TECHNICAL MEMORANDUM

827

Implicit and explicit cross-correlations in coupled data assimilation

P. Laloyaux[1], S. Frolov[2], B. Ménétrier[3] and M. Bonavita[1]

Research Department

[1]ECMWF, Reading, United Kingdom [2]Naval Research Laboratory, Monterey, United States [3]CNRM UMR 3589, Météo-France/CNRS, Toulouse, France

Published in Quarterly Journal of the Royal Meteorological Society

September 7, 2018

European Centre for Medium-Range Weather Forecasts
Europäisches Zentrum für mittelfristige Wettervorhersage
Centre européen pour les prévisions météorologiques à moyen terme

Series: ECMWF Technical Memoranda

A full list of ECMWF Publications can be found on our web site under:
http://www.ecmwf.int/en/research/publications

Contact: library@ecmwf.int

European Centre for Medium-Range Weather Forecasts
Shinfield Park, Reading, RG2 9AX, England

Abstract

The European Centre for Medium-Range Weather Forecasts has recently developed an implicit coupling approach in the CERA system, where atmospheric 4D-Var and ocean 3D-Var are synchronized using multiple outer iterations in the incremental variational formulation. Since this original work on the outer loop coupling, it has been unclear just how closely does the CERA system with implicit coupled covariances approximates a strongly coupled DA with ensemble based coupled covariances.

A series of single-observation experiments in a perfect twin framework have been conducted to evaluate the effectiveness and the efficiency of the outer loop coupling and the amplitude of the implicit cross-correlations between atmospheric and ocean temperature. They are compared against a coupled assimilation system that explicitly specifies cross-correlations from an ensemble of coupled forecasts. We find that both the implicit outer loop and the explicit ensemble-based methods produce equally skillful estimates of the coupled state in the middle of a 24 hour assimilation window. Our estimates of the rate of the outer loop convergence suggest that atmospheric and ocean states synchronise within the first 6 to 10 hours of the assimilation window. We conclude, that the outer loop coupling is effective when the window is *long* enough that original imbalances in the atmospheric and ocean increments can synchronise within the length of the assimilation window. On the other hand, we suggest that explicit coupling is preferable for data assimilation systems with short assimilation windows (e.g. 6 hours or less).

1 Introduction

Coupled Data Assimilation (CDA) aims at providing a consistent state estimate of the coupled model, in which observations in one fluid (e.g. ocean surface temperature observations) can correct for errors in the coupled fluid (e.g. the state of the atmosphere). Laloyaux, Balmaseda, Dee, Mogensen & Janssen (2016) proposed to use the outer loop of the atmospheric and oceanic incremental variational formulation as a way to synchronise the information between two otherwise uncoupled assimilation systems. This aims to produce a more consistent and better balanced analysis between the atmosphere and the ocean. In this outer loop coupling, the nonlinear trajectories are produced by the coupled model but the increments are computed independently for the ocean and the atmosphere. If the outer loop is performed only once, assimilation is considered weakly coupled as the coupled model is only used to produce the first-guess trajectory and the potentially unbalanced initial state is only synchronised during the model forecast (Smith et al. 2015). However, if more than one outer loop iteration is used, then some degree of coupling between atmosphere and ocean increments is achieved.

The European Centre for Medium-Range Weather Forecasts (ECMWF) has implemented the coupled model constraint at the outer loop level of the incremental variational formulation as a part of the coupled atmosphere-ocean assimilation system (CERA) (Laloyaux, Balmaseda, Dee, Mogensen & Janssen 2016). This system has been used to produce two different reanalyses at ECMWF which include the atmosphere, ocean, land, ocean waves and sea ice components of the Earth system. The first reanalysis (CERA-20C) reconstructs the past climate and weather over the 20th century (1901-2010) focusing on low-frequency climate variability (Laloyaux et al. 2018). CERA-20C reanalysis does not assimilate the full observing system at any time, but only surface pressure and marine wind observations in the atmosphere, as well as subsurface temperature and salinity profiles in the ocean. The atmosphere has a 125km horizontal grid resolution with 91 vertical levels, from near the surface to 0.01 hPa. The ocean has a one-degree horizontal grid with 42 vertical levels going down to 5350 m with a first layer thickness of 10 meters. The second reanalysis (CERA-SAT) is the second reanalysis dataset spanning 8 years between 2008 and 2016 (Schepers 2018). It is a proof-of-concept for a coupled reanalysis which assimilates the full observing system available in the modern satellite age. Based on a 65km atmospheric

grid and a quarter of a degree ocean grid (eddy-permitting) with a first layer thickness of one meter, CERA-SAT also evaluates the impact of higher resolution to represent correctly the coupled processes in the atmosphere-ocean analysis.

Besides the reanalysis activities, ECMWF has started to explore the potential of coupled assimilation to initialise the numerical weather forecast. Such an approach has the potential to make better use of satellite measurements and to improve the quality of the forecasts. For example, it has been found that the ocean temperature in the mixed layer is improved in the CERA system when scatterometer data are assimilated, while the impact is neutral in an uncoupled system (Laloyaux, Thépaut & Dee 2016). Coupled assimilation should generate a reduction of initialisation shocks in coupled forecasts by better accounting for interactions between the components. It should also lead to the generation of a consistent Earth-system state for the initialisation of forecasts across all timescales. Some work has been already done in that direction and it has been found that forecasts initialised by separate oceanic and atmospheric analyses do exhibit initialisation shocks in lower atmospheric temperature, when compared to forecasts initialised using the CERA system. Changes in the ocean component can lead to sea surface temperature shocks of more than 0.5K in some equatorial regions during the first day of the forecast (Mulholland et al. 2015). A similar assessment has been conducted at the Met Office with their coupled assimilation system where the coupled model is used to compute the first-guess trajectory (Lea et al. 2015).

