5.84
22	NUKU HIVA	9S	140W	6.24	5.15	4.82	4.85	3.71	3.73
23	LA LIBERTA	2S	81W	7.09	5.34	5.41	5.35	5 52	5.61
24	CALLAO	12S	80W	6.45	4.36	4.11	4 18	3.92	4.08
25	DAVAO	7N	126E	6.13	5.16	5.06	4.94	5.06	4.30
26	ACAPULCO	14N	100W	5.81	6.48	5.90	6.02	5.77	5 85
27	WAKE	19N	167E	9.43	8.98	8.09	8.03	7 90	8.06
28	JOHNSTON	17N	170W	7 98	7.69	7.66	7.73	7.99	7.69
29	ISLA CEDRO	28N	115W	5.72	4.84	4.89	4.96	4.80	4.85
30	TRUK MOEN	7N	152E	6.60	3.89	4 24	4.28	4.18	4.63
31	KWAJALEIN	9N	168E	6.09	4.53	4.02	4.38	4.50	4.37
32	GUAM	13N	145E	7.79	6.49	6 14	6.21	6.18	6.20
33	PAGO PAGO	14S	171W	6.08	5 48	5.68	5 01	4.98	5.06
34	HILO	20N	155W	5.89	4.77	5.09	5.15	5.70	5.78

	Results of KF Runs (EOFs Retained / Variance Explained)				
Unf	93 /99%	32 /90%	17 /80%	9 /70%	5 /60%
RMS deviation, all stations	6 92	5 78	5.74	5.67	5.73	5.81
No. of stations better than unf by 1.0 cm		21	21	22	21	21
No. of stations worse than unf by 1.0 cm		1	1	0	0	0

In each case the data compared to were withheld in the assimilation. The reduced KF assimilations were run with 93, 32, 17, 9, and 5 EOFs retained in the representation of the error covariance matrices. "Unf" indicates the unfiltered (no data assimilation) results.

included in our standard data source, the Integrated Global Ocean Services System (IGOSS) Sea Level Program in the Pacific, which is maintained at the University of Hawaii SLDC. We chose not to use Fanning and Jarvis because the available time series were very spotty after 1983 and we were not confident of the data quality. Not all of the retained stations have continuous records for the entire period. Table 2 presents comparisons with observations at the 34 tide gauge stations in terms of rms differences. Each column for the KF assimilations presents the results of 34 separate runs, in each of which one station is withheld. All comparisons are against withheld data, as is usual in a cross-validation procedure. The data assimilation clearly helps: at almost every location the

rms error is reduced by 1 cm or more and the correlation is improved by 0.1 or more. The filtered results are worse than the unfiltered run at only one station, Rarotonga, far from the equator in the data sparse South Pacific.

There is some tendency for errors to be largest at stations close to 20°N or 20°S (also see Figure 7). This may be due to the phase speed errors of our model at such extra-equatorial latitudes [Cane and Patton, 1984], an explanation suggested to us by Mitchum [1994]. In a comparison of tide gauges with TOPEX data, he deduced a dominance of Rossby wave propagation at these latitudes. Consequently, small inaccuracies in propagation speeds result in sizable errors in sea level.

Figure 5. Time series of sea level height at selected stations. Assimilation values are from Kalman filter runs with 93 EOFs retained. The solid line is the KF result with the data withheld at the station, and the dashed line is the result with all data assimilated. Unfiltered model output (dotted line) and observations (stars) are also shown. Shown are (a) Nauru, (b) Kwajalein, (c) Pago Pago, and (d) Yap. See text for discussion.

Figure 5 illustrates the results at a few selected stations. Shown are Nauru, in the equatorial waveguide, where a mediocre simulation in the unfiltered run is greatly improved; Kwajalein, where the improvement in a somewhat better unfiltered result is less substantial; Pago Pago, where the assimilation makes only a slight improvement in a mediocre unfiltered result; and Yap, where a good unfiltered simulation is hardly changed. An examination of these figures and similar ones for the remaining 30 gauge sites disclose few universal rules. Generally, the results with data withheld (solid line) appear to be close to those with all data assimilated (dashed line). The model rarely reaches the extremes of the observed fluctuations, though Yap is an exception. Table 2 shows a tendency for better results at stations closer to the equator and for the worst results to be at stations poleward of 15°. (Isla Cedro, in the coastal extension of the waveguide, is a notable exception.) The possible reasons are legion. To begin with, the wind-driven simulation shows little skill in these higher latitudes. The nature of the dynamics means that the correlation scales are shorter in higher latitudes, so points there are less influenced by remote data than are low-latitude locations. Furthermore, there are fewer stations in these higher latitudes, so there are fewer neighboring observation points to influence the assimilation.

A surprising and disconcerting feature revealed in Table 2 and mirrored in the correlation statistics is how little difference there is when different numbers of EOFs are retained in the KF calculations. We postpone most discussion for the conclusion section and just point out the most salient features here. Though we did no formal test, the typical differences among the columns appear to be insignificant, an impression reinforced by the way the best scores at each location are distributed among the columns. There are only three stations (Santa Cruz, Tarawa, Hilo) where the rms error decreases monotonically as the number of retained EOFs is increased. This could just be a random occurrence. Or, since Baltra and Honolulu are quite close to Santa Cruz and Hilo, respectively, it is possible that their influence is captured better by the greater local structure allowed by increasing the numbers of the EOFs. (Though the favor is not reciprocated: results at Honolulu and Baltra are not always improved by increasing the number of EOFs.)

By a slight (and surely insignificant) margin the best overall performance, measured either by rms difference or correlation, was obtained for just 17 EOFs, which account for 80% of the variance in the original model run. The 17 seem to provide enough structure to capture whatever information from the 34 tide gauges our KF procedures manage to use effectively. Using fewer EOFs makes the analyzed fields smoother in space and time; smoothing out noise will raise correlations and lower rms errors. Apparently, not enough information is added by increasing the number of EOFs to do more than merely compensate for this smoothing effect. Since it represents local structure more completely, increasing the number of EOFs does bring the analysis closer to assimilated observations. However, the results of Table 2 clearly caution that this does not guarantee improvement elsewhere.

