

Initialization of an ENSO Forecast System Using a Parallelized Ensemble Filter

S. ZHANG, M. J. HARRISON, A. T. WITTENBERG, AND A. ROSATI

Geophysical Fluid Dynamics Laboratory, Princeton University, Princeton, New Jersey

J. L. ANDERSON

NCAR, Boulder, Colorado

V. BALAJI

Geophysical Fluid Dynamics Laboratory, Princeton University, Princeton, New Jersey

(Manuscript received 16 March 2004, in final form 21 March 2005)

ABSTRACT

As a first step toward coupled ocean–atmosphere data assimilation, a parallelized ensemble filter is implemented in a new stochastic hybrid coupled model. The model consists of a global version of the GFDL Modular Ocean Model Version 4 (MOM4), coupled to a statistical atmosphere based on a regression of National Centers for Environmental Prediction (NCEP) reanalysis surface wind stress, heat, and water flux anomalies onto analyzed tropical Pacific SST anomalies from 1979 to 2002. The residual part of the NCEP fluxes not captured by the regression is then treated as stochastic forcing, with different ensemble members feeling the residual fluxes from different years. The model provides a convenient test bed for coupled data assimilation, as well as a prototype for representing uncertainties in the surface forcing.

A parallel ensemble adjustment Kalman filter (EAKF) has been designed and implemented in the hybrid model, using a local least squares framework. Comparison experiments demonstrate that the massively parallel processing EAKF (MPPEAKF) produces assimilation results with essentially the same quality as a global sequential analysis. Observed subsurface temperature profiles from expendable bathythermographs (XBTs), Tropical Atmosphere Ocean (TAO) buoys, and Argo floats, along with analyzed SSTs from NCEP, are assimilated into the hybrid model over 1980–2002 using the MPPEAKF. The filtered ensemble of SSTs, ocean heat contents, and thermal structures converge well to the observations, in spite of the imposed stochastic forcings. Several facets of the EAKF algorithm used here have been designed to facilitate comparison to a traditional three-dimensional variational data assimilation (3DVAR) algorithm, for instance, the use of a univariate filter in which observations of temperature only directly impact temperature state variables. Despite these choices that may limit the power of the EAKF, the MPPEAKF solution appears to improve upon an earlier 3DVAR solution, producing a smoother, more physically reasonable analysis that better fits the observational data and produces, to some degree, a self-consistent estimate of analysis uncertainties. Hybrid model ENSO forecasts initialized from the MPPEAKF ensemble mean also appear to outperform those initialized from the 3DVAR analysis. This improvement stems from the EAKF's utilization of anisotropic background error covariances that may vary in time.

1. Introduction

The El Niño–Southern Oscillation (ENSO) over the tropical Pacific has been recognized as the earth's dominant climate fluctuation on interannual time scales (Rasmusson and Wallace 1983; Glantz et al. 1991). Un-

derstanding ENSO is a key to understanding the global climate anomaly. Theories exist to explain the ENSO phenomenon, in which the converging part attributes ENSO to the dynamic coupling between the atmosphere and ocean in the equatorial Pacific region (Zebiak and Cane 1987; Suarez and Schopf 1988; Neelin 1991; Sun and Liu 1996; Neelin et al. 1998; Fedorov and Philander 2001). The dynamic coupling refers to a positive feedback loop between surface wind stress, sea surface temperature (SST), and ocean upwelling. Over the

Corresponding author address: Dr. S. Zhang, NOAA/GFDL, Princeton University, P.O. Box 308, Princeton, NJ 08542.
E-mail: snz@gfdl.noaa.gov

tropical Pacific Ocean, the surface winds are driven by SST gradients (Lindzen and Nigam 1987) and changes in SST gradients affect the strength of surface winds. Since upwelling is driven by the surface winds, changes in the strength of the surface winds affect the strength of upwelling, which in return affects the SST distribution.

Modeling the ENSO mechanism of air–sea interaction contains many uncertainties since the coupled feedback process makes the climate system highly sensitive to errors in simulations, such as inaccurate parameterizations of clouds, radiation, convection, and turbulent mixing, etc. Therefore, viewing the modeled ocean–atmosphere coupling system as a continuous stochastic dynamical process may be more appropriate than a single deterministic process, in which the climate evolution is described by the vector stochastic differential equation (Jazwinski 1970),

$$d\mathbf{x}_t/dt = \mathbf{f}(\mathbf{x}_t, t) + \mathbf{G}(\mathbf{x}_t, t)\mathbf{w}_t. \quad (1)$$

Here, \mathbf{x}_t is an n -dimensional vector representing the coupled model state at time t (n is the size of the model state), \mathbf{f} is an n -dimensional vector function, \mathbf{w}_t is a white Gaussian process (uncorrelated in time) of dimension r with mean 0 and covariance matrix $\mathbf{S}(t)$, while \mathbf{G} is an $n \times r$ matrix. The first and second terms of the right-hand side in Eq. (1) respectively represent the deterministic modeling and uncertainties in a coupled system.

In ocean modeling, surface temperatures are typically damped toward the analyzed SSTs, and estimated fluxes of momentum, heat, and water are applied as the surface forcings. Unfortunately, restoring SST may only change the top layer structure, rather than building up the whole vertical thermal structure, and the estimated wind stress, heat flux, and water flux have errors (Wittenberg 2004). Incorporating these inaccurate surface forcings into a biased model cannot validly prevent the drift of the modeled ocean state from climatology. In addition, in eddy-resolving models, ocean data assimilation is expected to introduce mesoscale eddies and nonlinear dynamical features (Ezer and Mellor 1994), which are inherently unpredictable in nonassimilated models. Therefore ocean modeling needs ocean data assimilation (ODA), which reconstructs the historical series of the ocean evolution using model dynamics to extract information from all observations available. An ODA procedure attempts to produce consistent ocean states that serve as initial conditions for model forecasts. On the other hand, with diagnostics, the ODA reconstructed historical series of the ocean states with three-dimensional structure aids further understanding

of dynamics and physics of ocean evolution and may improve ocean modeling.

The traditional ODA methods that include the three-dimensional variational data assimilation (3DVAR) approach (Derber and Rosati 1989) and the four-dimensional variational data assimilation approach (4DVAR; Galanti et al. 2003; Weaver et al. 2003) solve a single estimate of ocean state by minimizing a defined distance measurement between the analysis and observations (3DVAR), or between the modeled and observational trajectories (4DVAR). In these traditional approaches, the prior specified background error covariance is usually flow independent and time invariant, and therefore may be unable to properly describe the uncertainties referred to in the second term in Eq. (1).

An ensemble filter uses finite samples to estimate the probability density function (PDF) of the system state, solving the data assimilation problem by computing the product of modeled and observational PDFs. The background error covariance between state variables is directly derived from the model dynamics, using a Monte Carlo approach. The error covariances are therefore flow dependent and time varying (Zhang and Anderson 2003). This aspect of the ensemble filter is well suited to the tropical Pacific Ocean, where flow structures are highly anisotropic and strongly dependent on the seasonal cycle and interannual (ENSO) fluctuations.

At the current state of the art, ensemble filters assume consistency of the prior state PDF (estimated by Monte Carlo samples of the model) and the real-world PDF. Under this framework, ocean data assimilation is in many ways a very different problem than atmospheric data assimilation. Whereas the atmosphere is highly sensitive to initial conditions (slightly different atmospheric states can be expected to spread out from one another very strongly after just a few days or weeks), the ocean tends to be more stable and to evolve more slowly. Model biases can therefore emerge as a strong source of error in an ocean assimilation. Moreover, because ocean observations are generally sparse and irregular in space and time, ocean model biases can grow to significant amplitude in data-void regions. Outside strongly eddying zones, tropical upper-ocean variability is driven primarily by interactions with the atmosphere. Where the effects of the surface forcing are less intense (e.g., in the deep ocean or away from the equator), oceanic variability tends to be quite weak.

It is the combination of these aspects—weak ensemble spread, sparse data, and strong model biases—that make the ocean problem a challenge for an ensemble filter. Our approach attempts to deal with each of these problems. To enhance the ensemble spread

and better sample the covariance structure of the ocean model, the ocean model is coupled to a stochastic atmosphere. This provides a prototype system for 1) representing the uncertainty of the atmospheric forcing, and 2) truly coupled ocean–atmosphere data assimilation, in the limiting case where no atmospheric data are assimilated. Following the precedent of 3DVAR ODA, observations are allowed to impact state variables over a time window that includes a number of model time steps before and after the time of the observation. Under certain conditions, this is believed to lead to assimilations that are smoother in time and reduce the magnitude of undesirable shocks generated by sparse observations. Further work may be necessary to incorporate this feature into the Bayesian theoretical framework of ensemble filters.

The current version of the GFDL data assimilation system is based on a 3DVAR scheme. Similar systems have been used in experimental climate predictions at the Geophysical Fluid Dynamics Laboratory (GFDL; Rosati et al. 1997) and operational forecasts at the National Centers for Environmental (NCEP; Behringer et al. 1998) for over a decade. This study documents efforts to develop a next-generation system based on the ensemble adjustment Kalman filter (EAKF; Anderson 2001, 2003). The system is implemented in a prototype coupled ocean–atmosphere ENSO model, based on the GFDL Modular Ocean Model Version 4 (MOM4) coupled to a statistical atmosphere. The purpose of this study is not to produce a fully operational forecast system, nor to provide an exhaustive analysis of the assimilation state estimate. Rather the intent is a proof-of-concept for the EAKF technique applied to coupled data assimilation for initialization of seasonal-to-interannual climate forecasts. We describe the assimilation methodology, its implementation for parallel machines, and an evaluation of some key assimilation metrics, including a comparison with the current 3DVAR assimilation. We include among these assimilation metrics the skill of forecasts initialized from the assimilation solutions, since the primary intent of the assimilation will be to provide initial conditions for coupled model predictions. A more comprehensive investigation of the assimilation quality will be performed once the system is implemented in a fully coupled ocean–atmosphere GCM.

The paper is organized as follows. After a brief description of the assimilation methods (ensemble filter and 3DVAR) used in this study in section 2, section 3 describes how the hybrid coupled model sets the coupled model prototype that represents the forcing uncertainties in air–sea interaction. Section 4 presents




FIG. 1. Schematic diagram of how a sequential ensemble filter uses an observation y_o to update the ensemble. See context of section 2a(1).

the parallel design of the EAKF and discusses the impact of sequential adjustment in ensemble-based filters on parallel analysis. Section 5 examines the assimilation results, comparing to the existing 3DVAR scheme, and the forecast verification is given in section 6. Finally a summary and discussion are given in section 7.

2. Assimilation methods

a. Ensemble filter

1) SEQUENTIAL IMPLEMENTATION

A variety of ensemble filtering algorithms have been developed for atmospheric and oceanic assimilation applications. These algorithms can be understood as Monte Carlo approximations to the Bayesian filtering problem (Jazwinski 1970). As pointed out by Houtekamer and Mitchell (2001), individual scalar observations can be assimilated sequentially when the observational error distribution for each is independent. If sets of observations with correlated observational error distributions are used, as application of a singular value decomposition (Anderson 2003) continues to allow sequential assimilation for observations.

Anderson (2003) points out that the impact of an observation on the set of model state variables can also be computed sequentially as long as all state variables are updated before the forward operator for the next scalar observation is computed. In this context, an ensemble filter can be described without loss of generality by describing the impact of a single scalar observation on a single state vector element.

Figure 1 schematically illustrates how a sequential ensemble filter is implemented. In step 1, an ensemble of model states is integrated forward in time from the time of the previous set of observations, t_k , to the next time at which observations are available, t_{k+1} . In step 2, the forward observation operator, H , is applied to each model state prior estimate to obtain an ensemble prior estimate of an observed scalar quantity, the dark solid ticks in step 3. The value of the observation from the

instrument, y_o (gray tick at step 3), and the observational error distribution (gray curve superposed at step 3), which is a function of the observing system, must be combined with the prior ensemble estimate to get an improved analysis estimate. Step 4 shows that updated values (thin dark dashed on the y axis) can be associated with each of the prior ensemble estimates. An innovation, or increment, is associated with each prior ensemble estimate at the end of step 4. Finally, corresponding increments for a given model state variable are obtained by linearly regressing the observation increments onto the state variable using the prior ensemble joint distribution for the observation variable and the state variable. This impact of the scalar observation is computed for each state variable in turn. When all state variables are updated, the algorithm is repeated for the next scalar observation from time t_{k+1} . When all observations have been applied, the state is advanced forward to the next time at which observations are available.

Almost all ensemble filter algorithms that have been applied in atmospheric and oceanic applications can be described by Fig. 1. The differences between the algorithms are normally confined to the detail of step 4, computing the observation increments given a prior estimate, the observation, and the observational error distribution. Here, the EAKF (Anderson 2001) is used. The EAKF is one of a class of deterministic square root filters (Tippett et al. 2003; Bishop et al. 2001; Whitaker and Hamill 2002), all of which would be expected to give qualitatively similar results in this application. Other nondeterministic ensemble filters, such as the original ensemble Kalman filter of Evensen (1994) as corrected by Houtekamer and Mitchell (1998) would likely give qualitatively different results.

2) COVARIANCE FILTERING

The algorithm outlined in section 2a(1) assumes that the assimilating model is perfect and that the ensemble size is large enough to fit a PDF well. In practice, an assimilating model will have biases that may cause the analysis ensemble members to systematically drift away from reality. This drift tends to be greatest in those locations where observations are sparse in space and time. This can induce problems in the filter—as the observations begin to look increasingly “unlikely” under the (erroneous) assumption that they were drawn from the ensemble PDF, they are given increasingly less weight in the analysis distribution, further worsening the bias. This *filter divergence* is especially pernicious in regions where variability is small compared to model systematic error; in these regions the analysis ensemble tends to give little weight to observations that depart from the ensemble. Filter divergence can cause further problems where observations appear after being absent for awhile; the sudden shift in the ensemble solutions (in localized regions near the observational locations) can induce large gradients in the physical fields, giving rise to spuriously strong currents and numerical instabilities in the model. Small ensembles like the ones used here are even more prone to problems of this kind.

To help control these problems, we define a time window to smooth the impact of observations in time. To reduce the computational burden and enhance parallelism, and minimize sampling error, we further use a distance-dependent weighting function to constrain the spatial extent of the observational impacts (Hamill et al. 2001). The resulting weighting function takes the form (Gaspari and Cohn 1999)

$$\Omega(a, d) = \begin{cases} -\frac{1}{4}\left(\frac{d}{a}\right)^5 + \frac{1}{2}\left(\frac{d}{a}\right)^4 + \frac{5}{8}\left(\frac{d}{a}\right)^3 - \frac{5}{3}\left(\frac{d}{a}\right)^2 + 1, & 0 \leq d \leq a; \\ \frac{1}{12}\left(\frac{d}{a}\right)^5 - \frac{1}{2}\left(\frac{d}{a}\right)^4 + \frac{5}{8}\left(\frac{d}{a}\right)^3 + \frac{5}{3}\left(\frac{d}{a}\right)^2 - 5\left(\frac{d}{a}\right) + 4 - \frac{2}{3}\left(\frac{d}{a}\right)^{-1}, & a < d \leq 2a; \\ 0, & d > 2a. \end{cases}$$

Here d is either a Euclidean spatial distance (horizontal or vertical), or a time difference, between the model grid point and the observation location, and a controls the observational impact window. The horizontal a is set to be 1000 km so that the observational impact radius is 2000 km. To change the shape of the weighting function near the equator, a *cosine* factor multiplying the difference of grid point and observation

latitudes scales the horizontal weight (therefore the horizontal weight contours are ellipses). The vertical a is set to be 20 m, and each observation is only allowed to impact at most two neighboring levels (one on each side).