We can consider that an outer-loop coupling system is an approximation of a fully coupled 4D-Var system which tries to find an approximation to the same optimal solution by setting the coupled adjoint model and the cross-fluid error covariance at the initial time of the assimilation window to zero. However, outer loop coupling is not the only way to implement a coupled assimilation system. It is possible to relax this approximation by specifying a coupled background error covariance matrix which includes off-diagonal terms to represent the cross-correlation between the atmosphere and the ocean. It is possible to specify this cross covariance using an ensemble of coupled state which samples the coupled flow-dependent background error distribution. However the limited ensemble size will introduce sampling noise and covariance localisation needs to be applied. This approach shows potential for coupled assimilation, but it involves a lot of scientific and technical challenges to be implemented with an operational forecasting model (Frolov et al. 2016, Sluka et al. 2016, Wada & Kunii 2017).

Since the original work of Laloyaux, Balmaseda, Dee, Mogensen & Janssen (2016) on the CERA system, it has been unclear just how much synchronisation does the outer loop achieve or, in other words, how closely does the outer loop coupling approximate a coupled assimilation system where forecast error covariances between the two systems are explicitly modelled. To address this question, a series of single observation experiments in a simplified framework have been conducted assimilating a Sea Surface Temperature (SST) measurement. The impact of the ocean temperature assimilation on the reduction of the air temperature error has been evaluated in CERA and in an Ensemble Kalman Filter that models the coupled covariances explicitly. Different locations in the Pacific ocean and different dates have been selected to study the sensitivity of cross-correlations to the ocean mixed layer depth.

The structure of the paper is as follows. The mathematical formulation of data assimilation and the importance of the background error covariance matrix is reviewed in Section 2. The CERA algorithm is reviewed in Section 3 and the impact of the implicit cross-correlations generated by the outer loop coupling on the convergence of the atmospheric and oceanic variables is illustrated. Atmospheric and ocean increments synchronise during the nonlinear model integrations computed in the CERA algorithm. Section 4 shows how an ensemble of coupled forecasts can be used to estimate explicit cross-correlation and the importance of a proper localisation technique. The series of single observation experiments used to compare the implicit and explicit cross-correlation are described in Section 5. Section 6 shows how closely CERA approximate the state estimate obtained with an explicitly modelled ocean-atmosphere

error covariance. The speed of the response of the atmosphere to changes in the ocean temperature is studied for the CERA system in Section 7. We conclude by discussing the deficiencies of the assimilation methods that we encountered during this study and by discussing how our findings can be applied to coupled data assimilation systems currently in development at other operational centrers.

2 Background error covariance in data assimilation

Data assimilation is a procedure, in which a model integration is adjusted on the basis of actual observations. Many assimilation methods are similar in that observations, $\mathbf{y}$, are combined with a model-generated state estimate, $\mathbf{x}^b$, by minimising a functional

$$J(\mathbf{x}) = (\mathbf{x}^b - \mathbf{x})^{\mathrm{T}} \mathbf{B}^{-1} (\mathbf{x}^b - \mathbf{x})^{\mathrm{T}} + (\mathbf{y} - \mathscr{H}(\mathbf{x}))^{\mathrm{T}} \mathbf{R}^{-1} (\mathbf{y} - \mathscr{H}(\mathbf{x}))$$

with respect to the model state $\mathbf{x}$ (Daley 1991, Kalnay 2003, Rabier 2005). The function $\mathscr{H}$ represents the relationship between model state and observations. By incorporating in $\mathscr{H}$ a forecast model integration and using observations distributed over a time window, (2) becomes a four-dimensional cost function (4D-Var). Otherwise, (2) is a three-dimensional cost function (3D-Var). The matrices $\mathbf{B}$ and $\mathbf{R}$ represent the covariance operators associated with the background and the observation errors, respectively. These covariance matrices are actually operators that are applied, and never stored explicitly in either variational or ensemble DA schemes. The solution $\mathbf{x}^a$ which minimises the functional (2) satisfies the nonlinear equation

$$\mathbf{x}^a = \mathbf{x}^b + \mathbf{B} \left(\frac{\partial \mathscr{H}}{\partial \mathbf{x}} \right)^{\mathrm{T}}_{\mathbf{x}=\mathbf{x}^a} \mathbf{R}^{-1} (\mathbf{y} - \mathscr{H}(\mathbf{x}^a)) \tag{1}$$

obtained by setting the gradient of $J(\mathbf{x})$ to zero (Dee 2005). Even though it is not possible to evaluate the solution $\mathbf{x}^a$ from this nonlinear equation for operational models, it shows that the correction of the background $\mathbf{x}^b$ lies in the subspace spanned by $\mathbf{B}$ as it is the last operator that appears in the computation of the analysis increment in (1). $\mathbf{B}$ has a correlation part $\mathbf{C}$ (a non-diagonal matrix of correlations between the elements of $\mathbf{x}$) and a variance part Σ^2 (a diagonal matrix of variances of the elements of $\mathbf{x}$). The covariance matrix $\mathbf{B}$ is formed by multiplying respective columns and rows of $\mathbf{C}$ by the square roots of the variances

$$\mathbf{B} = \Sigma \mathbf{C} \Sigma^T \tag{2}$$

The correlation matrix $\mathbf{C}$ spreads out information in space (vertically and horizontally) and to other variables. This effect can be easily seen in a single observation experiment which directly observes one element of the state vector. In this case, the observation operator $\mathscr{H}$ becomes a row vector $\mathbf{H}$ of zeroes apart from the element of the state vector that is being observed (k-th element), which has value 1. The observation vector $\mathbf{y}$ contains only one measurement y and its error variance $\mathbf{R}$ is a scalar noted σ_o^2. With a linear operator $\mathbf{H}$, solving the equation (1) is equivalent to solving the Kalman filter equation (Gratton et al. 2011) and the solution of (2) is given by

$$\mathbf{x}^a = \mathbf{x}^b + \mathbf{B}\mathbf{H}^{\mathrm{T}} \left(\mathbf{H}\mathbf{B}\mathbf{H}^{\mathrm{T}} + \mathbf{R} \right)^{-1} \left(\mathbf{y} - \mathbf{H}\mathbf{x}^b \right) \tag{3}$$