6.2. Error Estimates

A putative virtue of the Kalman filter of which we are particularly enamored is its ability to supply error estimates. An estimate of the analysis error covariance P^a at each time is readily obtained in the course of the assimilation (compare (6)). However, these estimates are theoretical and are based on the assumptions we made in modeling the data error R and system noise Q. So before examining the maps of the estimated error variance, diag(P^a), we first verify these estimates against the actual differences obtained at the observation points. There are a number of differences one might consider:

1. On the basis of the unfiltered (no assimilation) state u^u and the observations w^o we may estimate the difference covariance

$$\Delta^u = <(w^o - Hu^u)(w^o - Hu^u)^T>_t \qquad (30a)$$

as

$$\Delta^u = HP^uH^T + R \qquad (30b)$$

where P^u is the estimate of the unfiltered model error (i.e., from (13) with no data assimilation) and we have assumed that this model error and the observational error are uncorrelated.

2. Using the forecasts in an assimilation run, u^f, we may estimate the difference covariance

$$\Delta^f = <(w^o - Hu^f)(w^o - Hu^f)^T>_t \qquad (31a)$$

as

$$\Delta^f = HP^f H^T + R \qquad (31b)$$

where P^f is the estimate of the forecast error (i.e., from KF equation (13)) and we have again made use of the assumption that this model error and the observational error are uncorrelated.

3. Using the analysis in an assimilation run, u^a, we may estimate the difference covariance between the observations and the analysis at points where data are assimilated

$$\Delta^a = <(w^o - Hu^a)(w^o - Hu^a)^T>_t \qquad (32a)$$

after some manipulation as

$$\Delta^a = R - HP^a H^T \qquad (32b)$$

where P^a is the estimate of the analysis error derived from the KF formalism (i.e., (14)). The formula (32b) accounts for the fact that the same observational error affects both the observation and the analysis [cf. *Miller*, 1990, p. 11,466].

4. At stations where data are withheld the differences with data Δ^w would be formally the same as (32a), but the theoretical estimate would be similar to (31b) since the analysis error is independent of the observational noise at these points:

$$\Delta^w = HP^a H^T + R. \qquad (33)$$

All these Δ are full difference covariance matrices, but we will only look at the square roots of their diagonal elements, the rms differences between the observations and the model-simulated variables. Since all the formulas above are written in terms of the reduced state space variables, consistency requires that the direct comparisons ((30a), (31a), (32a)) use the projection of the model fields into the reduced space of EOFs. Thus the observational error R used in the theoretical estimates ((30b), (31b), (32b), (33)) should account for the

Figure 7. Rms differences with observations as a function of distance from the equator. The KF run used 17 EOFs, and all results are truncated to 17 EOFs. Actual values are indicated by open circles, theoretical estimates are indicated by crosses. (a) Unfiltered run. (b) Analysis at stations with data withheld.


Figure 6. Actual and estimated differences between sea level observations and model output from unfiltered and KF runs. Differences are defined by (30)-(33). The Kalman filter uses 93 EOFs, and the output is reduced to 93 EOFs.
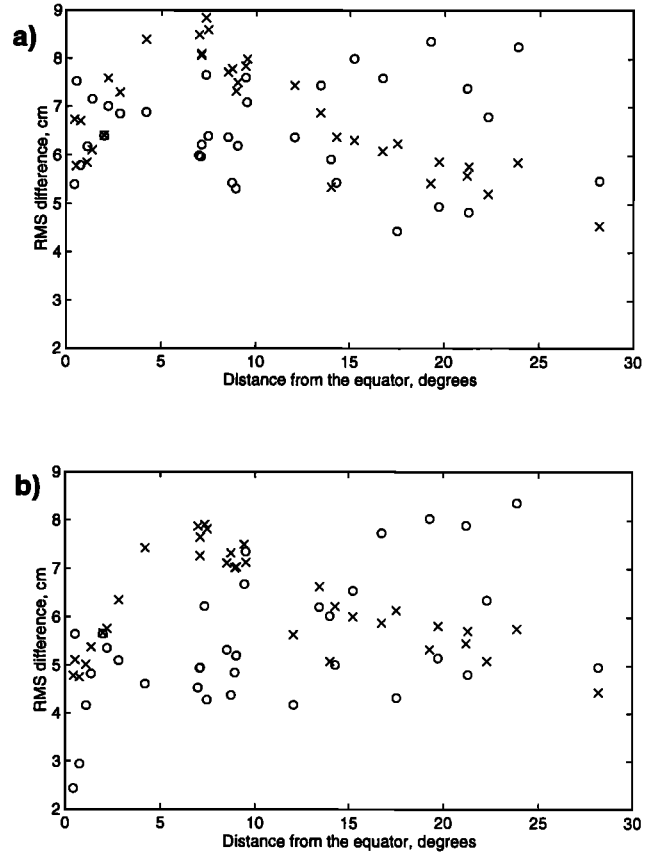
additional sampling error, i.e., the lessened ability of the true EOFs to represent the observations compared with the true full state space. Recall that this is one possible interpretation for R' (compare the discussion at the end of section 3).

Actual and estimated errors for the 93 EOF case are shown in Figure 6. (Instead of the mean over the run, the estimates at the end of the run are shown. This was done because we had not saved all estimates from all 34 runs with data withheld. In all the cases we have examined the estimates converge quickly, as expected, so the final estimate is very close to the mean.) In general, we judge the estimates to be fair to good predictors of the actual error. In most cases the estimate predicts the actual values to within a centimeter or two. While some scatter is to be expected, there are a number of points where the estimated and actual differences are quite different. Among them are several from the most telling comparison, the case of withheld data.