In theory, the information contained in individual observations would be assimilated only once; the model would be expected to correctly propagate the state PDF

in time. Unfortunately, ocean models typically show large biases and little ensemble spread, and subsurface temperature observations are sparse and infrequent. In the present case, using a too-short time window produces an unacceptable assimilation bias when and where observations are absent. To constrain the ensemble to the observations without inducing large shocks, we smooth the impact of the observations in time by setting a to 5 days by weighting the covariance using the distance of an observation from the center of the window. This value is consistent with previous three-dimensional variational ocean data assimilations (Derber and Rosati 1989; Harrison et al. 1996). Increasing the width of the time window greatly increases the assimilation cost by effectively increasing the number of observations. An additional effect of the time window is analogous to reducing the observational error associated with the observation since time windows may have overlap that observations may be repeatedly used with different weights in windows nearby. This results in an exaggerated reduction in the spread of the assimilated ensemble while possibly leading to a more aggressive fit of the observation. While the tighter fit may be an advantage when using a model with large systematic error, the reduced spread acts to reduce the impact of later observations. Based on the representation of data, the model bias, and the ensemble size, how to select an optimal time-window length will remain a topic for further research. Finally, the independent products of weighting functions for horizontal, vertical, and time are accounted as a covariance factor into assimilation computation.

Normally, EAKFs are able to applied in a multivariate fashion with an observation of any type being allowed to impact all close state variables, given that cross covariances between different physical variables can be easily estimated by the ensemble samples. The ensemble sampling size determines the accuracy of the estimated cross covariance. As the first step of efforts to implement the ensemble filter into the coupled assimilation with a relatively small ensemble size, the EAKF applied here is univariate; that is, observations of temperature are only allowed to impact temperature variables, which is expedient to compare with the existing 3DVAR system. Following up, the multivariate filtering is expected to minimize imbalances in the assimilated state since the correlative relations found in the prior ensemble state estimates are maintained to some extent in the state increments. It is expected that future implementations of an EAKF without the univariate modification would lead to more balanced assimilations that might eliminate the need for the time window

while leading to an overall improved assimilation. This will be explored in future research.

b. Brief description of GFDL 3DVAR ODA system

The original 3DVAR system was set up by Derber and Rosati (1989) and certain modifications were performed by Harrison et al. (1996). Because of the uncertainty of estimating the cross-covariance matrix and the requirement of defining an observation operator between different physical variables in the multivariate 3DVAR system, the GFDL 3DVAR still is a univariate system. Instead we are engaged to develop the ensemble filter outlines in section 2a(1), which can be naturally expanded to conduct the multivariate assimilation. For the purpose of comparison and contrast in this study, what follows provides a brief description of main characters of the GFDL 3DVAR scheme.

Taking the standard 3DVAR objective functional form, which includes a background term and an observational term (e.g., Kalnay 2003), operating on the field of the temperature correction instead of the temperature itself, the objective function minimized is

$$J = \frac{1}{2} \Delta \mathbf{T}^T \mathbf{B}^{-1} \Delta \mathbf{T} + \frac{1}{2} [H(\Delta \mathbf{T}) - \Delta \mathbf{T}_0]^T \mathbf{R}^{-1} [H(\Delta \mathbf{T}) - \Delta \mathbf{T}_0], \quad (2)$$

where $\Delta \mathbf{T}$ is an N component correction vector of temperature referred to the first guess (background), \mathbf{B} is the $N \times N$ background error covariance matrix, $\Delta \mathbf{T}_0$ is a K component difference vector between the observations and the interpolated first-guess temperature at the observation location, \mathbf{R} is the $K \times K$ observational error covariance matrix (only considering variance, R is diagonal), and H is a simple bilinear interpolation mapping operator from model space to observation space.

In minimization of the functional (2) by a preconditioned conjugate gradient algorithm (Gill et al. 1981; Navon and Legler 1987), avoiding the expensive computational cost of directly inverting the \mathbf{B} matrix, each analysis step instead approximately approaches the solution by an iterative procedure.

The key quantities in the preconditioned conjugate gradient algorithm are the gradient of the functional with respect to the correction field ($\mathbf{g} = \nabla_{|\Delta \mathbf{T}} J$) and $\mathbf{h} = \mathbf{B} \mathbf{g}$, a scaled gradient vector by the background covariance. For the functional defined in (2), the gradient is given by

$$\mathbf{g} = \mathbf{B}^{-1} \Delta \mathbf{T} + H^T \mathbf{R}^{-1} [H(\Delta \mathbf{T}) - \Delta \mathbf{T}_0]. \quad (3)$$

The procedure first starts from setting the correction field $\Delta \mathbf{T}^{(1)} = 0$, and then $\mathbf{g}^{(1)} = -H^T \mathbf{R}^{-1} \Delta \mathbf{T}_0$, and $\mathbf{h}^{(1)} =$

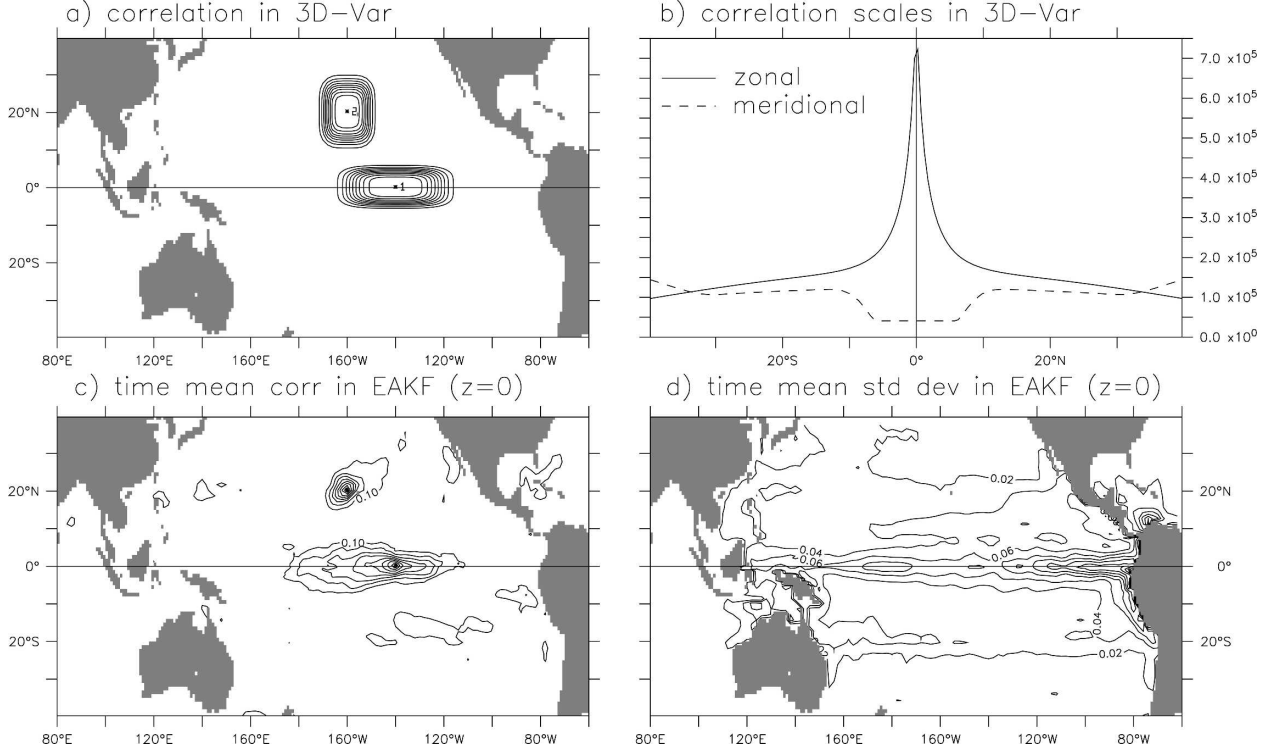


FIG. 2. The correlation structures with respect to the reference points (0°N , 140°W , marked by “1”) and (20°N , 160°W , marked by “2”) in the (a) 3DVAR and (c) EAKF (at the surface), (b) the correlation zonal (solid) and meridional (dashed) scales in the 3DVAR, and (d) the time mean standard deviation in the EAKF (at the surface). The contour intervals are 0.1 for (a) and (c) and 0.02 for (d).

$\mathbf{B}\mathbf{g}^{(1)}$. If also initializing the initial search direction [$\mathbf{d}^{(0)}$ and $\mathbf{e}^{(0)}$] and $\beta^{(1)}$ to be zero, the algorithm reaches the solution by the following iterative procedure:

$$\begin{aligned}
 \mathbf{d}^{(n)} &= -\mathbf{h}^{(n)} + \beta^{(n)}\mathbf{d}^{(n-1)}, \\
 \mathbf{e}^{(n)} &= -\mathbf{g}^{(n)} + \beta^{(n)}\mathbf{e}^{(n-1)}, \\
 \mathbf{f}^{(n)} &= \mathbf{e}^{(n)} + \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}\mathbf{d}^{(n)}, \\
 \alpha^{(n)} &= \frac{[\mathbf{g}^{(n)}]^T\mathbf{h}^{(n)}}{(\mathbf{d}^{(n)})^T\mathbf{f}^{(n)}}, \\
 \mathbf{g}^{(n+1)} &= \mathbf{g}^{(n)} + \alpha^{(n)}\mathbf{f}^{(n)} \\
 \Delta\mathbf{T}^{(n+1)} &= \Delta\mathbf{T}^{(n)} + \alpha^{(n)}\mathbf{d}^{(n)} \\
 \mathbf{h}^{(n+1)} &= \mathbf{B}\mathbf{g}^{(n+1)}, \\
 \beta^{(n+1)} &= \frac{[\mathbf{g}^{(n+1)}]^T\mathbf{h}^{(n+1)}}{[\mathbf{g}^{(n)}]^T\mathbf{h}^{(n)}}, \quad (4)
 \end{aligned}$$

where n is the iteration counter, initially set equal to one.

The background error covariance matrix \mathbf{B} is constructed by multiplying a uniform background variance

(σ_b^2) to an equivalent correlation model (implemented by repeating a Laplacian smoother using the zonal scale, x_L , and the meridional scale y_L) as

$$\rho(\mathbf{r}) = e^{-(r_x/x_L)^2 - (r_y/y_L)^2}, \quad (5)$$

where r_x and r_y are the zonal and meridional distance of the grid point to the observation location respectively. The elliptic property of correlation structure is controlled by x_L and y_L , which are plotted in Fig. 2b. For example, roughly 700 km x_L and 50 km y_L at the equator account for the effect of the well-known narrow correlation scale along the east–west near the equator (marked by “1” in Fig. 2a) while the correlation structure around 20°N(S) appears roughly isotropic due to the approximately equal x_L and y_L over there (marked by “2” in Fig. 2a). These correlation structures have the similar property as the estimates from the time mean of the temperature correlation at the surface in the EAKF (Fig. 2c).

While the time mean of prior error variance estimated in the EAKF has a spatial distribution (the SST standard deviation is shown in Fig. 2d) the background error variance used in the 3DVAR, σ_b^2 , is uniformly set to be $0.05[(^{\circ}\text{C})^2]$. This value is selected from tuning

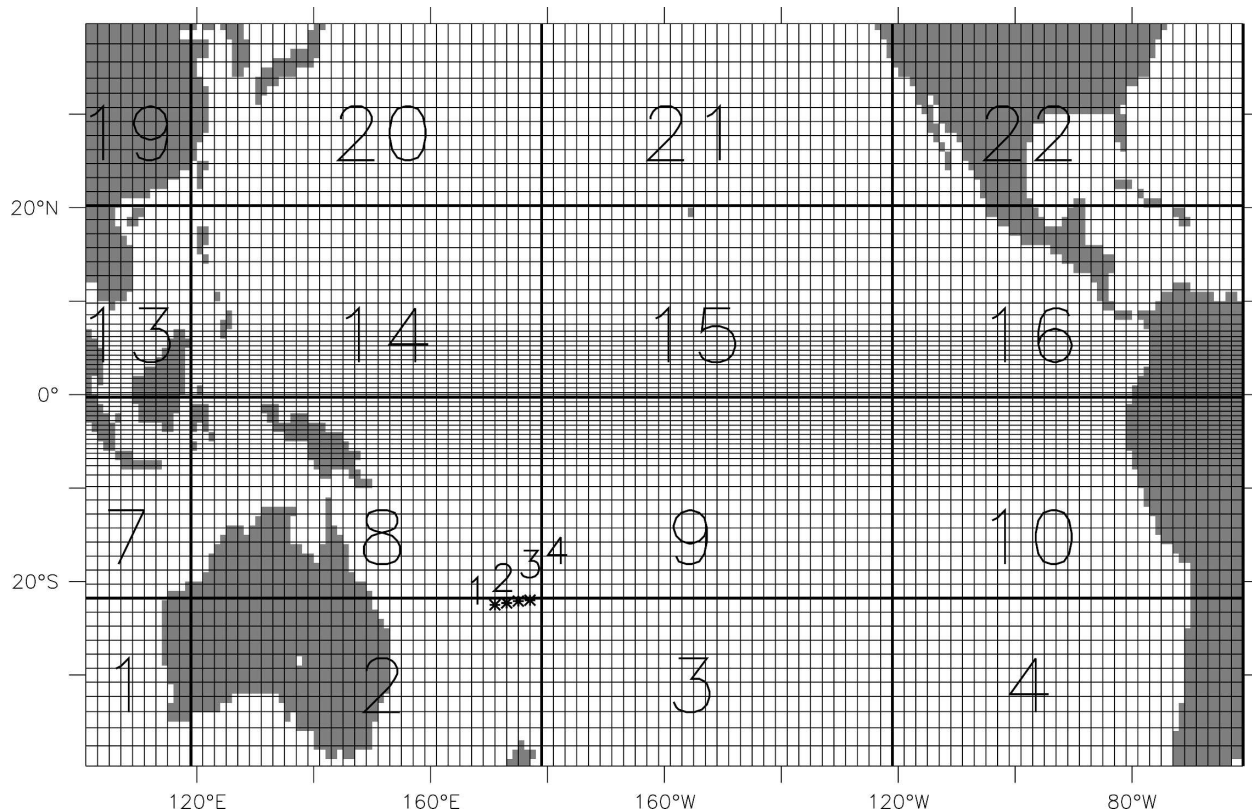


FIG. 3. The model grid configuration over the tropical Pacific basin. The number in each box is the PE index. Asterisk represents profiles used in section 4b and the numbers 1–4 denotes the profile index used in that section.

experiments to make the analysis reasonably close to observations without causing a too-large dynamical imbalance. More discussions on the spatial and temporal variation of prior error covariance in the ensemble filter will be made in section 5b.