With a single measurement observing the k-th element of $\mathbf{x}$, equation (3) can be written for each l-th element of the state vector $\mathbf{x}$ as:

$$x_l^a = x_l^b + \frac{\mathbf{B}_{lk}}{\mathbf{B}_{kk} + \sigma_o^2} \left(y - x_k^b \right) \tag{4}$$

where the scalar $\mathbf{B}_{kk}$ is the background error variance of the element k and the scalar $\mathbf{B}_{lk}$ is the background error covariance between the element k and l. Even though the observation has been made of element k,

the information has been spread to the other position l with a magnitude proportional to the covariance value $\mathbf{B}_{lk}$ (Bannister 2008).

One possibility to implement coupled data assimilation is to treat the atmosphere and ocean as a single system minimising the functional (2) where the state vector $\mathbf{x}$ contains the atmospheric and ocean states. In this approach, the full coupled $\mathbf{B}$ will include off-diagonal terms to represent the cross-covariances between the atmosphere and the ocean. This will ensure the transfer of information between the two components in the analysis, such as an ocean observation has immediate impact on the atmospheric state estimate and, conversely, assimilation of an atmospheric observation affects the ocean state. This approach presents significant scientific challenges as the cross-correlation between the different geophysical variables at any grid point needs to be estimated (Frolov et al. 2016, Smith et al. 2017, 2018). It also involves a lot of technical developments to merge together the different elements of the atmospheric and ocean assimilation systems which have been typically developed separately for many years. For these reasons, ECMWF has developed a coupled atmosphere-ocean assimilation system called CERA which keeps separate atmosphere and ocean background errors but introduces the coupling via the forecast model integrations computed at the outer-loop level of the incremental variational approach.

3 Correlations and convergence in CERA

ECMWF has been developing a coupled atmosphere-ocean model to produce medium-range, monthly and seasonal forecasts. The coupled model relies on the Integrated Forecast System (IFS) which is an atmospheric model developed in cooperation with Météo-France (ECMWF 2016). It is coupled to several other Earth system components including the ocean model NEMO (Madec 2008). The atmospheric and ocean components are integrated into a single executable with a common time step loop, sequentially calling each component and regridding fields as needed on the different model grids (Mogensen, Keeley & Towers 2012). The CERA-20C resolution (125km horizontal atmospheric grid and 110km horizontal ocean grid) is used in all the experiment reported here (Laloyaux et al. 2018). The model for the evolution of the coupled atmosphere-ocean system can be represented by the equation:

$$\begin{bmatrix} \bar{\mathbf{x}}_{i+1} \\ \underline{\mathbf{x}}_{i+1} \end{bmatrix} = \mathscr{M}_{i+1,i} \begin{bmatrix} \bar{\mathbf{x}}_{i} \\ \underline{\mathbf{x}}_{i} \end{bmatrix} \tag{5}$$

where $\mathscr{M}_{i+1,i}$ is the integration of the nonlinear coupled forecast model between time t_i and t_{i+1} from the atmospheric state $\bar{\mathbf{x}}_i$ and the ocean state $\underline{\mathbf{x}}_i$. A bar above a character is used to denote the upper fluid (i.e. the atmosphere) and a bar below a character is used to denote the lower fluid (i.e. the ocean).

The CERA system is based on a common 24-hour assimilation window shared by the atmosphere and the ocean to ingest simultaneously atmospheric and ocean observations. It is built upon the incremental variational approach which consists of an iterative solution to (2) with a pair of nested loops (Courtier et al. 1994, Veerse & Thépaut 1998). The outer loop integrates the coupled nonlinear forecast model over the 24-hour assimilation window, producing a four-dimensional state estimate($\bar{\mathbf{x}}^k$,$\underline{\mathbf{x}}^k$) and the observation misfits ($\delta\bar{\mathbf{y}}^k$,$\delta\underline{\mathbf{y}}^k$). This is summarised in Algorithm 1 at lines 3 and 4. The inner loop minimises a linearised version of the variational cost function (2) for the control variable increments, using a preconditioned conjugate-gradient-like (CG) method (Tshimanga et al. 2008). This minimisation is done separately in CERA for the atmosphere and the ocean (line 5) using their original background error matrices ($\bar{\mathbf{B}}$ and $\underline{\mathbf{B}}$) and their original linearised observation operators ($\bar{\mathbf{H}}$ and $\underline{\mathbf{H}}$). For the atmosphere, a tangent linear approximation of the forecast model is included in $\bar{\mathbf{H}}$. In the ocean, the minimisation of the linearized cost function is computed by the NEMOVAR system. It is 3D-Var configuration where no ocean dynamic is taken into account in the operator $\underline{\mathbf{H}}$ (Mogensen, Balmaseda & Weaver 2012). Finally,

the initial conditions are updated with the increments (line 6) and the next outer iteration can start by integrating the coupled nonlinear model from the new initial conditions. The CERA system computes a fixed number of outer iterations (usually 2 or 3) and a final coupled model integration is performed (line 8) to produce the analysis trajectory valid over the 24-hour window.