The mean of the estimates for the unfiltered case and the assimilation compared to withheld data are within a few percent of the actual means, while the mean estimates of the forecast and the case with all data assimilated are each about 17% higher than the actual means. Agreement in the mean is more or less guaranteed for the unfiltered case by the procedure used to choose the error model parameters of (23). The estimates for the KF runs are not so constrained. It is especially pleasing to see agreement with the actual values in
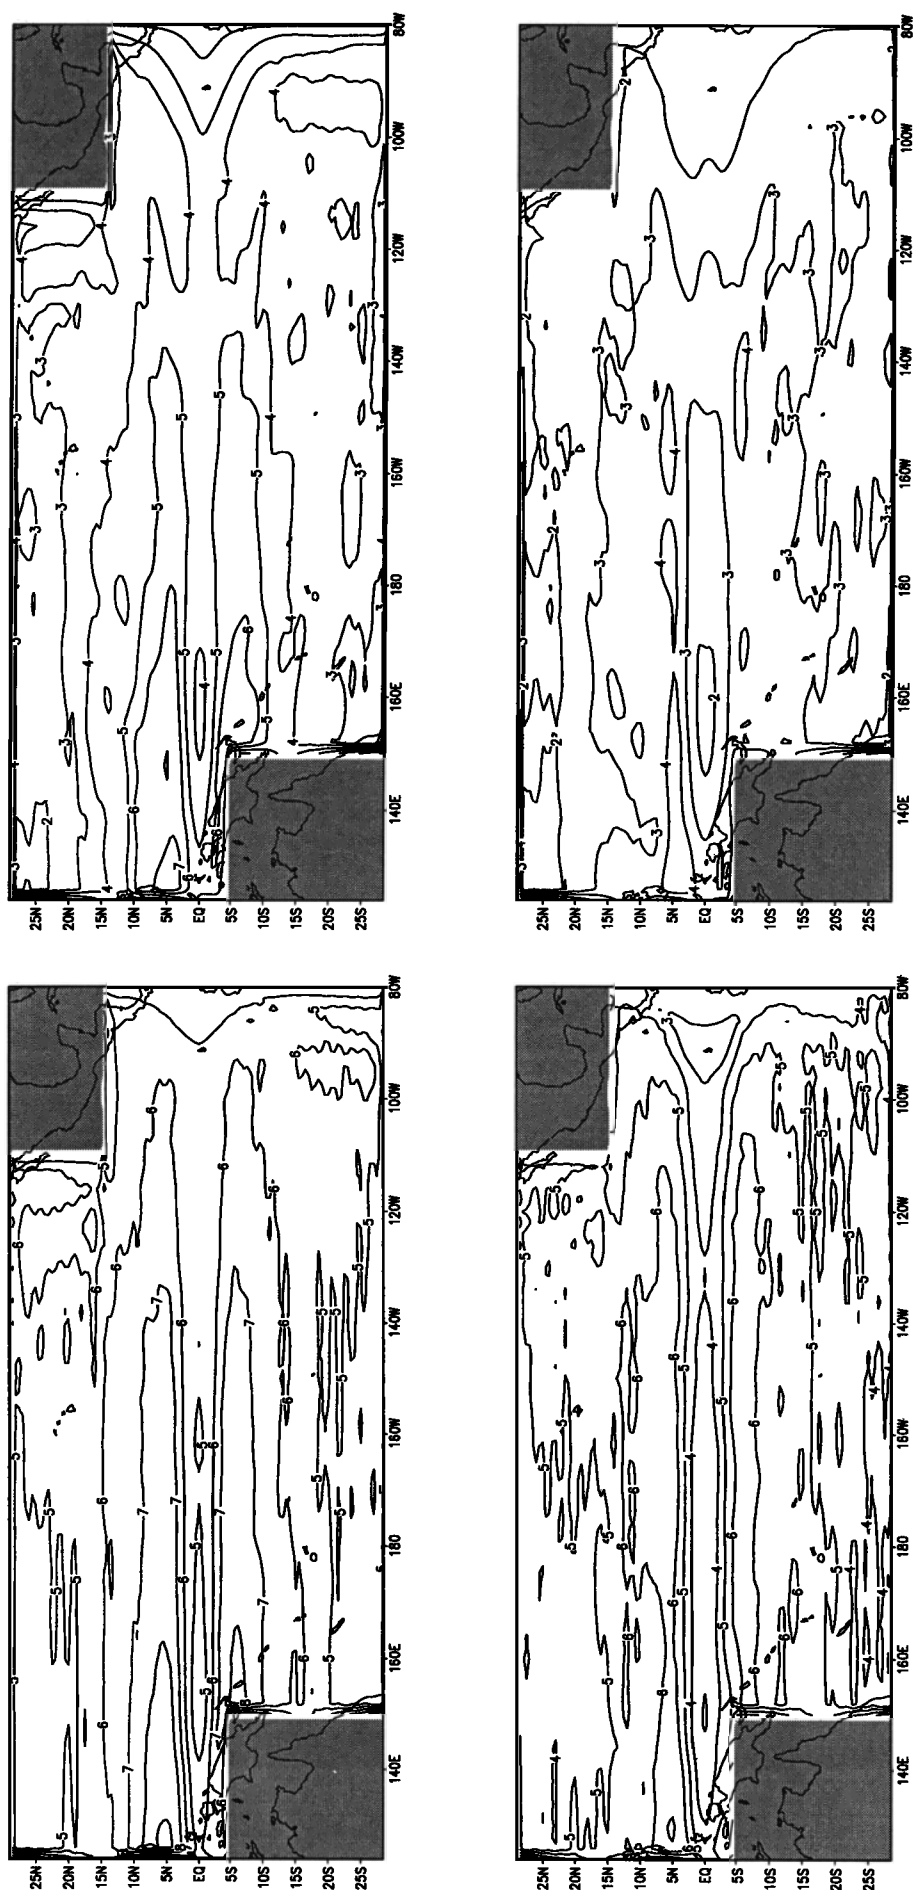
Figure 8. Maps of estimated rms errors in sea level for the 93 EOF case. (top left) Unfiltered simulation error at the grid scale (2° x 0.5°). (top right) Unfiltered simulation error field projected onto 93 EOFs. (bottom left) Analysis error at the grid scale. (bottom right) Analysis error field projected onto 93 EOFs.
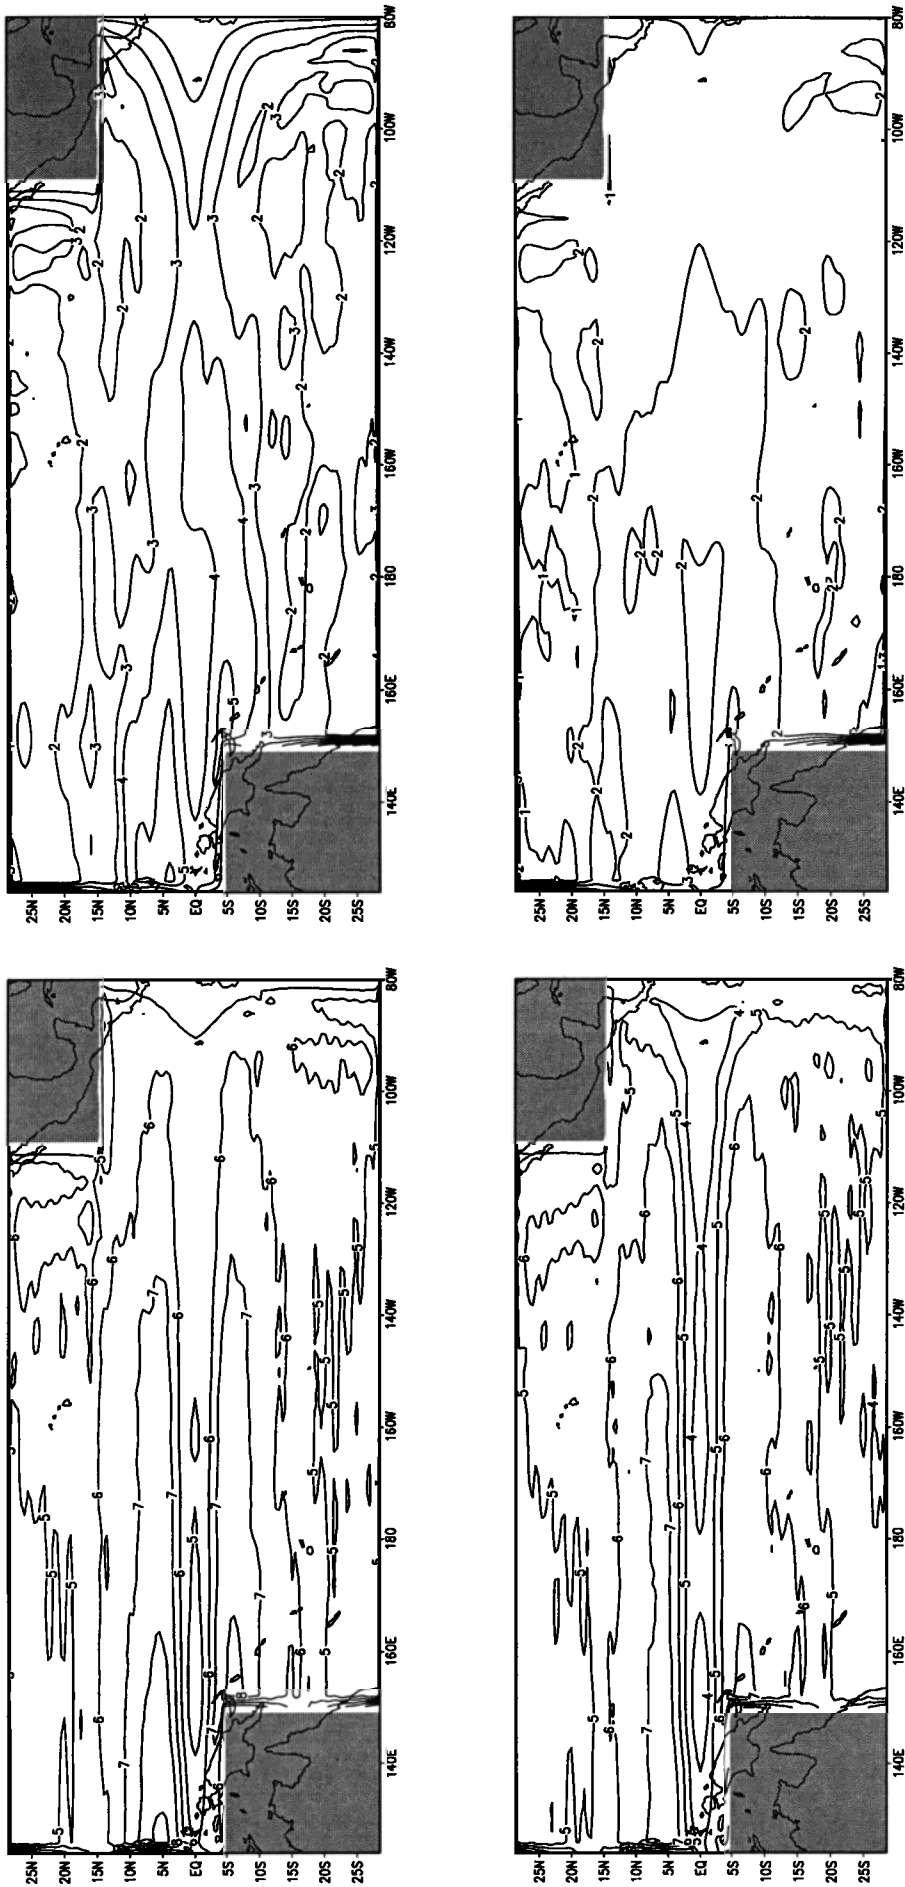
Figure 9. Maps of estimated rms errors in sea level for the 17 EOF case. (top left) Unfiltered simulation error at the grid scale (2° x 0.5°). (top right) Unfiltered simulation error field projected onto 17 EOFs. (bottom left) Analysis error at the grid scale. (bottom right) Analysis error field projected onto 17 EOFs. All maps are for December 1992.