3. The hybrid coupled model

As mentioned earlier, ocean data assimilation presents a special challenge for an ensemble filter, because of small ensemble spread, substantial model biases, and sparse observations. To enhance the ensemble spread and better sample the covariance structure of the ocean model, the ocean model is coupled to a stochastic atmosphere model. This additionally provides a prototype system for 1) representing the uncertainty of the atmospheric forcing, and 2) truly coupled ocean–atmosphere data assimilation, in the limiting case where no atmospheric data are assimilated.

a. The ocean model

The ocean model is the GFDL MOM4 (Griffies et al. 2003). For this study, the model is configured with 25

fixed depth levels, with 15-m grid spacing above 150 m. The horizontal grid spacing is 0.5° latitude near the equator, telescoping to 5° near the poles, and uniform 2 longitude. This gives a total of $180 \times 96 \times 25 = 432\,000$ grid points. The model grid configuration over the tropical Pacific basin is shown in Fig. 3. The model has an explicit free surface with explicit freshwater surface fluxes, a quicker advection scheme (Holland et al. 1998), nonlocal K-profile parameterization (KPP) vertical mixing (Large et al. 1994), and Laplacian horizontal diffusion and friction (Griffies and Hallberg 2000). Penetration of shortwave radiation into the surface layers is parameterized in terms of ocean color, using a prescribed climatology of Sea-viewing Wide Field-of-View Sensor (SeaWiFS)-measured chlorophyll concentrations that varies in space and time (Sweeney et al. 2005). The model has a 1-h time step and uses leapfrog time differencing with a Robert–Asselin time filter. Consistent with the time differencing, the analysis described in section 2a(1) uses a two-time-level adjustment (Zhang et al. 2004). Although the ocean model is global, a sponge poleward of 45° relaxes temperature and salinity toward the Levitus and Boyer (1994) climatology with an e -folding time of 30 days.

b. The statistical atmosphere

The statistical atmosphere is similar to those in Harrison et al. (2002) and Wittenberg (2002). The model attempts to capture the observed relationships between anomalous monthly mean tropical Pacific SST and surface fluxes (vector wind stress, shortwave and longwave radiation, sensible heat flux, evaporation, and precipitation) during 1979–2002 as represented in the NCEP/Department of Energy (DOE) Atmospheric Model Intercomparison Project (AMIP-II) reanalysis (NCEP2). The data are averaged onto a 2° latitude by 5° longitude grid covering the tropical Pacific Ocean (20°S–20°N, 120°E–70°W). A 12-month climatology is then computed and subtracted from the total fields to give monthly mean SST and flux anomalies. A statistical model is postulated for each flux anomaly time series:

$$\mathbf{Y} = \mathbf{XW} + \mathbf{E}, \quad (6)$$

where $\mathbf{Y}_{n \times q}$ is a matrix consisting of the n observed monthly means of the q -element flux anomaly field, $\mathbf{X}_{n \times p}$ is the corresponding matrix for the p -element SST anomaly (SSTA) field, $\mathbf{W}_{p \times q}$ are time-independent weights multiplying the SSTAs, and $\mathbf{E}_{n \times q}$ are stochastic shocks. We assume a priori that the flux shocks are normally and independently distributed in time, with zero mean and a variance that is stationary in time.

To obtain a set of predictors, we compute the SSTA/flux covariance matrix \mathbf{C} and perform a singular value decomposition:

$$\mathbf{C} = \frac{\mathbf{X}'\mathbf{Y}}{n-1} = \tilde{\mathbf{A}}\mathbf{D}\tilde{\mathbf{B}}', \quad (7)$$

where tildes denote nondimensional matrices; $\mathbf{D}_{r \times r}$ is a diagonal matrix, $r \equiv \min(p, q)$, whose diagonal elements are the singular values of \mathbf{C} ; and $\tilde{\mathbf{A}}_{p \times r}$ and $\tilde{\mathbf{B}}_{q \times r}$ are unitary matrices whose columns are the left (SST) and right (flux) singular vectors of \mathbf{C} . The SSTA weights are estimated by regressing the observed flux anomalies onto this set of predictors, namely, the SSTA singular vector expansion coefficients that explain the greatest fraction of squared covariance between the observed flux anomalies and SSTAs:

$$\hat{\mathbf{Y}}_N = \mathbf{X}\hat{\mathbf{W}}_N, \quad (8)$$

$$\tilde{\mathbf{W}}_N = \tilde{\mathbf{A}}_N(\tilde{\mathbf{A}}_N'\mathbf{X}'\mathbf{X}\tilde{\mathbf{A}}_N)^{-1}\tilde{\mathbf{A}}_N'\mathbf{X}'\mathbf{Y}, \quad (9)$$

$$\hat{\mathbf{E}}_N = \mathbf{Y} - \hat{\mathbf{Y}}_N, \quad (10)$$

where $\hat{\mathbf{Y}}_N$ and $\hat{\mathbf{E}}_N$ are the deterministic and residual stresses estimated from N predictors. A predictor is included only if it is an essential part of a group of three or fewer predictors that, together, significantly improve

TABLE 1. Number of SST predictors retained in each statistical surface flux model, and the percent of observed monthly mean anomaly variance captured by regression onto these predictors.

Flux field	No. of SST predictors	Variance % captured by regression
Vector wind stress	7	22.3
Shortwave radiative flux	3	20.3
Longwave radiative flux	5	12.4
Sensible heat flux	6	22.6
Evaporation	3	14.7
Precipitation	3	15.8

the model at more than half the grid points. Improvement at a grid point is deemed significant if a two-tailed F test on the change in residual sum of squares indicates less than 1% probability of that change occurring by chance. Table 1 shows the number of predictors obtained for each flux field, and the percent anomaly variance captured by each regression model.

For monthly mean stresses inside 20°S–20°N, the regression onto SSTAs explains less than 25% of the monthly stress anomaly variance, where this variance is computed over the set of all months and spatial grid points. The signal-to-noise ratio increases near the equator: the regression onto SSTAs captures nearly 50% of the variance for zonal wind stress anomalies averaged over 5°S–5°N. The signal-to-noise ratio also increases with time scale: the regression model captures nearly 75% of the variance for zonal stresses averaged over 5°S–5°N and filtered to retain only periods greater than a year. To represent the residual fluxes, we first note that the residuals and their principal components decorrelate rapidly, typically within 2 months or less (Wittenberg 2002). Thus rows of $\hat{\mathbf{E}}_N$ that are more than a few months apart are effectively independent realizations of stochastic fluxes. A straightforward way to include these in the model is to simply replay the time series of $\hat{\mathbf{E}}_N$ beginning in a random initial year and cycling back to the start of the time series whenever it reaches the end. Unlike the red noise approach of Wittenberg (2002), this provides only 24 (1979–2002) independent years of stochastic forcing; however, space-time correlations, propagating features, seasonal changes in variance, and cross correlations among variables are all preserved, making this an attractive option in an ensemble assimilation where the model's dynamical memory is constrained by observations.

The atmospheric forcing can thus be viewed as consisting of two parts: a slowly evolving “deterministic” part that depends on large-scale sea surface temperatures, and a highly chaotic (essentially stochastic) part that evolves independently of the ocean state.

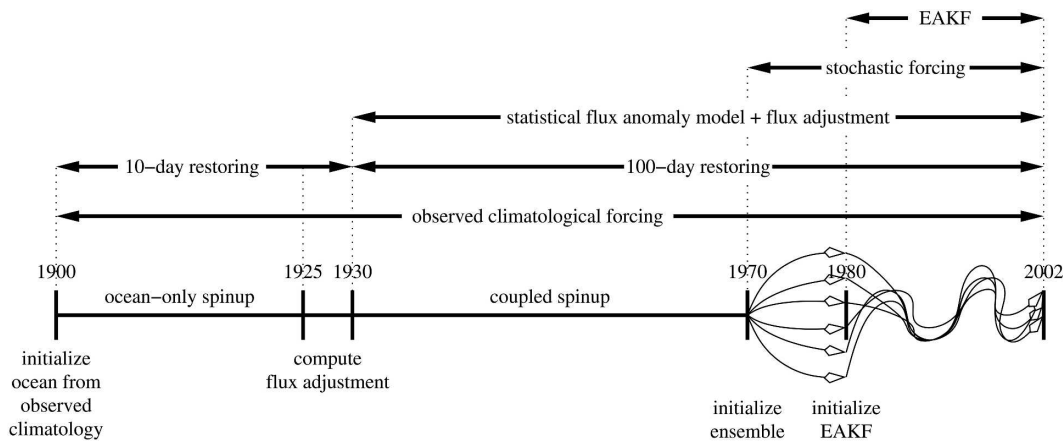


FIG. 4. Schematic of the model spinup and assimilation. The ocean model is initialized on 1 Jan 1900 from Levitus and Boyer (1994) climatological temperature and salinity. It is then integrated for 30 yr forced by observed climatological fluxes, with additional restoring terms that damp the model SST and SSS toward observed climatological values with an e -folding time of 10 days. The monthly climatologies of these two restoring terms are computed using the last 5 yr of this run. These climatological “flux adjustments” are then prescribed, the ocean model is coupled to the statistical flux anomaly model, the SST and SSS restoring is weakened, and the model is integrated in coupled mode for another 40 yr. Starting from six identical copies of the state at 1 Jan 1970, the model is integrated for another 10 yr with each ensemble member forced by a different realization of the stochastic fluxes. The ensemble states at 1 Jan 1980 then compose the initial conditions for the EAKF.

c. Flux adjustment and the ensemble spinup

The spinup of the coupled model is illustrated in Fig. 4. First the ocean model is initialized from Levitus and Boyer (1994) climatological temperature and salinity. The ocean model is then integrated for 30 yr, forced by climatological fluxes from the NCEP2 reanalysis, with additional restoring terms that damp the model SST and sea surface salinity (SSS) toward the climatological values with an e -folding time of 10 days over an upper-ocean cell of 10-m thickness. The monthly climatologies of these two restoring terms are computed using the last 5 yr of this run. These climatological “flux adjustments” are then prescribed, the ocean model is coupled to the statistical flux anomaly model, and the SST and SSS restoring is weakened to have an e -folding time of 100 days. This approach permits the coupled model to maintain a realistic climatology without significantly damping interannual variability.

Next, the flux-adjusted hybrid coupled model is integrated for 40 yr without any stochastic forcing to obtain the initial condition for the ensemble spinup. The model has self-sustained, irregular oscillations when the wind stress noise forcing is active. In the absence of noise forcing, the model is stable, but increasing the air–sea coupling (by increasing the strength of the statistical wind stress feedback) renders the model linearly unstable such that it sustains regular oscillations with a period of 3.3 yr. Starting from six identical copies of this initial state, the model is ensemble integrated for 10 yr

with each ensemble member forced by a different realization of the residual fluxes (the integrations are initialized at midnight 1 January, and each member feels a stochastic forcing beginning at midnight 1 January of a different residual year). The different ensemble states following this spinup compose the initial conditions for the EAKF experiments in the remaining sections.

4. Parallelization of the EAKF

a. Domain decomposition

In some circumstances, parallelizing the ensemble filter may be required to reduce computational time and memory usage. There are several possible algorithms for parallelizing the filter. First, if many observations are available at each observation time, the sequential algorithm can be recasted in a matrix form. The application of the forward operator (which is now a vector function) and the matrix inversion required to compute the impact of the observations on state vector elements can then be performed using parallel algorithms. This is an example of a naturally scaling exact algorithm, but it might not be particularly efficient on parallel systems with relatively slow interprocessor communication.

Here, an approximate algorithm of the compute domain/data domain strategy of Anderson (2001) is used to parallelize the sequential filter, making use of the fact that the impact of observations is localized to a small set of “nearby” state variables. The model grid is

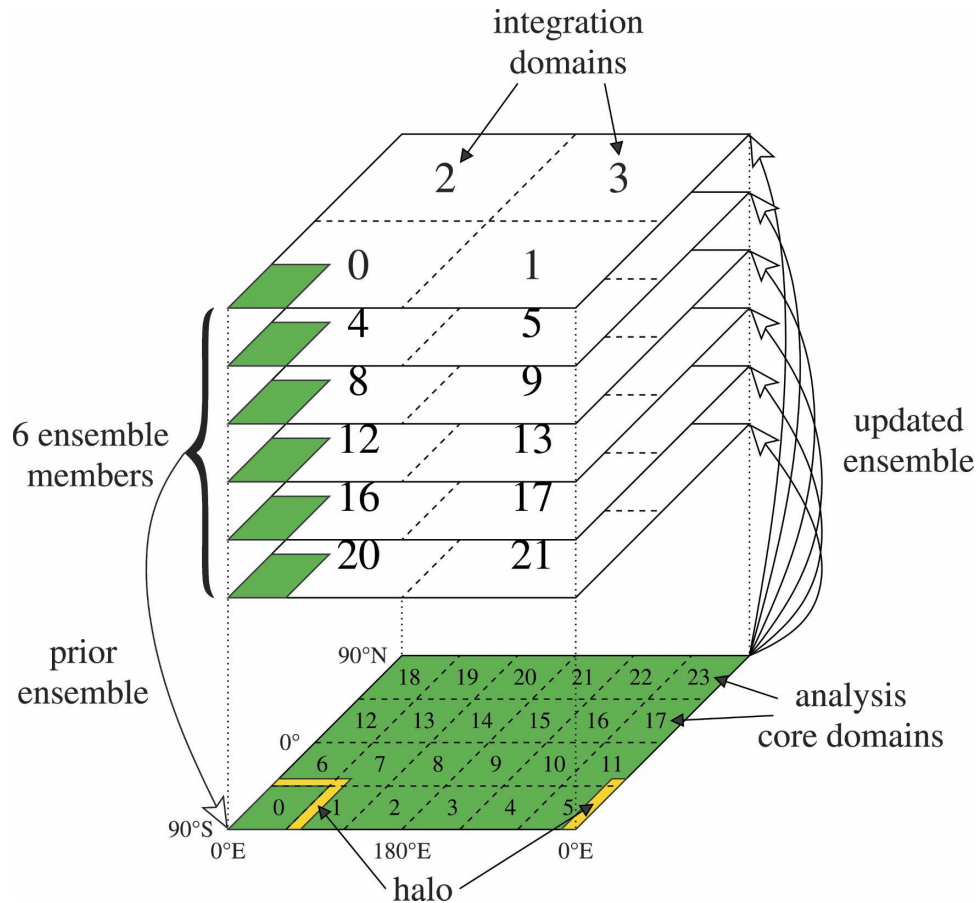


FIG. 5. Domain decomposition of a scalar field in the parallelized ensemble filter. The six ensemble members are integrated forward in time, in parallel, using four processors each (one for each quarter of the globe). At analysis time, the members are synchronized and the “prior” ensemble at each observational point is broadcast to all 24 analysis processors. For each physical field, each analysis processor then uses the observations to sequentially update the six-element ensemble vectors at each grid point in its core domain (green) and halo (yellow). Once all nearby observations have been assimilated, the updated ensemble vectors in the core domains are transmitted back to the integration processors, completing the cycle.

partitioned horizontally into a set of computational domains, each surrounded by a halo of additional grid points. The compute plus halo regions are referred to as an analysis domain. When a set of observations becomes available, the appropriate parts of the prior state ensembles are copied to each of the analysis domains. An observation is assimilated in a particular analysis domain only if all the state variables required for its forward operator (given the bilinear interpolation used here, this is simply a set of four adjacent grid points) are available in the analysis domain. On each analysis domain, all of the appropriate observations are assimilated sequentially and the state in the analysis domain is updated as appropriate before the next observation is assimilated. However, no communication between analysis domains is performed during the assimilation

cycle. Points near the edge of the analysis domain will not be appropriately impacted by observations that are just outside of the analysis domain. The net result is that the prior ensembles used within each analysis domain will have an erroneous ensemble spread and may cause analysis errors. However, appropriate choices of the computational and halo sizes can minimize the errors associated with this effect. This approach is similar to the local ensemble filter of Ott et al. (2004).