Algorithm 1 CERA system

1: **Initialisation from the background**
$$\begin{bmatrix}\bar{\mathbf{x}}_0^0\\ \underline{\mathbf{x}}_0^0\end{bmatrix} = \begin{bmatrix}\bar{\mathbf{x}}^b\\ \underline{\mathbf{x}}^b\end{bmatrix}$$

2: **for k=0,1,...,N do**

3: **Integrate the nonlinear model**
$$\begin{bmatrix}\bar{\mathbf{x}}_{i+1}^k\\ \underline{\mathbf{x}}_{i+1}^k\end{bmatrix} = \mathscr{M}_{i+1,i}\begin{bmatrix}\bar{\mathbf{x}}_i^k\\ \underline{\mathbf{x}}_i^k\end{bmatrix}$$

4: **Compute the observation departures**
$$\begin{bmatrix}\delta\bar{\mathbf{y}}^k\\ \delta\underline{\mathbf{y}}^k\end{bmatrix} = \begin{bmatrix}\bar{\mathbf{y}}\\ \underline{\mathbf{y}}\end{bmatrix} - \begin{bmatrix}\bar{\mathscr{H}}\bar{\mathbf{x}}^k\\ \underline{\mathscr{H}}\underline{\mathbf{x}}^k\end{bmatrix}$$

5: **Minimise the linearised problems**
$$\bar{J}^k(\delta\bar{\mathbf{x}}) = \left(\delta\bar{\mathbf{x}} - \left(\bar{\mathbf{x}}^b - \bar{\mathbf{x}}_0^k\right)\right)^{\mathrm{T}}\bar{\mathbf{B}}^{-1}\left(\delta\bar{\mathbf{x}} - \left(\bar{\mathbf{x}}^b - \bar{\mathbf{x}}_0^k\right)\right) + \left(\bar{\mathbf{H}}^k\delta\bar{\mathbf{x}} - \delta\bar{\mathbf{y}}^k\right)^{\mathrm{T}}\bar{\mathbf{R}}^{-1}\left(\bar{\mathbf{H}}^k\delta\bar{\mathbf{x}} - \delta\bar{\mathbf{y}}^k\right)$$
$$\underline{J}^k(\delta\underline{\mathbf{x}}) = \left(\delta\underline{\mathbf{x}} - \left(\underline{\mathbf{x}}^b - \underline{\mathbf{x}}_0^k\right)\right)^{\mathrm{T}}\underline{\mathbf{B}}^{-1}\left(\delta\underline{\mathbf{x}} - \left(\underline{\mathbf{x}}^b - \underline{\mathbf{x}}_0^k\right)\right) + \left(\underline{\mathbf{H}}^k\delta\underline{\mathbf{x}} - \delta\underline{\mathbf{y}}^k\right)^{\mathrm{T}}\underline{\mathbf{R}}^{-1}\left(\underline{\mathbf{H}}^k\delta\underline{\mathbf{x}} - \delta\underline{\mathbf{y}}^k\right)$$

6: **Update initial condition**
$$\begin{bmatrix}\bar{\mathbf{x}}_0^{k+1}\\ \underline{\mathbf{x}}_0^{k+1}\end{bmatrix} = \begin{bmatrix}\bar{\mathbf{x}}_0^k\\ \underline{\mathbf{x}}_0^k\end{bmatrix} + \begin{bmatrix}\delta\bar{\mathbf{x}}^k\\ \delta\underline{\mathbf{x}}^k\end{bmatrix}$$

7: **end for**

8: **Produce the analysis trajectory**
$$\begin{bmatrix}\bar{\mathbf{x}}_{i+1}^N\\ \underline{\mathbf{x}}_{i+1}^N\end{bmatrix} = \mathscr{M}_{i+1,i}\begin{bmatrix}\bar{\mathbf{x}}_i^N\\ \underline{\mathbf{x}}_i^N\end{bmatrix}$$

In the CERA system, the transfer of information between the two components is not done through a coupled background error matrix, but using the coupled model as a constraint during the nonlinear model integrations. At every coupling time step during the nonlinear model integration, the ocean model transfers the sea surface temperature and the surface currents estimate to the atmosphere. These are used to update radiation (solar and non-solar) and evaporation minus precipitation fields which will be sent back to the ocean model. The coupled model constraint induces implicit cross-correlations between the atmosphere and the ocean within the assimilation process which ensures a consistent estimate of the coupled state. By setting the gradient of the linearised problem $J^k(\bar{\mathbf{x}})$ to zero, the optimal increment is given by

$$\delta\bar{\mathbf{x}}^k = \left(\bar{\mathbf{x}}^b - \bar{\mathbf{x}}_0^k\right) + \bar{\mathbf{B}}\bar{\mathbf{H}}^{k\mathrm{T}}\left(\bar{\mathbf{H}}^k\bar{\mathbf{B}}\bar{\mathbf{H}}^{k\mathrm{T}} + \bar{\mathbf{R}}\right)^{-1}\left(\delta\bar{\mathbf{y}}^k - \bar{\mathbf{H}}^k\left(\bar{\mathbf{x}}^b - \bar{\mathbf{x}}_0^k\right)\right)$$

where $\bar{\mathbf{x}}^b$ is the background and $\bar{\mathbf{x}}_0^k$ is the initial condition estimated at the k-th outer iteration. This equation shows that the atmospheric increment computed at each outer iteration depends on the observation misfit $\delta\bar{\mathbf{y}}^k$ which has been computed by the coupled model using the updated atmospheric and ocean initial condition ($\bar{\mathbf{x}}_0^k$ and $\underline{\mathbf{x}}_0^k$). The operator $\bar{\mathbf{H}}^k$ which linearises $\bar{\mathscr{H}}$ around $\bar{\mathbf{x}}^k$ also depends on the nonlinear model trajectory. This coupling is illustrated in Figure 1 where a single SST observation valid in the middle of the 24 hour window is assimilated. Figure 1 shows the temperature vertical profiles through

the fluid interface for the different CERA iterates at the observation time when two outer iterations are computed. Starting from the background, the CERA system produces a sequence of iterates to fit the ocean observation. The misfit with the observation is reduced in the ocean at each outer iteration and the error in the atmospheric state is decreased thanks to the implicit cross-correlation generated by the coupled model constraint.