the mean of the estimates E^w for the error in the analysis at points where data are not assimilated since this is the analysis error we most want to know. However, the estimated E^w can be correct if the analysis and observational errors are both incorrect as long as their sum is correct (see (33)). That the estimates of differences for points where data are assimilated and for forecasts are both high suggests that the estimate of R may be too high. The discrepancies are somewhat worse for the 17 EOF case and increase as the number of EOFs is reduced. The discrepancies could be reduced by lowering our estimate of R'.

One might conclude that despite the great uncertainty of our underlying noise estimates, the KF estimates have captured the main features of the analysis errors. Regardless, a significant structural problem is revealed in Figure 7, which shows the actual and expected errors as a function of distance from the equator (absolute latitude). We have shown the 17 EOF case; the others are similar. The actual errors scatter, with no systematic latitude dependence. The estimates have a definite latitudinal dependence with a maximum at about 7°. Since the observational noise was taken to be the same at all points, this discrepancy points to flaws in our model for the system noise Q.

While the above cautions against reliance on the details of the estimated error, our estimates do provide a roughly correct idea of the error magnitude, erring slightly to the high side. Maps of estimated errors are provided in Figures 8 and 9. They clearly show the improvement that the analysis provides compared to the unfiltered simulation. This improvement is more striking when one looks at the smoother fields that result by projecting onto the EOFs. The lower error for 17 EOFs reflects the greater degree of smoothing; the amplitude of the signal is lower as well. The errors at the grid scale are estimated by adding P' to P^a, where P' is the error due to the discarded modes estimated from the Monte Carlo calculation, see (19). Note that the grid scale estimates are a bit smaller for the KF with 93 EOFs than for the one with 17.