Figure 5 shows the domain decomposition and communications for a case with 24 processing elements (PEs) and the six ensemble members used in this study. There are two types of domain decomposition in the horizontal: integration domains and the analysis domains described in the last paragraph. Integration domains are used to advance ensemble members in time;

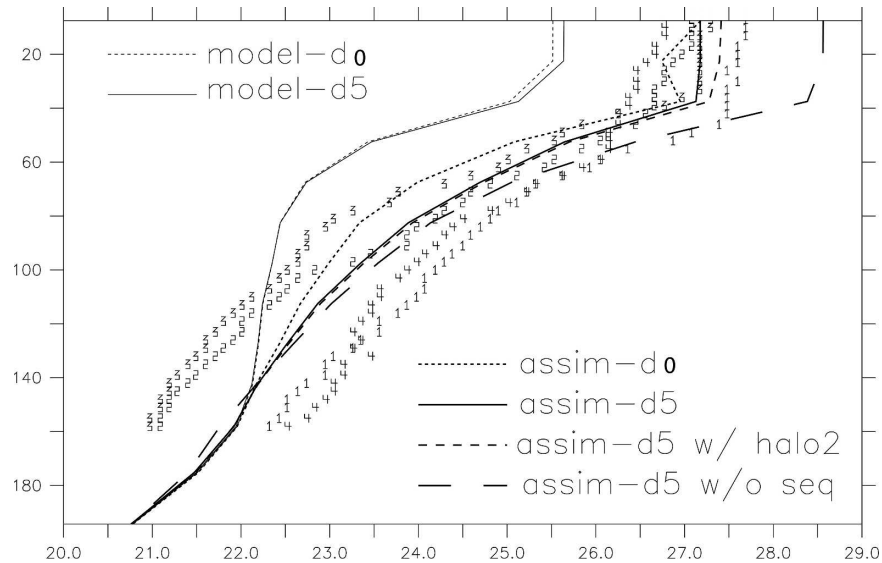


FIG. 6. The adjustments of the model profile at (22°S , 175°E) (dotted for day 0 and solid for day 5) by four observational temperature profiles located at (22.5°S , 171°E) (denoted by 1), (22.3°S , 173°E) (denoted by 2), (22.1°S , 175°E) (denoted by 3), and (22°S , 177°E) (denoted by 4), through one-step (thick dotted) and five-step (thick solid) global sequential analysis, five-step parallelized analysis with the halo size as two- (dashed) and five-step nonsequential analysis (long dashed).

each PE works on a part of the globe for one ensemble member.

b. Choosing a halo size to ensure the sequential computation in EAKF

In choosing a halo width, there is a trade-off between parallelism and analysis quality. With no halo, the EAKF is highly parallel but may suffer from reduced quality near the edges of analysis domains. With a halo that encompasses the entire planet, each PE conducts an identical global sequential analysis with no edge effects—but then the algorithm is no longer parallel. The challenge is to choose a halo that provides the optimal balance of quality and parallelism. This can be done by choosing a halo large enough to encompass all observations that affect the core analysis region (which is the only region communicated back to the integration PEs).

Designing a simple analysis domain layout is relatively straightforward since observations are only being assimilated within Tropics in the current experiments. A global assimilation would be presented with more difficult problems as the model grid became more dense away from the equator. The parallelization is also assisted in the present study by the tightly localized regions of significant correlations between observations and state variables that result from the use of a very small ensemble (see Fig. 9 and associated discussion).

The meaningful impact of observations is confined to a very few neighboring grid points in the horizontal and so is ideally suited to the parallelization method chosen. The use of large ensembles that are able to extract weaker observations to state correlations could also lead to a much more difficult implementation, where more strategy (e.g., Fukumori 2002, the partitioning technique) can be considered.

To demonstrate the impact of the halo size on the EAKF, an assimilation of four profiles during 1–5 January 1980, shown by asterisks in Fig. 3 (22.5°S , 171°E), (22.3°S , 173°E), (22.1°S , 175°E), and (22°S , 177°E), is performed. The four profiles are located near the domain corners of PE2, PE3, PE8, and PE9, in a fairly inactive region of the southwestern tropical Pacific.

Figure 6 presents the adjustments of the model profile at (22°S , 175°E), in which the thin-dotted line (day 0) and the thin-solid line (day 5) show the change of the model profile in 5 days. One-step global sequential analysis (thick-dotted line) adjusts the model profile close to the observations, and after four more analysis steps, the adjusted model profile (thick-solid line) is refined to fit the observations (marked by 1, 2, 3, and 4 in Fig. 6 corresponding to the profile indexes in Fig. 3) very well. On the other hand, in the parallelized analysis, to show the importance of sequentially updating the ensemble estimate of observations, we first check how the assimilation performs if only the first guess of the

ensemble estimates for all observations is used. The long-dashed line represents the 5-day adjusted model profile, using only the first-guess ensemble estimates for observations and background covariance without sequential update. This shows that without updating the ensemble estimates, the observational constraint is greatly overestimated since the computation violates Bayes' rule. This overestimate causes the adjusted model profile to have a departure from observations to the other side. With no halo each observation gives a positive analysis increment that does not take into account that another observation may have already reduced the background/observation mismatch. If the halo size is set as 2, PE2 and PE8 can update the ensemble estimates for all observations but PE3 and PE9 can only update the ensemble estimates for profiles 3 and 4. Under this circumstance, the 5-day adjusted model profile (dash) is still a little overestimated. As the halo size increases to 4, the adjusted model profile is very close to the global sequential analysis, and increasing the halo size to 6, the model profile adjusted by the parallelized analysis is bitwise-identical to the global analysis.

The analysis above on four-profile assimilation results shows that choosing an appropriate halo size can ensure that observations strongly impacting a given point will know each other when an ensemble-based filter is parallelized. Typically, the halo size scale can be determined by the covariance localization described in section 5a.

5. Assimilation results for 1980–2002

a. Data, impact domain, and halo size

Considering the difference of zonal and meridional grid structure a 6-point longitude \times 10-point latitude halo is chosen in the parallel EAKF described in section 4. Observations used include profiles maintained by the National Oceanographic Data Center Global Temperature and Salinity Pilot Program (NODC/GTSPP), Tropical Atmosphere and Ocean (TAO) array, and Reynolds SST. No observations are assimilated outside of 30° latitude. For each grid point, the impacting observations are limited within a $\Delta\phi \times \Delta\lambda \sec\phi$ window, where $\Delta\phi$ and $\Delta\lambda$ are the latitudinal and longitudinal widths (20° in this study) and $\sec\phi$ is the latitudinal adjustment factor of the longitudinal width. The analysis domain with 6×10 halo structure covers most of the impacting observations [limited within 10° south–north (east–west) of a grid point]. First, the massively parallel processing EAKF (MPPEAKF) assimilation and the global sequential ensemble filtering assimilation (identical on each PE) are run from 1996 to 1999 to

check the quality of the parallel analysis. Results (not shown here) show there is no qualitative difference between the parallel analysis and the global sequential analysis and both assimilated SSTs are nearly identical to the Reynolds. The MPPEAKF tremendously reduces both computational cost and storage for assimilation comparing to the global sequential EAKF (around one-tenth as a factor for both in this case).

b. Examination of assimilation results

Both the MPPEAKF and the 3DVAR is run from 1 January 1980 to 1 December 2002 with a daily analysis interval, and the 3DVAR is relaxed to “observed” (Reynolds) SSTs and the MPPEAKF assimilates the SSTs. Figure 7 shows that the filtered ensemble of SSTs, ocean heat contents (averaged temperature over top 300 m), and thermal structures converge well through the constraint of observations, despite the imposed noise forcings. As with the equatorial Pacific (e.g., averaged over 2°S–2°N) SST anomalies in the 3DVAR (restoring SSTs), again the ones in the MPPEAKF analysis (ensemble mean; not shown here) are nearly identical to the Reynolds for the whole 23-yr period.

Figure 8 evaluates the MPPEAKF (ensemble mean) and 3DVAR temperatures at 140°W on the equator, which shows that the 3DVAR (a standard 3DVAR experiment uses “observed” NCEP products for fluxes, marked by “3DVAR^O”) subsurface structure noticeably departs from the observations (e.g., weaker 1986/87, 1991/92, and 1997/98 warm events, weaker 1987/88 and stronger 1998 cold events) while the MPPEAKF follows them much more closely. To examine the role played by the estimated fluxes from the coupled MPPEAKF assimilation, the 3DVAR is run using the fluxes from the coupled assimilation and corresponding results are shown as “3DVAR^E” in Fig. 8. Overall, the 3DVAR^E looks like a smooth version of the 3DVAR^O, which is not able to improve the assimilation, and in fact degrades the 3DVAR solution. The same conclusion can be made while the ocean-only simulation using the fluxes from the coupled assimilation (marked by “model^E” in Fig. 8) is compared to the simulation using the observational fluxes (marked by “model^O” in Fig. 8); that is, generally the model^E is a smooth version of model^O. These experiment results show that despite the errors in the EAKF simulated fluxes, the ensemble filter is able to perform better than the 3DVAR, which uses the “observed” fluxes.

The difference of the (top) 3DVAR and (middle) MPPEAKF assimilated temperature at (0°, 140°W) can be more clearly shown by the climatological seasonal cycle (subtracting the annual mean from the climatol-

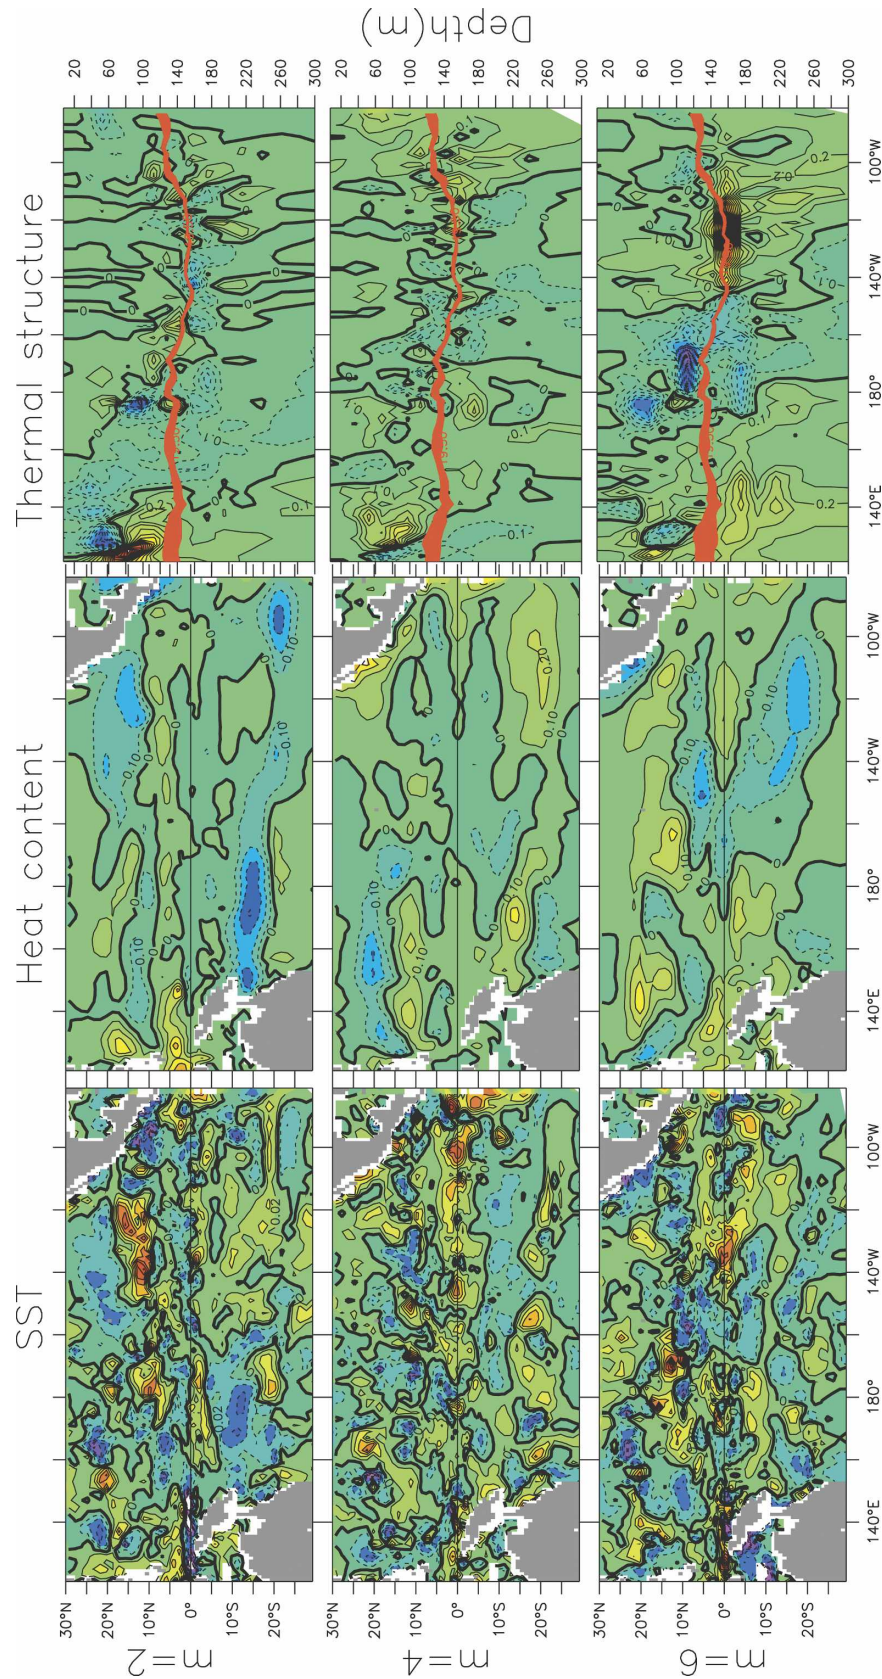


FIG. 7. The monthly mean SST, heat content (averaged temperature over top 300 m), and subsurface thermal structure of the equatorial Pacific for ensemble members 2, 4, and 6 (departure from the ensemble mean) in Nov 1997. The contour intervals are 0.02°C for SST, and 0.1°C for heat content and vertical temperature structure. In the rightmost column, isotherms from 19.5° to 20.5°C are indicated in red to represent the thermocline.