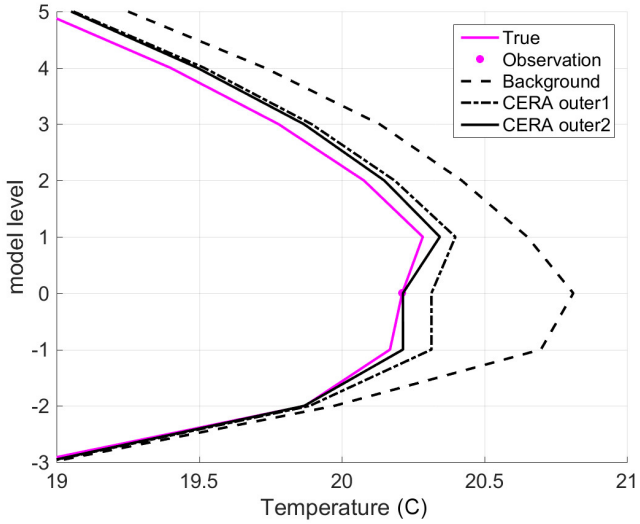

Figure 1: Temperature vertical profile showing the convergence of CERA iterates towards the true state (solid pink) when a single SST observation valid after 12 hours (pink dot) is assimilated. Starting from the background values, the CERA system produces a sequence of iterates to fit the ocean observation. Positive levels are in the atmosphere and negative levels are in the ocean, with the zero level being the ocean top layer where the observation is assimilated.

Due to computational budget and timeliness issues in reanalysis and NWP, the sequence of incremental problems has to be fixed (between 2 and 4 outer iterations). Each incremental problem is solved approximately, using an iterative conjugate-gradient-like (CG) method (Tshimanga et al. 2008) and performing only from 20 to 50 inner iterations. In Figure 1, the CERA algorithm converged after two outer iterations for the single observation experiment. In general, the convergence rate depends on the nonlinearity of the forecast model, the length of the assimilation window, and on the nonlinearity of the observations (Tremolet 2005). To illustrate this, the CERA system was run over a week, using 3 outer iterations and assimilating surface and subsurface conventional observations. Figure 2 shows the norm of the increment computed at each outer iteration normalised by the total increment. We found that the atmospheric variables of the coupled system require more outer iterations to converge (red lines in Figure 2). For example, the atmospheric heat and freshwater flux were still adjusted by CERA after the third outer iteration. The magnitude of the increments computed at the third outer iteration is still 30% of the total increment. This is not surprising as these variables depend on the whole atmospheric column and involve nonlinear processes. On the other hand, the deep ocean converged almost in one outer iteration (blue lines in Figure 2) as the magnitude of the ocean temperature increment at 500 meters changes by less than 10% at the second and third iterations. Oceanic fields in the deep layer (at 500 m depth) changes very slow. They are almost stationary in 24 hours and atmospheric influence is also small. Thus, 3DVAR without outer loops is enough for analysis of the deep layers if short assimilation window like 24 hours is adopted.

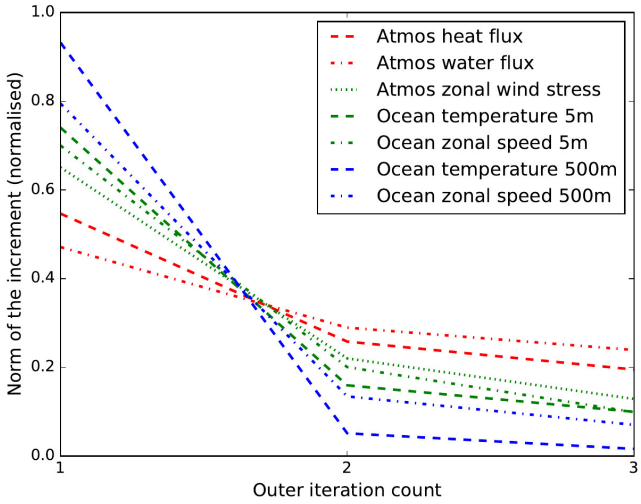

Figure 2: Norm of the increment computed at each outer iteration of the CERA system normalised by the total increment. Surface and subsurface conventional observations have been assimilated in this experiment.

4 Correlations in coupled ensemble forecast

Besides the CERA system, another possibility to implement a coupled data assimilation is to treat the two components as a single system minimising the functional (2) where the state vector **x** contains the atmospheric and ocean geophysical variables (Nozomi et al. 2008, Smith et al. 2015). This approach allows the introduction of an estimate of the fully coupled background error covariance matrix **B** which includes off-diagonal terms to explicitly represent the cross-correlation between the atmosphere and the ocean. These terms can be estimated from an ensemble of coupled forecasts which sample the coupled flow-dependent background error distribution. Having such an ensemble is challenging as only small size ensembles are affordable in operational environments.

An ensemble of such coupled models is available as a part of the CERA-20C reanalysis and was generated using the Ensemble of Data Assimilation (EDA) technique (Bonavita et al. 2016). The entire length of the CERA-20C record has 10 members but for the February 2005 and August 2005 that were used in this study, we increased the number of members to 25. In each member, observation are perturbed and different sea surface temperature and model physics are used. Observation errors are represented by perturbations sampled from the assumed observation error covariance matrix **R**. The sea surface temperature fields are not directly perturbed but different realisations are provided as part of the HadISST2 product (Titchner & Rayner 2014). Note that the 10 different available members in HadISST2 are used several times. The model physics is perturbed in a stochastic way, adding perturbations to the physical tendencies to simulate the effect of random errors in the physical parameterisations (Buizza et al. 1999, Palmer et al. 2009).