6.3. Sea Level Maps

The fruit of this work is a set of analyzed maps of sea level, combining tide gauge observations with an estimate from a wind driven ocean model. Maps for selected Decembers are shown in Figures 10-14. In each figure we show the unfiltered (no assimilation) sea level field and the analyzed fields from the reduced state space Kalman filter with 17 and 93 EOFs. Shown are three warm events (1982, 1986, 1991; Figures 10-12), one cold event (1988; Figure 13), and one non event (1990; Figure 14). Generally, the two KF analyses are similar at large scale, with the 93 EOF analysis providing a more detailed, less smooth field. In two of the warm events the assimilation makes an obvious systematic difference, moving the maximum from the coast in the unfiltered field to the central equatorial Pacific. This change suggests a systematic model problem, though it does not tell us whether to blame the model or the winds driving it. For the 1982 warm event, the available data are too sparse to have such a strong influence, although it is also possible that this very strong event truly had a different structure. The unfiltered version of the cold event (1988) is closer to the analyses, though centered slightly farther east. In the more or less normal year of 1990 the analysis again moves the equatorial maximum to the west. It also increases its amplitudes. Changes elsewhere are small.


Figure 10. Maps of sea level for December 1982: (top) unfiltered (no assimilation); (middle) analysis from the Kalman filter with 17 EOFs; (bottom) analysis from the Kalman filter with 93 EOFs.

7. Discussion

The results we have presented above indicate some success in the primary goal of this work: we have combined the available tide gauge data with a wind-driven ocean model to produce analyzed maps of sea level in the tropical Pacific. The maps improve on what the model alone could generate. We were able to demonstrate the extent of this improvement by comparisons with data withheld from the assimilations. The comparisons show substantial (≈ 1 cm rms) improvement on average, with the most consistent positive results in the region within 15° of the equator. There is also theoretical reason to believe that our maps are superior to those produced from the sparse array of tide gauges alone (e.g., the maps produced by the Sea Level Data Center at the University of Hawaii). Our assimilation procedure ensures plausible time continuity, and the model at the core of the procedure enforces physically motivated dynamical constraints on the analyzed fields.

We have demonstrated the reduced state space Kalman filter as a feasible data assimilation procedure in a real oceanographic problem. The Kalman filter is generally acknowledged as a potentially optimal data assimilation procedure, but one that is prohibitively costly because it calls
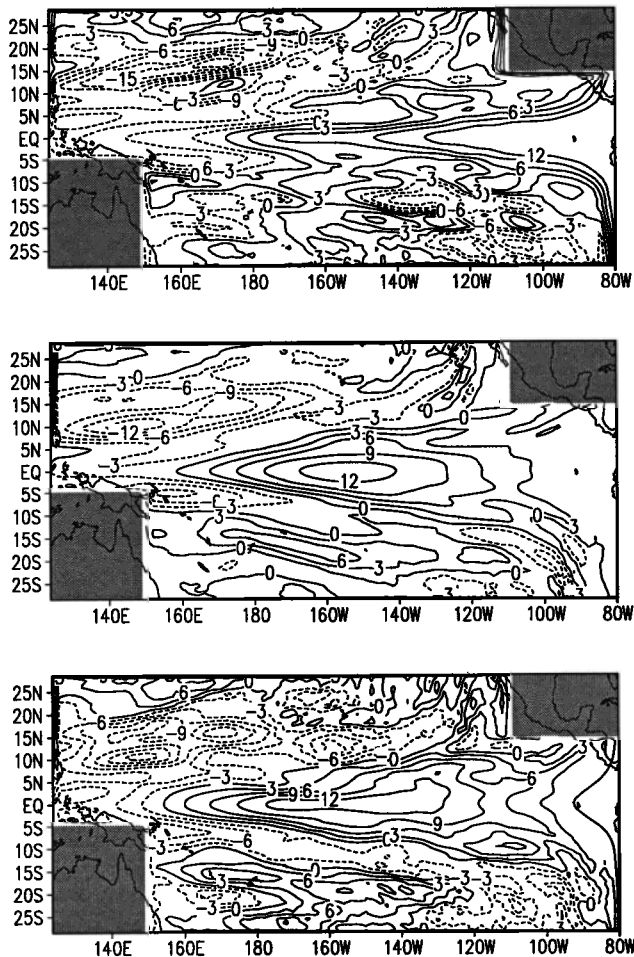
Figure 11. Maps of sea level for December 1986: (top) unfiltered (no assimilation); (middle) analysis from the Kalman filter with 17 EOFs; (bottom) analysis from the Kalman filter with 93 EOFs.

for the quite imperfect simulation models available for large-scale oceanography. Even if this were not true, it is hard to conceive of a real oceanographic problem where we know enough to fill the requisite system and observational noise matrices. Even if we did, the length of oceanographic assimilations is so short that detailed error models are unlikely to help much. This is a consequence of the slow convergence of the error covariance matrices with sample size, which makes it unlikely that even the correct long-term error statistics will match in any detail the statistics for the relatively brief periods we work with.

Finally, we argued that the detail may actually be harmful, because it encourages the assimilation to locally overfit the too sparse data. The deleterious effects of too much structure have been pointed out by many authors of data assimilation studies. *Bennett and Budgell* [1987] showed that the Kalman filter will create unrealistic local features if the noise model has too much power at small scales. The review by *Busalacchi* [1996] shows the need for smoothness constraints to be a recurrent theme in the application of variational (adjoint) methods to tropical oceanography [cf. *Long and Thacker*, 1989a, b; *Moore et al.*, 1987; *Sheinbaum and Anderson*, 1990a, b]. In addition to dissatisfaction with the structure of the solutions obtained, *Long and Thacker* [1989b] found that
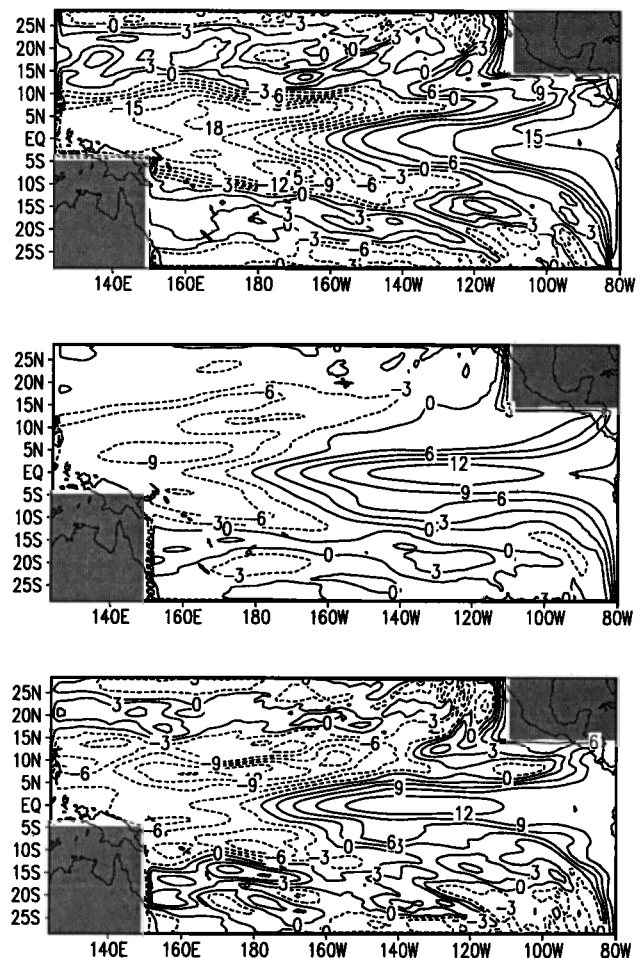
for the calculation of error covariance matrices whose dimension is the square of the size of the model state space.