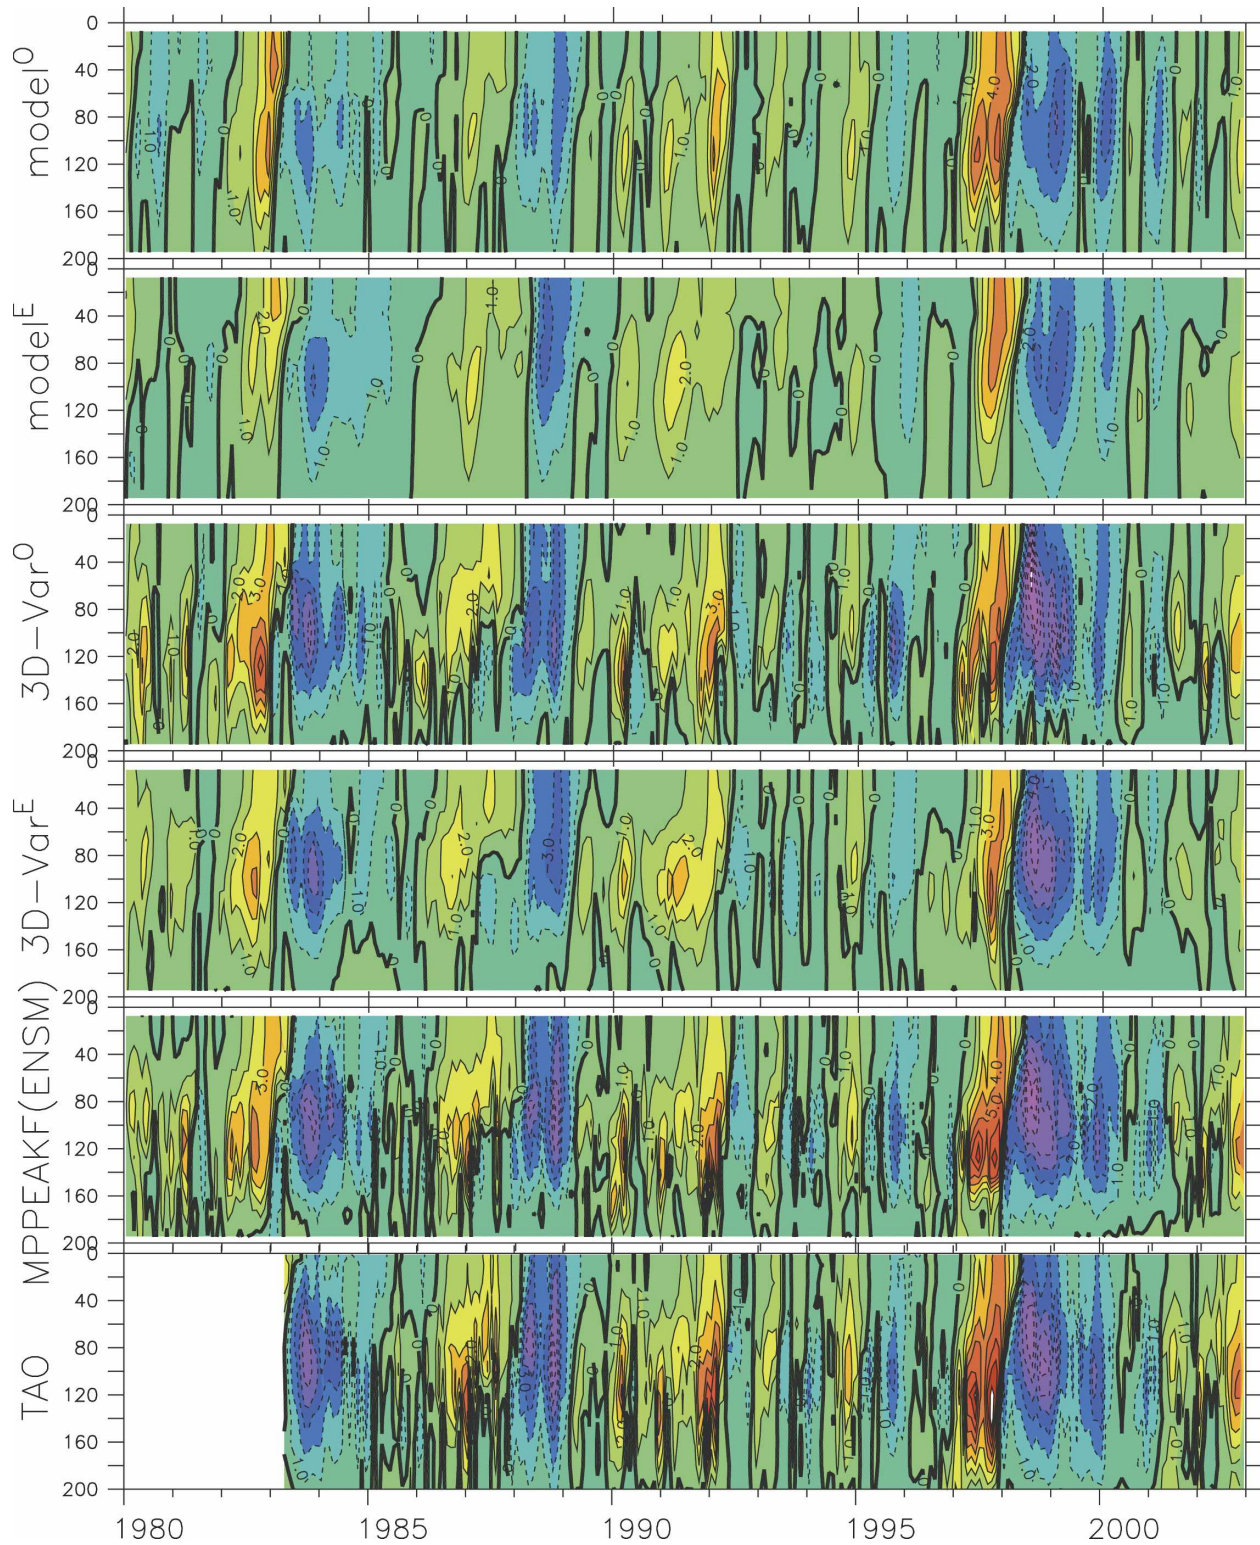


FIG. 8. The time series of the anomalies of the temperature profile at (0°, 140°W) of the ocean-only simulation using the observational fluxes (model^O) and the MPPEAKF coupled assimilation fluxes (model^E), the 3DVAR assimilation using the observational fluxes (3DVAR^O), and the MPPEAKF coupled assimilation fluxes (3DVAR^E), parallelized EAKF assimilation (MPPEAKF, ensemble mean), and TAO. The contour interval is 1°C.

ogy) shown in Fig. 9. The MPPEAKF follows the observational seasonal cycle much better than the 3DVAR. The causes of the differences between the 3DVAR and the ensemble filter are complex and cannot be completely isolated by this study: It is possible that the spatially and temporally varying aspect of the covariance provided by the ensemble filter is an important factor in data assimilation (Zhang and Anderson 2003), but this study cannot categorically confirm this hypothesis, which requires further research work. In addition, using the full multivariate aspects of the EAKF would have further obscured this comparison.

An example of the anisotropic and temporally varying nature of the background covariance used in the MPPEAKF is shown in Figs. 10 and 11. Figure 10 presents the variation of (a), (b) time mean standard deviation and (c), (d) correlation scales (the symbol “*” marks the reference points) of the EAKF in zonal vertical at the (a), (c) equator and (b), (d) meridional vertical at 140°W . Figure 11 presents the time series of the 120-m temperature standard deviation and the surface temperature autocorrelation about a point (0°N , 123°W) over the east equatorial Pacific (160° – 80°W). The maximum standard deviation in the second half of 1997 may reflect the 1997/98 warm event. The difference of the west–east bound autocorrelation of the surface temperature about the reference point may reflect, to some degree, the wind stress direction that organizes the warm/cold phase of the surface water. A complete understanding of the estimated background error covariance of ocean state variables evaluated by an ensemble-based filter is very important to understanding the model dynamics, but this topic is beyond the scope of this study. Figure 11 also highlights the limitations of the extremely small, six-member ensembles used here. Two random samples drawn from a normal distribution have an expected correlation of nearly 0.4 due to sampling noise. This is reflected in Fig. 11 (right) where the minimum time mean values of correlation to a given longitude are bounded below by about 0.4. Sampling error also impacts larger correlations so that only time mean values close to 0.8 in Fig. 11 indicate that signal is dominating noise. As can be seen, the meaningful impact of observations is localized to only a few grid points surrounding the observation. Future work will examine the impacts of increasing the ensemble size in order to improve the signal-to-noise ratio of extracted information from observations.

To better evaluate the MPPEAKF assimilation, we compare the Equatorial Undercurrent (EUC) (Fig. 12) and upwelling (Fig. 13) of the equatorial Pacific of MPPEAKF (left, ensemble mean) and 3DVAR (right) assimilation, on the onset (upper) and mature (lower)

phases for the 1997/98 warm event. For the onset phase (July 1997), the MPPEAKF shows a strong westerly burst in the top layer and dominant easterlies below 80 m, while the 3DVAR has some localized westerly or easterly centers throughout top 300 m. For the mature phase, the central Pacific westerly at the top layer of the MPPEAKF weakens and transfers toward easterly while the EUC of the 3DVAR keeps stronger localized westerly and easterly centers throughout the whole layer of top 300 m. Comparing to the MPPEAKF upwelling magnitude of a couple of meters per day (Fig. 13, left), the upwelling of the 3DVAR (Fig. 13, right) is far too strong, in some localized centers exceeding 10 m per day.

The stronger EUC and upwelling in 3DVAR may be due to the prior specified background covariance between the model/observational temperature profiles, which may overestimate the observational constraint, as shown in Fig. 14 where the temperature correction in both MPPEAKF and 3DVAR assimilations are presented. From both (a)–(d) time mean and (e) time series, some stronger localized analysis correction in 3DVAR than MPPEAKF are observed in Fig. 14. Although some instantaneous flow signal may be considered in minimizing a defined cost function, the 3DVAR, due to the homogeneous and flow-independent nature of the prior specified background covariance, may produce some localized temperature gradients, so that the EUC and upwelling derived from the analyzed temperature gradient may not be consistent with the dynamics. This analysis is consistent with the comparison of the EUC climatology at (0° , 140°W) of the 3DVAR (top) and MPPEAKF (middle) assimilations shown in Fig. 15, in which the climatology of the TAO current profiles (bottom) also is plotted for reference. Figure 15 shows that the 3DVAR produces a stronger EUC, which loses the seasonal cycle phase at the top layer, while the MPPEAKF produces a much weaker EUC with consistent seasonal cycle phases.

c. Examination on time series of analyzed ensemble

One of the key advantages of an ensemble filter is the estimation of the analysis uncertainty (PDF). Through the examination of the time series of the analyzed ensemble versus observations, one can evaluate the performance of the assimilation. As an example, here we conduct a verification-assimilation experiment in which some observational profile is excluded and then we check the consistency of the guessed profile by the assimilation process with the observed profile. For example, withholding the profiles at (0° , 140°W) we can make an assimilation guess for profiles at this point. Experiment results show (not shown here) that gener-

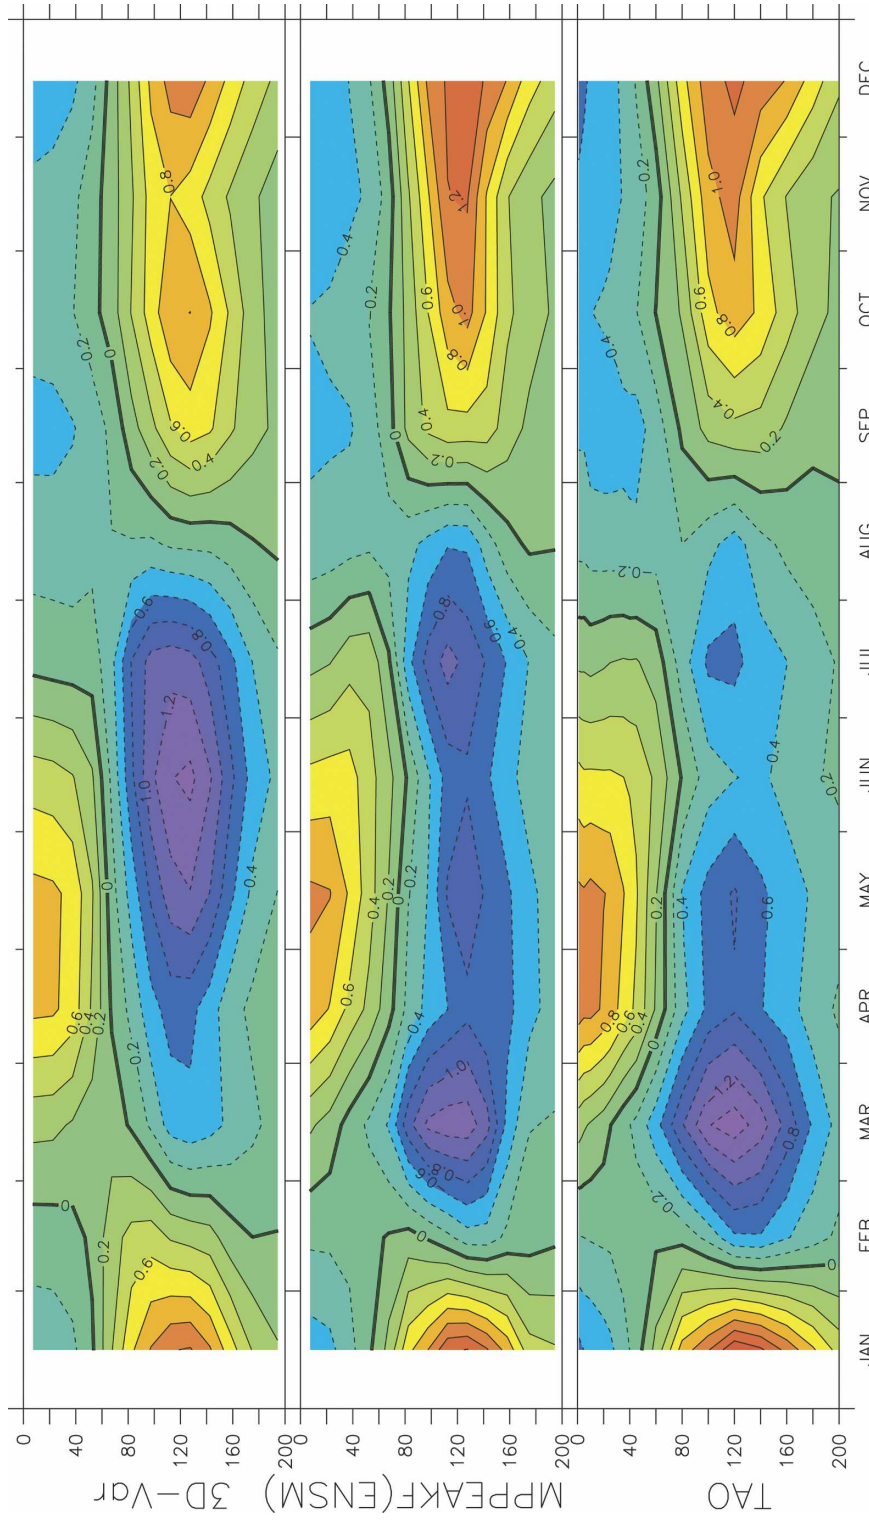


FIG. 9. The climatological seasonal cycle (subtracting annual mean from climatology) of the temperature profile at (0°N, 140°W) of the (top) 3DVAR, (middle) parallelized EAKF assimilations, and (bottom) TAO, during 1980–2002.