Figure 3 shows the correlation map between SST and surface air temperature among ensemble members averaged over February 2005 (a) and August 2005 (c). It confirms seasonal and regional patterns of correlations between air temperatures and SST that were first documented in (Feng et al. 2018). Specifically, we see year-round strong correlation in the Tropical East Pacific (TEPAC) region and weak correlations in the tropical West Pacific. Feng et al. (2018) attribute this pattern to persistent shallowing of the mixed layer depth in the Tropical East Pacific region and presence of strong convective weather in the Warm Pool area of the Western Pacific that disrupts the ocean-atmospheric correlations. Similar to Feng et al. (2018), we also observe a strong seasonal cycle of correlations in the mid-latitudes, where correlations

are high in the local summer when the mixed layer is shallow and the correlations are low in winter when the mixed layer is deep. However, the magnitude of the correlation in our studies are lower than in Feng et al. (2018) because our correlations are instantaneous (or averages of instant correlations) while they computed correlations of monthly averaged perturbations. We believe that their estimates are unrealistic for use in coupled data assimilation. Additionally, the panels (b) and (d) of Figure 2 show that instantaneous correlations have local variations and high temporal variability, due to intermittent interactions of ocean and atmospheric features and due to ensemble sampling errors.

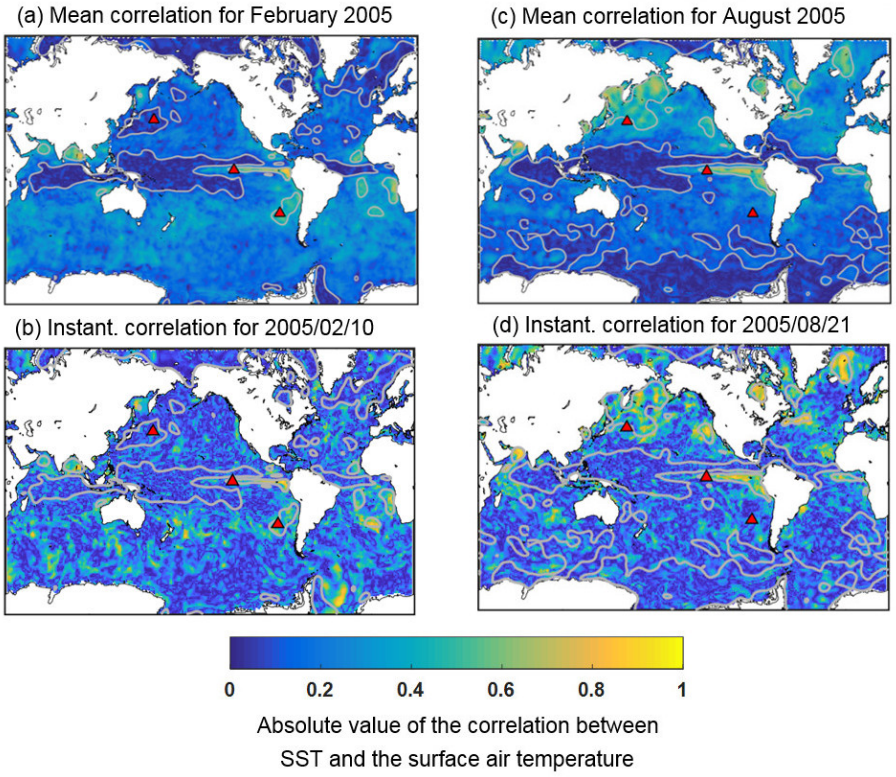

Figure 3: Absolute value of the correlation between SST and surface air temperature among ensemble members. (a) Average of daily correlations for February 2005; (b) daily correlation for 2005/02/10; (c) average of daily correlations for August 2005; (d) daily correlation for 2005/08/21. The isolines of 0.1 and 0.4 monthly average correlations are shown in gray. Red triangles show locations of the single observation experiments.

4.1 Ensemble localization

As the ensemble size used to estimate the cross-correlation is limited, techniques are needed to filter the sampling noise. The most popular is covariance localisation, which relies on a Schur (i.e., element by element) product of the sample covariance matrix with a positive definite correlation matrix. Merging the theories of linear filtering and covariance sampling, Ménétrier et al. (2015*a*,*b*) derived a method to diagnose localisation functions that reduce the sampling noise present in any ensemble-based covariance matrix, while preserving the signal of interest. The left panel of Figure 4 shows the vertical cross-correlations between SST and the atmosphere-ocean column at the location (130°W, 0°N) on 21st August 2005. The raw background error correlations based on the 25 member ensemble forecasts are in blue, while the localised background error correlations from Ménétrier et al. (2015*a*,*b*) are in red. The localisation function has been plotted separately on the right panel of Figure 4. The raw ensemble corre-

lations were able to estimate the correlations between the temperatures in the coupled ocean-atmosphere state. Specifically, we found a moderate correlation of 0.4 between SST and the lower levels of the atmosphere. However, the raw correlations appear noisy because of the limited ensemble size (see for example the correlations between SST and temperature errors in the upper atmosphere). The spurious correlations appear to be strongly attenuated with the localisation function but not fully removed in the resulting localised error correlations (see for example, the noise in the correlation between SST and the deep ocean below 1000 meters). Indeed, the localisation function diagnosed from Ménétrier et al. (2015*a*,*b*) is also noisy because of the limited ensemble size (25 members). To fully remove these spurious correlations, several solutions could be considered and combined:

- averaging the localisation diagnostic over several cycles to get a cleaner localisation function,
- fitting the localisation diagnostic with a smooth parametric function, for example using the static coupled localization of (Frolov et al. 2016),
- developing a hybrid localisation method that would use the extended localisation theory of Ménétrier & Auligne (2015) to linearly combine a localised ensemble-base covariance matrix with a climatological background error covariance matrix.