Our approach is to reduce the size of the space used for the covariance matrices needed to compute the Kalman gain. The work reported here required literally hundreds of experiments, each assimilating data once a month for 18 years. It is perfectly feasible to run them on a workstation in a week's time. It is true that while the model we used had a fairly high resolution (2° x 0.5°), it is far simpler and faster than an ocean GCM. Nonetheless, there is nothing about our procedure that does not carry over to models of arbitrary complexity. The limiting cost factor would be the ability to run the model itself, not the KF procedure. While the method is quite general, it is true that the tropical oceans are an especially favorable environment since, as is well known, variability there is dominated by a relatively few large-scale structures. However, we believe the approach will prove broadly applicable in oceanography. An encouraging example is FMR's application of a similar method to an idealized midlatitude jet.

We argued that the reduction in the size of the filter matrices is driven by more compelling considerations than computational cost. *Dee* [1995] points out that the formal assumptions of Kalman filter theory almost surely do not hold

Figure 12. Maps of sea level for December 1991: (top) unfiltered (no assimilation); (middle) analysis from the Kalman filter with 17 EOFs; (bottom) analysis from the Kalman filter with 93 EOFs.
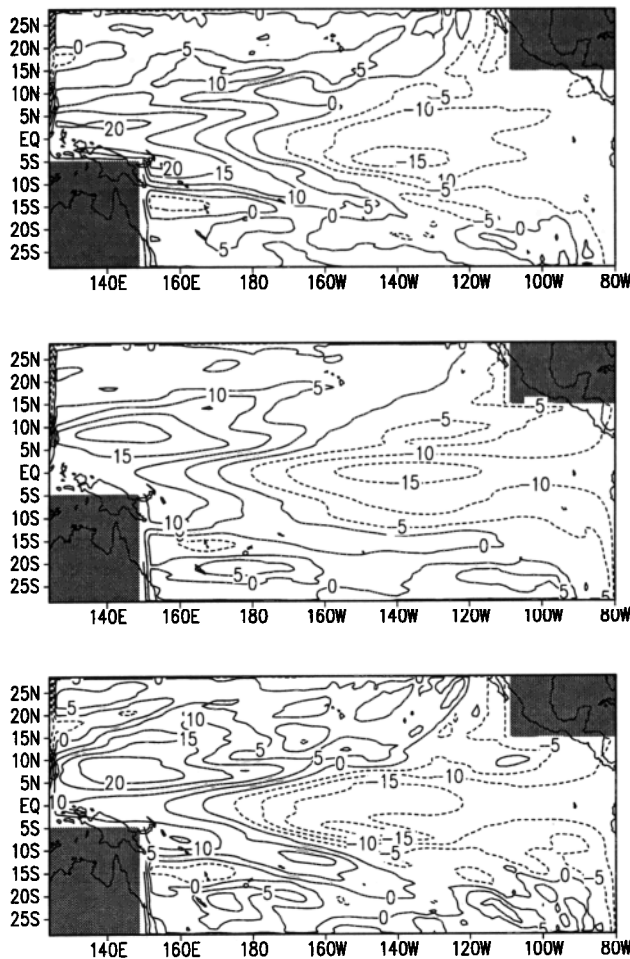
Figure 13. Maps of sea level for December 1988: (top) unfiltered (no assimilation); (middle) analysis from the Kalman filter with 17 EOFs; (bottom) analysis from the Kalman filter with 93 EOFs.

influence of the data (as with *Bennett's* [1992] representors), but that has not been done here. The issue may be investigated further by choosing a problem with far more data to assimilate. *Reverdin et al.* [1996] have used the same machinery to assimilate 86,000 XBTs and bathythermographs in the same tropical region. Results with 17 modes were noticeably better than with fewer modes, lending modest support to this explanation. A planned assimilation of altimetry data should shed added light on this issue.

That the reduced state space Kalman filter did as well as the grid point one supports our argument that our very limited knowledge of the error structure makes it pointless to use the full state space in the filter. Unfortunately, the results here go beyond what we desire: apparently, the error models we use (and perhaps the dynamic model as well) are so poor that going beyond a very few degrees of freedom extracts little additional information from the data.

Attention is best directed to the model and error estimates, not to expanding the filter. The reduced state space approach is helpful here, because it makes it feasible to perform many assimilation experiments. It also helps by reducing the number of parameters that must be determined to fit the error model. Adaptive approaches [e.g., *Dee et al.*, 1985] become realistic possibilities. Earlier, we discussed some diagnostics
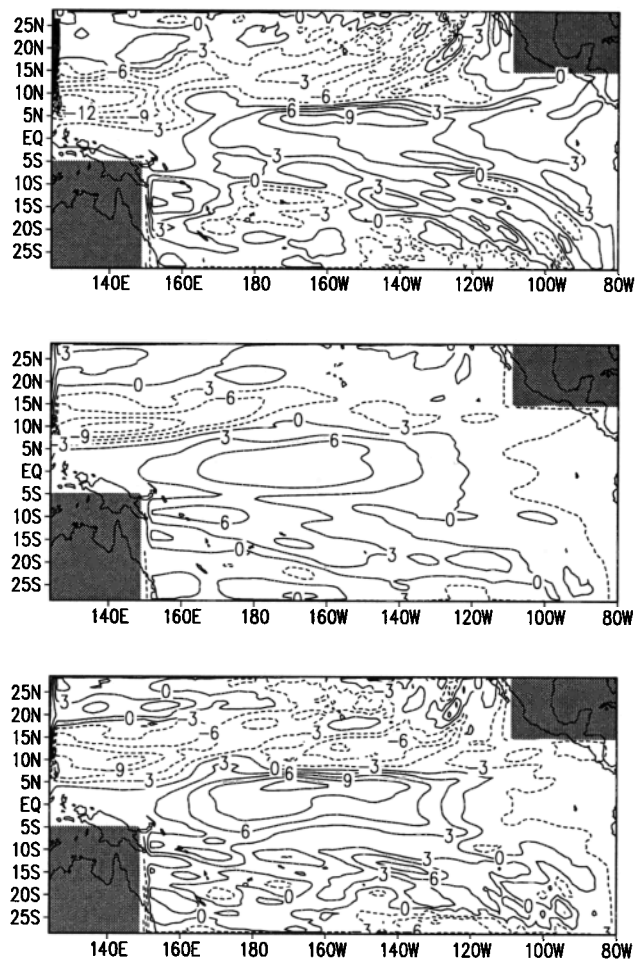

adding structure made the search for the optimal solution poorly conditioned. Ameliorating tactics tried in the variational context include adding a curvature penalty to the cost function, weighting the large scales of the prior estimates, and restricting the iteration count. The approach we take here is built on acceptance of the inherent impossibility of accurately reconstructing details of the ocean state from the limited data available.