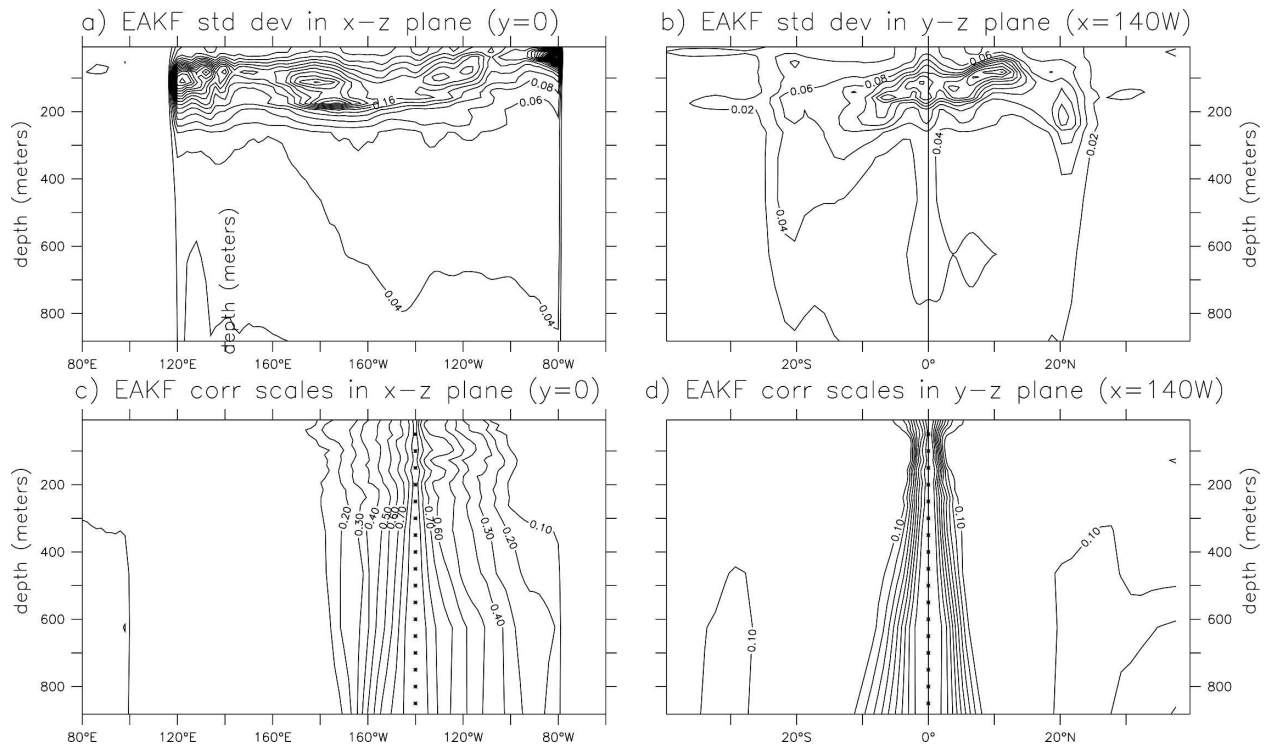


FIG. 10. The (a), (c) zonal-vertical (at the equator) and (b), (d) meridional-vertical (at 140°W) variation of the (a), (b) time mean standard deviation and (c), (d) correlation scales (the symbol “*” marks the reference points) in the EAKF, averaged over 23-yr assimilation period.

ally the MPPEAKF assimilation guesses are qualitatively equivalent to the full assimilation and both are tracking the TAO observations above the thermocline very well, while free ensemble forecasts diverge from the observations. This phenomenon means that the MPPEAKF assimilation procedure can correct the model bias and coherently fill the data gap according to the model dynamics in a reasonably dense observational network [around (0°, 140°W); basically only TAO profiles are available during this period]. How to use this kind of assimilation-guess experiment to validate an assimilation scheme is under investigation in a follow-up study that requires a large number of experiments to make significant statistics.

The equivalence of the assimilation-guessed and analyzed temperature allows us to examine the time series of analyzed ensemble versus observations for the whole period of 1980–2002, as shown in Fig. 16. Figure 16 shows that during this period the assimilated temperature ensemble members (blue-dotted) are again tracking the TAO observations very well above the thermocline, while the free ensemble forecasts (red-dashed) only oscillate with an annual cycle, having a big spread above the thermocline and a small spread in deep water. The performance of the assimilation in

deep water (lower-right panel) is more interesting: although the model spread is small the filter is, to some extent, still able to correct the model bias according to observations. During 1996–98 since both expendable bathythermographs (XBTs) and TAO data are frequently gapped (especially TAO observations for this particular region) in deep water (below 200 m) the analyzed temperature at 500 m (lower-right panel) stays close to the free ensemble forecasts. After that period when TAO observations in deep water are available the analyzed temperature is adjusted back. This phenomenon means the observational data in deep water are important for the estimation of the three-dimensional ocean states that are, perhaps, of central importance for ocean climate prediction. However, when the model spread becomes increasingly small in deeper water, how to efficiently extract the observational information to reduce the model bias is another research issue.

6. Impact on coupled forecasts

We next examine the usefulness of the MPPEAKF for initializing ENSO forecasts. We focus on the forecasts of initialized from the *ensemble mean* solution, and without stochastic forcing during the actual fore-




FIG. 11. The time series of the (left) 120-m temperature ensemble std dev and (right) the surface temperature ensemble autocorrelation about $(0^\circ\text{N}, 123^\circ\text{W})$ over the east equatorial Pacific $(160^\circ\text{--}80^\circ\text{W})$ during 1996–99.

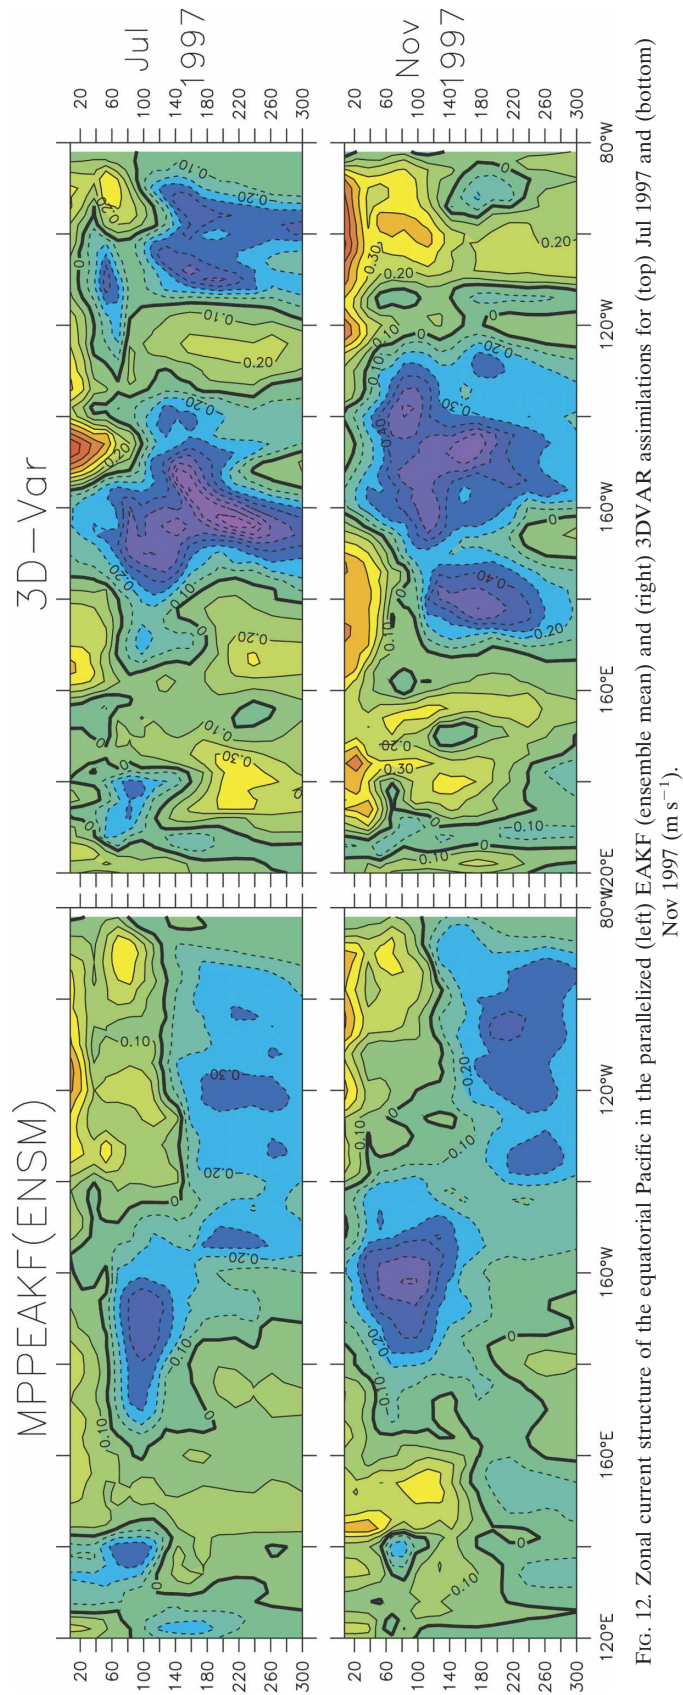


FIG. 12. Zonal current structure of the equatorial Pacific in the parallelized (left) EAKF (ensemble mean) and (right) 3DVAR assimilations for (top) Jul 1997 and (bottom) Nov 1997 (m s^{-1}).

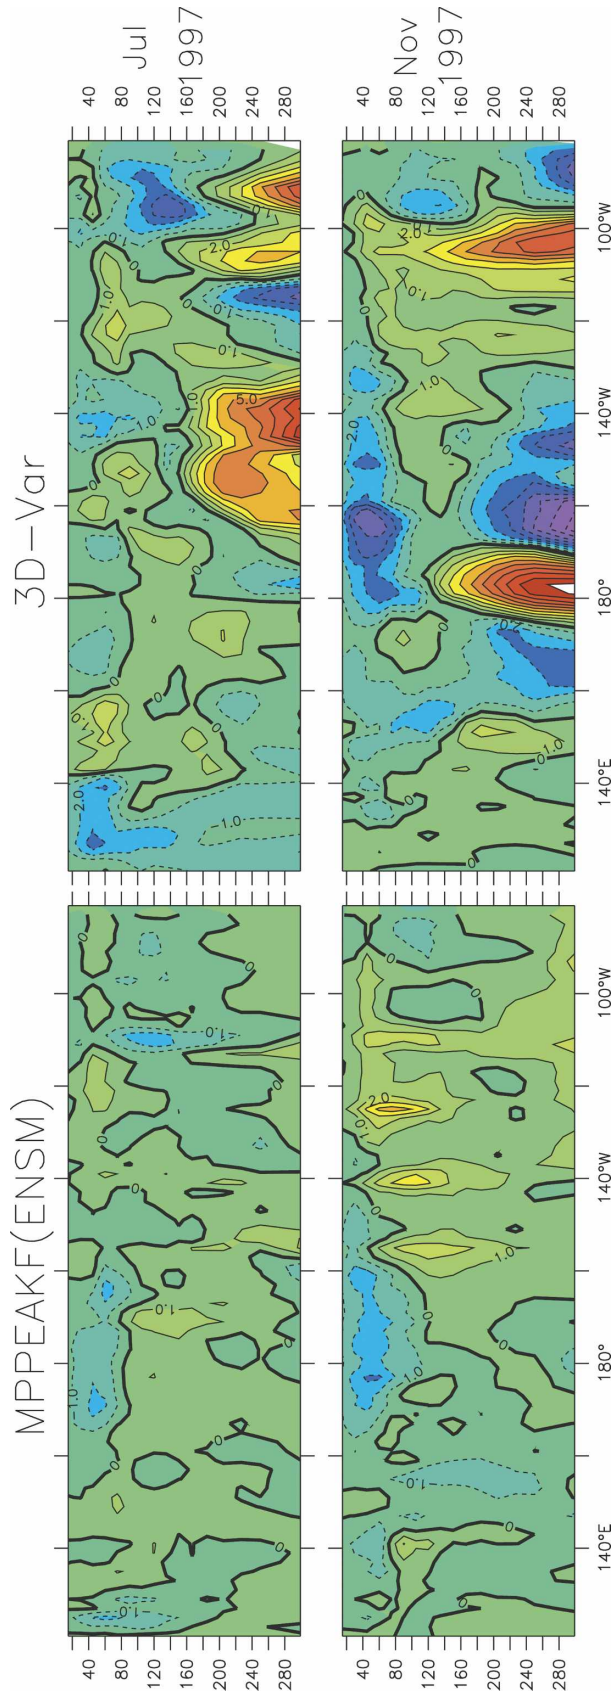


FIG. 13. The same as Fig. 12 except for upwelling (m day^{-1}).