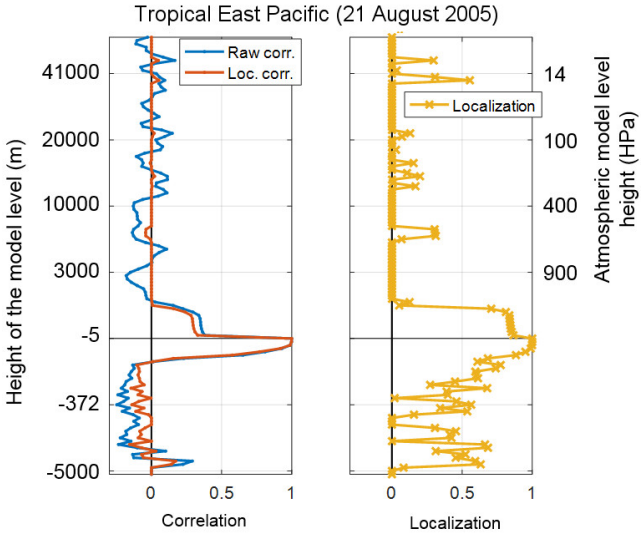

Figure 4: Correlations and localisation between SST and the atmosphere-ocean column at the location (130°W, 0°N) on 21st August 2005 (top) and on . Raw (blue) and localised (red) correlations are plotted on the left panel while the localisation function is plotted on the right panel. Correlations are plotted on model levels, which are then converted to either to height in meters (y-axis on the left) or into HPa for atmospheric levels (y-axis on the right).

When we examined the cross-correlation between SST and the atmosphere-ocean column, we found support for the relationship between the depth of the mixed layer and the coupling of errors in the atmosphere-ocean state. In the TEPAC region, and in the North Pacific (NPAC) and South Pacific (SPAC) regions at local summer-time, we found shallow mixed layer depth (about 50 metres in TEPAC and 25 metres in NPAC and SPAC) that translates to moderate correlations of SST errors and air temperature errors in the atmospheric boundary layer. The average correlation between SST and the first atmospheric layer was about 0.25 for the shallow mixed layer depth cases. However, the mixed layer depth was about 200 metres in the NPAC and 100 metres in the SPAC regions at local winter-time. This results in zero

correlation between SST and the atmospheric boundary layer (Figure 5). It is important to note that the raw ensemble correlations were in fact non-zero, but these non-zero sample correlations were completely obviated when the localisation function was applied.

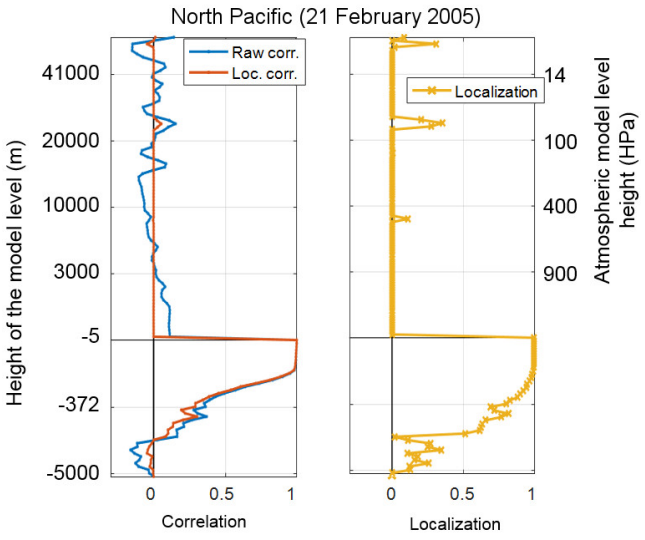

Figure 5: Same as in Figure 4 but for the location (160°E, 40°N) on 21st February 2005.

While the vertical profiles presented in Figures 4 and 5 focus only on the correlations between SST and atmosphere-ocean temperature, Figure 6 shows the full coupled vertical background correlation of temperature on the 21st August 2005 for the location (130°W, 0°N) when it is estimated from the raw ensemble (a) and localised (b). In the localised background error, temperature correlations between the atmosphere and the ocean are visible (see black arrows), as well as temperature correlation in the atmospheric surface boundary layer and the shallow ocean mixed layer.

5 Configuration of the single observation experiments

To compare the effectiveness of the implicit cross-correlation induced by the CERA system and presented in Section 3 with the explicit cross-correlation found in the ensemble of coupled forecast in Section 4, a series of single observation experiments has been conducted. The atmosphere-ocean forecast model is using a one hour coupling frequency. The resolution is set as in the CERA-20C reanalysis (125km horizontal atmospheric grid and 110km horizontal ocean grid). A perfect twin experiment setup is used to evaluate the error reduction. One member of the EDA system is assumed to be the truth run and another member of the EDA system is assumed to be the background. The SST measurement is taken from the truth after 12 hours and at 5 meter depth. No observation noise has been added and the standard deviation of the observation error ($\sigma_o = 0.01$) is significantly smaller than the one of the background error ($\sigma_b = 0.2$) which has been derived from the EDA system. This framework allows to see how the assimilation of the single SST observation will correct the background and reduce the error from the truth state.

Following Feng et al. (2018), we select three locations to conduct the single observation experiments. Each location is within a distinct area of average atmosphere-ocean coupling (Figure 2). The Tropical East Pacific (TEPAC) location (130°W, 0°N) is in the region of strong coupling. The North Pacific