In section 5 we demonstrated that the reduced state space KF produced as good an analysis as a full state space filter, one that kept a row and column for every variable at every model grid point. The more extensive study reported in section 6 demonstrated our thesis to the point of embarrassment: reducing the state space to only five modes gave results barely distinguishable from those obtained with 93. A priori, we had expected that tens of modes would suffice, but that many tens would also be necessary, and that the most elaborate oceanographic problem now imaginable might require hundreds. On the basis of the assimilations reported here, it would be hard to argue that more than 17 modes are needed.

A possible explanation is that with only 34 data points to assimilate, only a handful of structures are required by the filter gain. No doubt there is some truth to this. It would surely hold if the structures had been chosen to capture efficiently the

Figure 14. Maps of sea level for December 1990: (top) unfiltered (no assimilation); (middle) analysis from the Kalman filter with 17 EOFs; (bottom) analysis from the Kalman filter with 93 EOFs.

that revealed shortcomings of our error models (e.g., Figure 6) and some possible fixes (e.g., reducing **R'**). It is beyond the scope of this paper, but the subject of tuning error estimates merits a comprehensive study. Though it cannot automatically correct systematic model biases, the Kalman filter framework is useful for analyzing model errors, an essential step toward model improvement.

One of the oft cited virtues of the Kalman filter is that it produces estimates of analysis errors. In this work we were able to compare the theoretical error estimates produced by the KF machinery against actual errors. Agreement did not extend to the details; only the gross features matched. FMR also reported some discrepancies between error estimates and actual error. Discrepancies ought to be expected given that the KF estimates of analysis errors rely on specifications of poorly known observational and model system noise covariances. Our experience raises the question of how far one should trust aspects of the theoretical estimates that have not been verified in some way. This has to be done carefully, since it is likely that obvious discrepancies between theoretical and actual errors will be eliminated by being "used up" to tune specifications of the noise structure.

Our approach to the Kalman filter opens up a number of avenues that should be pursued. Learning how to use assimilations to improve the specification of error models is at the top of the list [cf. Chan et al., 1996]. In the near future we plan to apply this methodology to the coupled ocean atmosphere El Niño-Southern Oscillation forecasting model of Cane et al. [1986]. We would like to demonstrate the feasibility of using it with an ocean general circulation model, especially to assimilate altimetric sea level data. As with tide gauge measurements, altimetry provides no direct information about subsurface thermal structure. Data assimilation experiments have shown the value of surface observations for models with substantial vertical structure to be limited [e.g., Long and Thacker, 1989b] without the use of a priori empirical information to connect the surface and subsurface [e.g., Ezer and Mellor, 1992]. Our procedure is intrinsically multivariate; furthermore, it automatically fills in the horizontal gaps between orbital tracks. Thus it seems well suited to this task.

References

Bennett, A. F., *Inverse Methods in Physical Oceanography*, 346 pp., Cambridge Univ. Press, New York, 1992.

Bennett, A. F., and W. P. Budgell, Ocean data assimilation and the Kalman filter: Spatial regularity, *J. Phys. Oceanogr.*, *17*, 1584-1601, 1987.

Blumenthal, M. B., Predictability of a coupled ocean-atmosphere model, *J. Clim.*, *4*, 766-784, 1991.

Busalacchi, A. J., Data assimilation in support of tropical ocean circulation studies, in *Modern Approaches to Data Assimilation in Ocean Modeling*, edited by P. Malanotte-Rizzoli, pp. 235-270, Elsevier Sci., New York, 1996.

Busalacchi, A. J., and M. A. Cane, Hindcasts of sea level variations during 1982/83 El Niño, *J. Phys. Oceanogr.*, *15*, 213-221, 1985.

Busalacchi, A. J., and J. J. O'Brien, Interannual variability of the equatorial Pacific in the 1960s, *J. Geophys. Res.*, *86*, 10,901-10,907, 1981.

Cane, M. A., Modeling sea level during El Niño, *J. Phys. Oceanogr.*, *14*, 586-606, 1984.

Cane, M. A., Forecasting El Niño with a geophysical model, in *Teleconnections Linking Worldwide Climate Anomalies*, edited by M. H. Glantz, R. W. Katz, and N. Nicholls, chap. 11, pp. 345-369, Cambridge Univ. Press, New York, 1991.

Cane, M. A., and R. J. Patton, A numerical model for low-frequency equatorial dynamics, *J. Phys. Oceanogr.*, *14*, 1853-1863, 1984.

Cane, M.A., S. E. Zebiak, and S. C. Dolan, Experimental forecasts of El Niño, *Nature, 321*, 827-832, 1986.

Chan, N-H, J. B. Kadane, R. N. Miller, and W. Palma, Estimation of tropical sea level anomaly by an improved Kalman filter, *J. Phys. Oceanogr.*, *26*, 1286-1303, 1996.

Cohn, S. E., and D. F. Parrish, The behavior of forecast error covariances for a Kalman filter in two dimensions, *Mon. Weather Rev.*, *119*, 1757-1785, 1991.

Dee, D. P., Simplification of the Kalman filter for meteorological data assimilation, *Q. J. R. Meteorol. Soc.*, *117*, 365-384, 1991.

Dee, D. P., On-line estimation of error covariance parameters for atmospheric data assimilation, *Mon. Weather Rev.*, *123*, 1128-1145, 1995.

Dee, D. P., S. E. Cohn, A. Dalcher, and M. Ghil, An efficient algorithm for estimating noise covariance in distributed systems, *IEEE Trans. Autom. Control, AC 30*, 1057-1065, 1985.