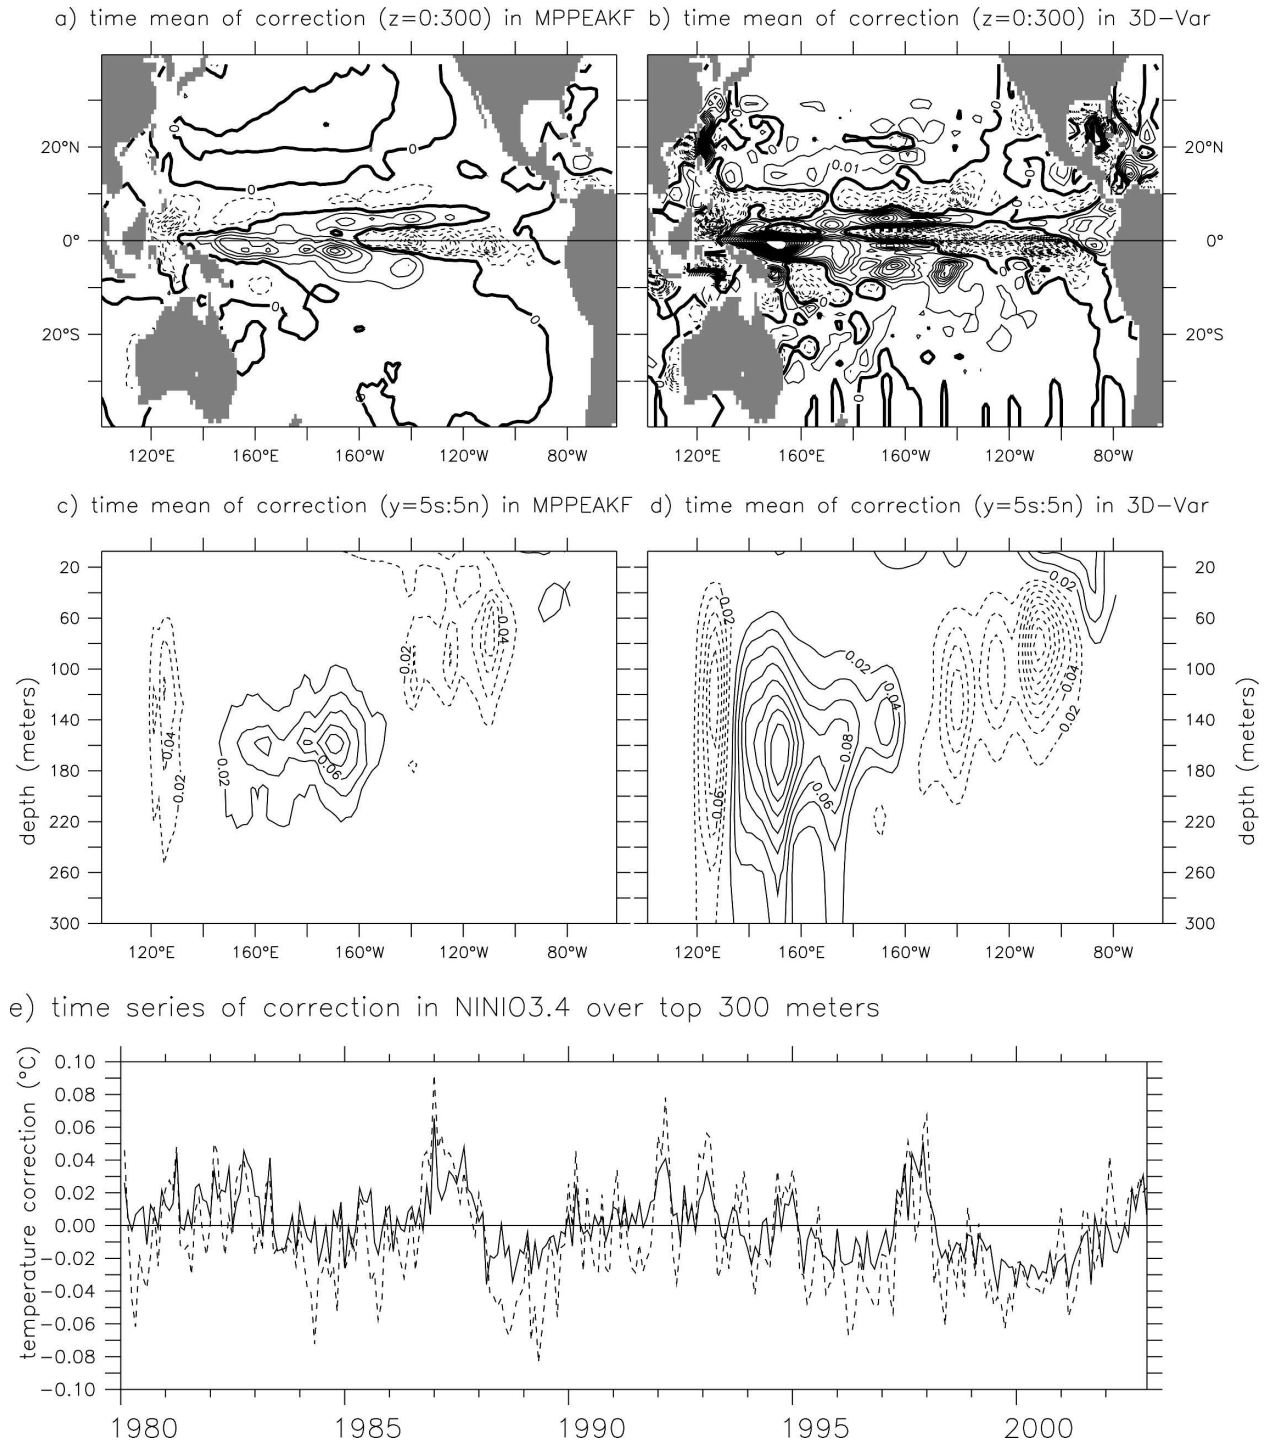


FIG. 14. (a)–(d) Time mean and (e) time series of temperature correction in the 1980–2002 assimilation period. The x - y plane is the average over the top ocean 300 m in (a) MPPEAKF and (b) 3DVAR; the x - z plane is the average of 5°S–5°N in (c) MPPEAKF and (d) 3DVAR; (e) the time series is for Niño-3.4 averaged over the top ocean 300 m. The contour interval is 0.01°C for (a) and (b) and 0.02°C for (c) and (d).

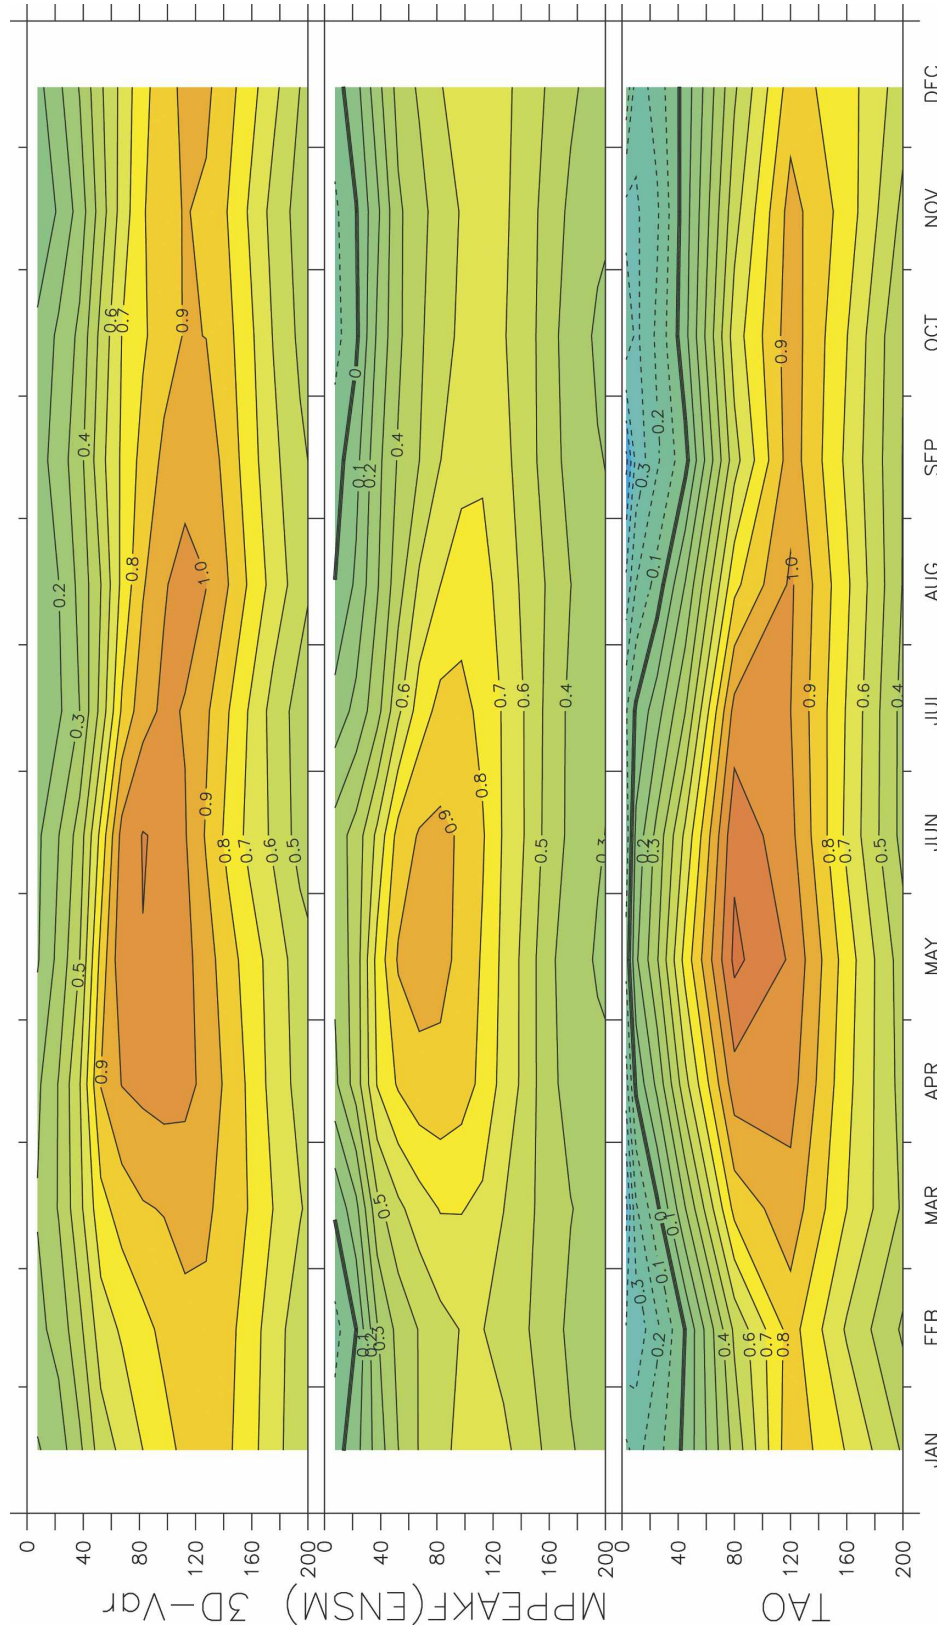


FIG. 15. The Equatorial Undercurrent climatology at (0°N, 140°W) for the (top) 3DVAR, (middle) parallelized EAKF assimilations, and (bottom) TAO profiles. Units are $m\ s^{-1}$.

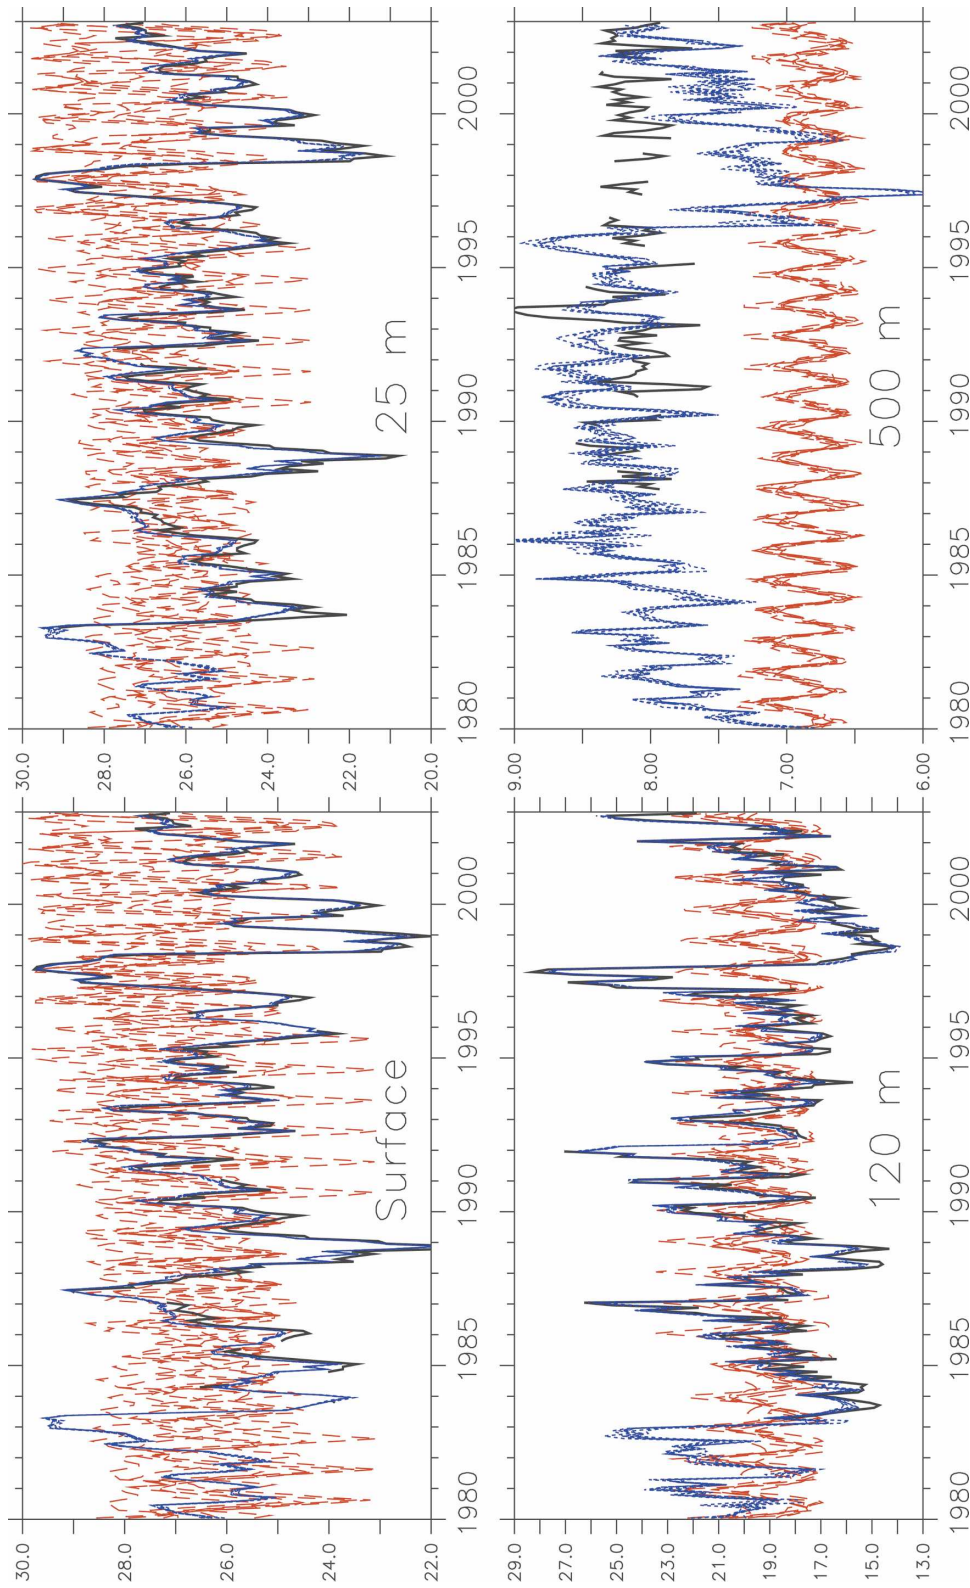


FIG. 16. Time series of monthly mean temperature ensemble samples ($^{\circ}\text{C}$) at (0° , 140°W) of the free ensemble forecasts (red dashed), MPPEAKF assimilation (blue dotted), and TAO profile at (0° , 140°W) (black) for (top left) sea surface, (top right) 25-, and (bottom left) 120-, and (bottom right) 500-m ocean during the 1980 ~ 2002 assimilation period. The ensemble assimilation/forecast is initialized by the spinup ensemble initial condition at 0000 UTC 1 Jan 1980.