Ezer, T., and G. L. Mellor, A numerical study of the variability and the separation of the Gulf Stream induced by surface atmospheric forcing and lateral boundary flows, *J. Phys. Oceanogr.*, *22*, 660-682, 1992.

Fu, L-L, I. Fukumori, and R. N. Miller, Fitting dynamic models to the Geosat sea level observations in the tropical Pacific Ocean, II, A linear, wind-driven model, *J. Phys. Oceanogr.*, *23*, 2162-2181, 1993.

Fukumori, I., and P. Malanotte-Rizzoli, An approximate Kalman filter for ocean data assimilation: An example with an idealized Gulf Stream mode, *J. Geophys. Res.*, *100*, 6777-6793, 1995.

Fukumori, I., J. Benveniste, C. Wunsch, and D. B. Haidvogel, Assimilation of sea surface topography into an ocean circulation model using a steady-state smoother, *J. Phys. Oceanogr.*, *23*, 1831-1855, 1993.

Ghil, M., Meteorological data assimilation for oceanographers, I, Description and theoretical framework, *Dyn. Atmos. Oceans*, *13*, 171-218, 1989.

Ghil, M., and P. Malanotte-Rizzoli, Data assimilation in meteorology and oceanography, *Adv. Geophys.*, *33*, 141-266, 1991.

Ghil, M., S. E. Cohn, J. Tavantzis, K. Bube, and E. Isaacson, Applications of estimation theory to numerical weather prediction, in *Dynamic Meteorology: Data Assimilation Methods*, edited by L. Bengtsson, M. Ghil, and E. Kallen, pp. 139-224, Springer-Verlag, New York, 1981.

Goldenberg, S., and J. J. O'Brien, Time and space variability of tropical wind stress, *Mon. Weather Rev.*, *109*, 1190-1205, 1981.

Gourdeau, L., S. Arnault, Y. Meynard, and J. Merle, Geosat sea-level assimilation in a tropical Atlantic model using Kalman filter, *Oceanol. Acta, 15*, 567-574, 1992.

Hernandez, L. L. M., and J. A. Calderon, Analysis of dynamic data assimilation for atmospheric phenomena. Effect of the model order, *Atmosfera, 4*, 145-164, 1991.

Ji, M., and T. M. Smith, Ocean model response to temperature data assimilation and varying surface stress: Intercomparisons and Implications for climate forecast, *Mon. Weather Rev.*, *123*, 1811-1821, 1995.

Jiang, S,. and M. Ghil, Dynamical properties of error statistics in a shallow-water model, *J. Phys. Oceanogr*, *23*, 2541-2566, 1993.

Long, R. B., and W. C. Thacker, Data assimilation into a numerical equatorial ocean model, I, The model and the assimilation algorithm, *Dyn. Atmos. Oceans, 13*, 379-412, 1989a.

Long, R. B., and W. C. Thacker, Data assimilation into a numerical equatorial ocean model, II, Assimilation experiments, *Dyn. Atmos. Oceans, 13*, 413-440, 1989b.

Lorenz, E. N., Empirical orthogonal functions and statistical weather prediction, *Sci. Rep. 1, Stat. Forecasting Proj. Mass. Inst. of Technol.*, Cambridge, 1956.

Miller, R. N., Tropical data assimilation experiments with simulated data: The impact of the tropical ocean and global atmosphere thermal array for the ocean, *J. Geophys. Res.*, *95*, 11,461-11,482, 1990.

Miller, R. N., and M. A. Cane, A Kalman filter analysis of sea level heights in the tropical Pacific, *J. Phys. Oceanogr.*, *19*, 773-790, 1989.

Miller, R. N., A. J. Busalacchi, and E. C. Hackert, Sea surface

topography fields of the tropical Pacific from data assimilation, *J. Geophys. Res.*, *100*, 13,389-13,425, 1995.

Mitchum, G. T., Comparison of TOPEX sea surface heights and tide gauge sea levels, *J. Geophys. Res.*, *99*, 24,541-24,553, 1994.

Moore, A. M., N. S. Cooper, and D. L. T. Anderson, Initialization and data assimilation in models of the Indian Ocean, *J. Phys. Oceanogr.*, *17*, 1965-1977, 1987.

Parrish, D. F. and S. E. Cohn, A Kalman filter for a two-dimensional shallow-water model: Formulation of preliminary experiments, *Off. Note 304*, 64 pp., Natl. Meteorol. Cent., Washington D. C., 1985.

Provost, C., and R. Salmon, A variational method for inverting hydrographic data,. *J. Mar. Res.*, *44*, 1-34, 1986.

Reverdin, G., A. Kaplan, and M. A. Cane, Simulations of XBT data in the tropical Pacific, *J. Geophys. Res.*, in press, 1996.

Rosati, A., K. Miyakoda, and R. Gudgel, The impact of ocean initial conditions on ENSO forecasting with a coupled model, *Mon. Weather Rev.*, in press, 1996

Sasaki, Y., Numerical variational analysis formulated under the constraints as determined by long wave equations and a low-pass filter, *Mon. Weather Rev.*, *98*, 884-898, 1970.

Sheinbaum, J., and D. L. T. Anderson, Variational assimilation of XBT data, I, *J. Phys. Oceanogr.*, *20*, 672-688, 1990a.

Sheinbaum, J., and D. L. T. Anderson, Variational assimilation of XBT data, II, Sensitivity studies and use of smoothing constraints, *J. Phys. Oceanogr.*, *20*, 689-704, 1990b.

Wyrtki, K., Teleconnections in the equatorial Pacific Ocean, *Science*, *180*, 66-68, 1973.

Wyrtki, K., El Niño - The dynamic response of the equatorial Pacific Ocean to atmospheric forcing, *J. Phys. Oceanogr.*, *5*, 572-584, 1975.

Xue, Y., M. A. Cane, S. E. Zebiak, and M. B. Blumenthal, On the prediction of ENSO: A study with a low-order Markov mode, *Tellus, Ser. 46A*, 512-528, 1994.

A. J. Busalacchi, NASA Goddard Space Flight Center, Laboratory for Hydrospheric Processes, Greenbelt, MD 20771.

M. A. Cane and A. Kaplan, Lamont-Doherty Earth Observatory of Columbia University, Palisades, NY 10964.

E. C. Hackert, Hughes STX Corporation, Lanham, MD 20706.

R. N. Miller, College of Oceanic and Atmospheric Sciences, Oregon State University, Corvallis, OR 97331.

B. Tang, Centre for Earth and Ocean Sciences, University of Victoria, P.O. Box 1700, Victoria, British Columbia V8W 262, Canada.