FIG. 17. Skill evaluation of 12-month hybrid coupled model forecasts initialized from the MPPEAKF ensemble-mean (green) and the 3DVAR assimilation (red), for SST anomalies averaged over the Niño-3.4 region (5°S – 5°N , 170° – 120°W). (top) Results for 12 forecasts initialized at midnight on 1 Jan 1991–2002. (bottom) Results for 11 forecasts initialized at midnight on 1 Jul 1991–2001. (left) The evolution of the forecast bias (forecast minus observations) for each month after initialization. (middle) Rmse for forecasts after bias correction; for reference, the dotted line shows the observed std dev of Niño-3.4 SST anomalies for each month. (right) Correlations between forecast and observed anomalies for each month. As a benchmark, solid black curves indicate forecasts made by simply persisting lead-zero SST anomalies unchanged into the future.

cast. Under the assumptions of the filter the ensemble mean provides the best linear unbiased estimate of the ocean state at each time. This single state is used to initialize a set of hybrid coupled forecasts, including the deterministic fluxes by modeled SSTs without residual parts (noise forcings), starting at midnight on 1 January of each year from 1991 to 2002 (12 forecast cases), and at midnight 1 July of each year from 1991 to 2001 (11 forecast cases). Complementary sets of forecasts are launched from the 3DVAR assimilation. Summary statistics for SST anomalies averaged over the equatorial central Pacific are shown in Fig. 17.

For January starts, the MPPEAKF initializations give a smaller forecast bias, slightly lower rms error, and a higher correlation with observed anomalies than the 3DVAR initializations, over the first few months of the forecasts. The MPPEAKF forecast bias is slightly worse by May–August, but otherwise the skill of the forecasts from the MPPEAKF is comparable to those from the 3DVAR. For July starts, the MPPEAKF forecasts have a slightly larger bias than the 3DVAR forecasts, but for the MPPEAKF the rms error is reduced and the anomaly correlation is improved for all lead times up to 11 months.

While there is room for improvement in the forecasts initialized from the MPPEAKF, it appears that for this forecast model the MPPEAKF ensemble mean produces forecasts with slightly better skill than those initialized from the 3DVAR analysis. Further improvements may be possible by 1) launching forecasts from the individual MPPEAKF ensemble members (instead of only the ensemble mean), and evaluating the PDF of those ensemble forecasts, or 2) turning on the stochastic forcing during the forecasts, and launching “stochastic ensembles” of forecasts for each initial condition. The ultimate goal is a combination of these, that is, probabilistic forecasts launched from probabilistic initial conditions.

7. Summary and future directions

A parallel ensemble filter has been implemented in a GFDL hybrid coupled model, which serves both as a prototype for representing uncertainties in the surface forcing and as a test bed for truly coupled data assimilation. A parallel scheme is designed and applied to a modified ensemble adjustment Kalman filter (EAKF) algorithm under a local least squares framework (Anderson 2003). The parallel scheme requires an analysis domain consisting of a core domain plus a halo for each processor element (PE). The analyzed ensemble (arranged in core domains) is transposed into the model integration domains of the individual ensemble members, so that the system can synchronously advance the ensemble and conduct a parallel analysis. A halo is used to retrieve the updated information on background covariances, for those observations outside the core domain that impact grid points within the core domain. When the halo is sufficiently large, the massively parallel processing EAKF (MPPEAKF) produces a solution with the same quality as a global sequential analysis.

The MPPEAKF is used to assimilate observed temperature profiles from 1980 to 2002, using six ensemble members that are forced by independent realizations of the stochastic (weatherlike) part of the surface fluxes. Despite the independent forcings and the crude parameterization of the atmospheric response to SSTs, the filter converges very well to the observed thermal structure of the ocean. All warm and cold events during 1980–2002, and the corresponding subsurface thermal and current structures, are reconstructed by the assimilation. Compared to the 3DVAR analysis, the ensemble filter produces a smoother solution that is more consistent with the observations, presumably due to the filter’s incorporation of temporally and spatially varying background covariances (Zhang and Anderson

2003). The MPPEAKF solution also provides a better initialization than the 3DVAR, judging from the improvement in forecast skill. Moreover, the ensemble filter has a potential to provide an estimate of the analysis uncertainty, which is not available through other approaches.

Because of the requirement of evaluating the prior distribution (ensemble model integration), typically an ensemble filter is much more expensive than 3DVAR. In this six-member case, the computational cost of the MPPEAKF is however roughly 4 times more than 3DVAR owing to the cost of minimization (three iterations in this case).

An improved model and a better estimate of the forcing uncertainties will likely improve the filter performance. Currently the observed temperatures directly impact only the model temperature—yet since the cross covariances among state variables are available through the ensemble (Zhang and Anderson 2003), it is worth asking whether a multivariate assimilation (including salinity, surface height, and currents) could improve the ocean state estimate. We may also ask whether an anisotropic covariance structure (e.g., as estimated by the ensemble filter) or the state dependence background error covariance structure (Behringer et al. 1998) could help improve the less expensive 3DVAR assimilation in cases where the temporal variation of the error covariance is not important. Other interesting issues are how to incorporate the vertical correlation structure of the observations (Wu and Purser 2002) and how the ensemble size affects the assimilation. Looking beyond the simple hybrid model test bed described here, we plan to apply the MPPEAKF to an ocean GCM forced by all available observational flux products, and also a fully coupled ocean–atmosphere GCM, to provide improved initializations for coupled ENSO forecasts.

Acknowledgments. The authors thank Drs. S. Griffies and T. Ezer for their comments on earlier versions of this manuscript. Thanks go to Drs. Z. Liang and F. Zeng for their generous help in data processing and visualization. The authors thank two anonymous reviewers for their thorough examination and comments that were very useful for improving the manuscript.

REFERENCES

- Anderson, J. L., 2001: An ensemble adjustment Kalman filter for data assimilation. *Mon. Wea. Rev.*, **129**, 2884–2903.
- , 2003: A local least squares framework for ensemble filtering. *Mon. Wea. Rev.*, **131**, 634–642.
- Behringer, D. W., M. Ji, and A. Leetmaa, 1998: An improved coupled model for ENSO prediction and implications for

- ocean initialization. Part I: The ocean data assimilation system. *Mon. Wea. Rev.*, **126**, 1013–1021.
- Bishop, C. H., B. J. Etherton, and S. Majumdar, 2001: Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Mon. Wea. Rev.*, **129**, 420–436.
- Derber, J., and A. Rosati, 1989: A global oceanic data assimilation system. *J. Phys. Oceanogr.*, **19**, 1333–1347.
- Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, **99**, 10 143–10 162.
- Ezer, T., and G. L. Mellor, 1994: Continuous assimilation of Geosat altimeter data into a three-dimensional primitive equation Gulf Stream model. *J. Phys. Oceanogr.*, **24**, 832–847.
- Fedorov, A. V., and S. G. Philander, 2001: Is El Niño changing? *Science*, **288**, 1997–2001.
- Fukumori, I., 2002: A partitioned Kalman filter and smoother. *Mon. Wea. Rev.*, **130**, 1370–1383.
- Galanti, E., E. Tziperman, M. J. Harrison, A. Rosati, and Z. Sirkes, 2003: A study of ENSO prediction using a hybrid coupled model and the adjoint method for data assimilation. *Mon. Wea. Rev.*, **131**, 2748–2764.
- Gaspari, G., and S. E. Cohn, 1999: Construction of correlation functions in two and three dimensions. *Quart. J. Roy. Meteor. Soc.*, **125**, 723–757.
- Gill, P. E., W. Murray, and M. H. Wright, 1981: *Practical Optimization*. Academic Press, 401 pp.
- Glantz, M. H., R. W. Katz, and N. Nicholls, Eds., 1991: *Teleconnections Linking Worldwide Climate Anomalies: Scientific Basis and Societal Impact*. Cambridge University Press, 535 pp.
- Griffies, S. M., and R. W. Hallberg, 2000: Biharmonic friction with a Smagorinsky-like viscosity for use in large-scale eddy-permitting ocean models. *Mon. Wea. Rev.*, **128**, 2935–2946.
- , M. J. Harrison, R. C. Pacanowski, and A. Rosati, 2003: A technical guide to MOM4. GFDL Ocean Group Tech. Rep. 5, NOAA/Geophysical Fluid Dynamics Laboratory, Princeton, NJ, 295 pp.
- Hamill, T. M., J. S. Whitaker, and C. Snyder, 2001: Distance-dependent filtering of background error covariance estimates in an ensemble Kalman filter. *Mon. Wea. Rev.*, **129**, 2776–2790.
- Harrison, M. J., A. Rosati, R. Gudgel, and J. Anderson, 1996: Initialization of coupled model forecasts using an improved ocean data assimilation system. Preprints, *11th Conf. on Numerical Weather Prediction*, Norfolk, VA, Amer. Meteor. Soc., 7.
- , —, B. J. Soden, E. Galanti, and E. Tziperman, 2002: An evaluation of air–sea flux products for ENSO simulation and prediction. *Mon. Wea. Rev.*, **130**, 723–732.
- Holland, W. R., J. C. Chow, and F. O. Bryan, 1998: Application of a third-order upwind scheme in the NCAR ocean model. *J. Climate*, **11**, 1487–1493.
- Houtekamer, P. L., and H. L. Mitchell, 1998: Data assimilation using an ensemble Kalman filter technique. *Mon. Wea. Rev.*, **126**, 796–811.
- , and —, 2001: A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Wea. Rev.*, **129**, 123–137.
- Jazwinski, A. H., 1970: *Stochastic Processes and Filtering Theory*. Academic Press, 376 pp.
- Kalnay, E., 2003: *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, 341 pp.
- Large, W. G., J. C. McWilliams, and S. C. Doney, 1994: Oceanic vertical mixing: A review and a model with a nonlocal boundary layer parameterization. *Rev. Geophys.*, **32**, 363–403.
- Levitus, S., and T. P. Boyer, 1994: *Temperature*. Vol. 4, *World Ocean Atlas 1994*, NOAA Atlas NESDIS 4, 117 pp.
- Lindzen, R. S., and S. Nigam, 1987: On the role of sea surface temperature gradients in forcing low-level winds and convergence in the Tropics. *J. Atmos. Sci.*, **44**, 2418–2436.
- Navon, I. M., and D. M. Legler, 1987: Conjugate-gradient methods for large-scale minimization in meteorology. *Mon. Wea. Rev.*, **115**, 1479–1502.
- Neelin, J. D., 1991: The slow sea surface temperature mode and the fast-wave limit: Analytic theory for tropical interannual oscillations and experiments in a hybrid coupled model. *J. Atmos. Sci.*, **48**, 584–606.
- , D. S. Battisti, A. C. Hirst, F.-F. Jin, Y. Wakata, T. Yamagata, and S. Zebiak, 1998: ENSO theory. *J. Geophys. Res.*, **103**, 14 261–14 290.
- Ott, E., B. R. Hunt, I. Szunyogh, A. V. Zimin, E. J. Kostelich, M. Corazza, E. Kalnay, and D. J. Patil, 2004: A local ensemble Kalman filter for atmospheric data assimilation. *Tellus*, **56A**, 415–428.
- Rasmusson, E. M., and J. Wallace, 1983: Meteorological aspects of the El Niño/Southern Oscillation. *Science*, **222**, 1195–1202.
- Rosati, A., K. Miyakoda, and R. Gudgel, 1997: The impact of ocean initial conditions on ENSO forecasting with a coupled model. *Mon. Wea. Rev.*, **125**, 754–772.
- Suarez, M. J., and P. S. Schopf, 1988: A delayed action oscillatory for ENSO. *J. Atmos. Sci.*, **45**, 3283–3287.
- Sun, D.-Z., and Z. Liu, 1996: Dynamic ocean–atmosphere coupling: A thermostat for the Tropics. *Science*, **272**, 1148–1150.
- Sweeney, C., A. Gnanadesikan, S. M. Griffies, M. J. Harrison, A. J. Rosati, and B. L. Samuels, 2005: Impacts of shortwave penetration depth on large-scale ocean circulation and heat transport. *J. Phys. Oceanogr.*, **35**, 1103–1119.
- Tippett, M. K., J. L. Anderson, C. H. Bishop, T. M. Hamill, and J. S. Whitaker, 2003: Ensemble square root filters. *Mon. Wea. Rev.*, **131**, 1485–1490.
- Weaver, A. T., J. Vialard, and B. L. T. Anderson, 2003: Three- and four-dimensional variational assimilation with a general circulation model of the tropical Pacific Ocean. Part I: Formulation, internal diagnostics and consistency checks. *Mon. Wea. Rev.*, **131**, 1360–1378.
- Whitaker, J. S., and T. M. Hamill, 2002: Ensemble data assimilation without perturbed observations. *Mon. Wea. Rev.*, **130**, 1913–1924.
- Wittenberg, A. T., 2002: ENSO response to altered climates. Ph.D. thesis, Princeton University, 475 pp.
- , 2004: Extended wind stress analyses for ENSO. *J. Climate*, **17**, 2526–2540.
- Wu, W. S., and R. J. Purser, 2002: Three-dimensional variational analysis with spatially inhomogeneous covariance. *Mon. Wea. Rev.*, **130**, 2905–2916.
- Zebiak, S. E., and M. A. Cane, 1987: A model El Niño–Southern Oscillation. *Mon. Wea. Rev.*, **115**, 2262–2278.
- Zhang, S., and J. L. Anderson, 2003: Impact of spatially and temporally varying estimates of error covariance on assimilation in a simple atmospheric model. *Tellus*, **55A**, 126–147.
- , —, A. Rosati, M. J. Harrison, S. P. Khare, and A. Wittenberg, 2004: Multiple time level adjustment for data assimilation. *Tellus*, **56A**, 2–15.