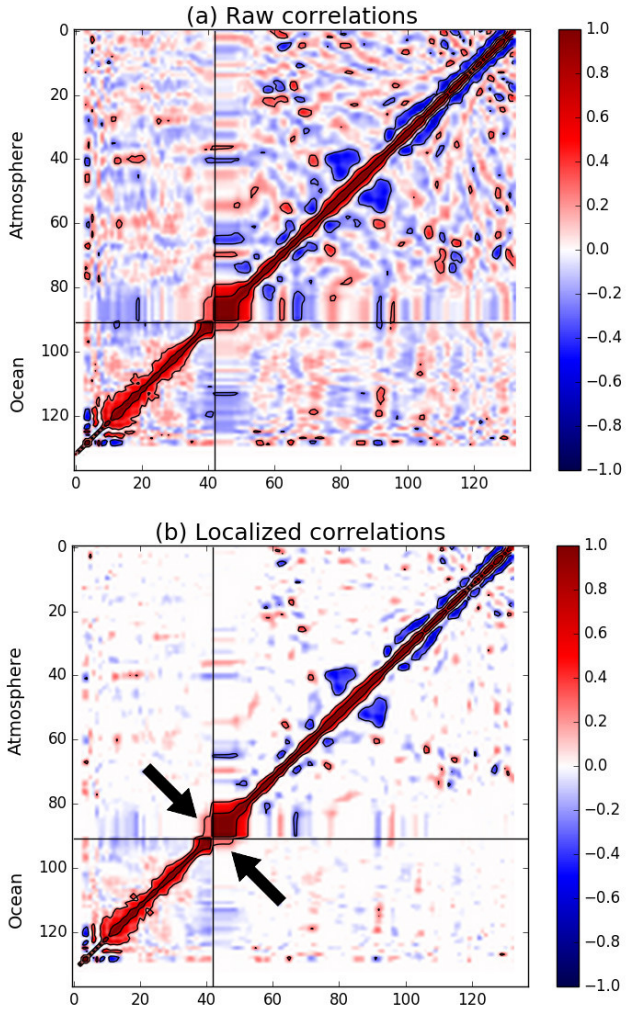

Figure 6: Temperature coupled background error on the 21st August 2005 for the location (130°W, 0°N) from the raw ensemble (a) and localised (b).

(NPAC) and the South Pacific (SPAC) locations are in mid-latitudes that exhibit seasonal cycle of coupling: one in the Kuroshio extension region (160°E, 40°N) and one in the South East Pacific (90°W, 30°S). For each location, we have selected three dates in February 2005 and three dates in August 2005 such as we have 18 single observation experiments in total.

To assess the effectiveness of the explicit cross-correlations to transfer the information from the SST measurement to the ocean and the atmosphere, the Kalman filter update (3) has been computed at the observation time using the same background trajectory as in the CERA system and the two flow dependent coupled covariance matrices presented in Figure 6 which are computed from 25 EDA members and localized using Ménétrier et al. (2015*a*,*b*).. To make the implementation of the Ensemble Kalman filter feasible, the state $\mathbf{x}$ contains only a vertical one-dimensional column including the atmospheric and the ocean temperature. The analysis computed with the CERA system and with the explicit cross-correlations computed using the Kalman filter can be compared at the observation time. Note that the CERA system has access to the nonlinear coupled dynamics of the model which include horizontal processes, while the Ensemble Kalman filter is solved for a vertical one-dimensional column in the middle of the 24 hour CERA assimilation window.

6 Comparisons between explicit and implicit coupled estimates

We found that CERA produced reliable estimates of the coupled state as measured by the reduction of the surface air temperature error in response to the assimilation of the SST observation. This has been summarised in Figure 7 which shows the atmospheric surface temperature RMS error for the background (dark blue) and the CERA system (light blue) for different locations and seasons. As expected, the error reduction in CERA is larger when the mixed layer depth is shallow as more coupling occurs in this situation.

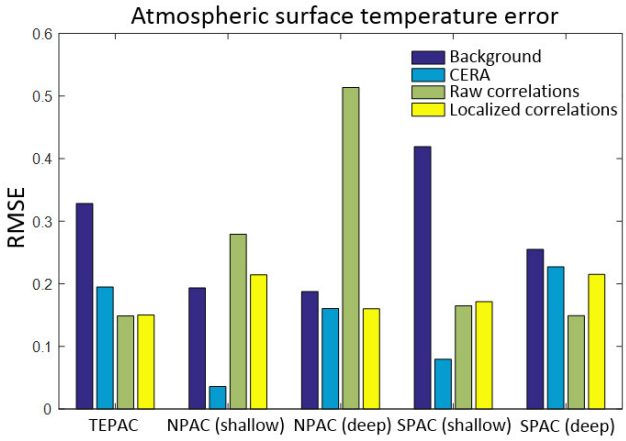

Figure 7: Atmospheric surface temperature RMS error in the middle of the assimilation window (after 12 hours). Errors are sorted by location and season (TEPAC all seasons, NPAC shallow, NPAC deep, SPAC shallow and SPAC deep) and by the type of the estimator (background (dark blue), CERA (light blue), raw explicit cross-correlation (green), and localised explicit cross-correlation (yellow)).

For the TEPAC region and the deep mixed layer regions, the errors in the CERA system is comparable to the error obtained with the explicit localised coupled background covariances (compare light blue and yellow bars). For example, the panel (a) of Figure 8 shows the results of one single observation experiment in the TEPAC region on the 31st August 2005. The ocean and atmospheric background errors are represented by the blue and red dotted lines respectively. The CERA analysis trajectory is plotted using star markers. The Kalman filter update has been computed at the observation time (after 12 hours) from the ocean and atmospheric background valid at that time. Square markers show the results with the raw ensemble correlations and the triangle marker with the localised correlations. Note that no blue triangle has been plotted as it will superpose on the blue square since the localisation value is equal to one at the top of the ocean (see Figure 4 and Figure 5). In response to the assimilation of the SST observation in CERA, the ocean SST was increased by 0.5 C (difference between the two blue lines). The air temperature also increased, achieving almost zero error at 12 hour mark when the SST observation was assimilated. The two Kalman filter estimates using the raw covariances (square) and localised (triangle) covariances agree closely with the CERA analysis.

For the shallow mixed layer regions, the error in the CERA system is smaller than the error with explicit background covariances (compare light blue, green and yellow bars in Figure 7). Looking at the panels (b) and (d) of Figure 8, the ocean analyses fit very well the SST observation with almost zero error for the different estimates. However, the raw ensemble correlation between SST and the atmosphere seems to be too strong and the response of the atmosphere to the correction of the SST error is too large. The SST was corrected by 0.5 C in panel (b) which also resulted in a 0.4 C correction in the air temperature while the

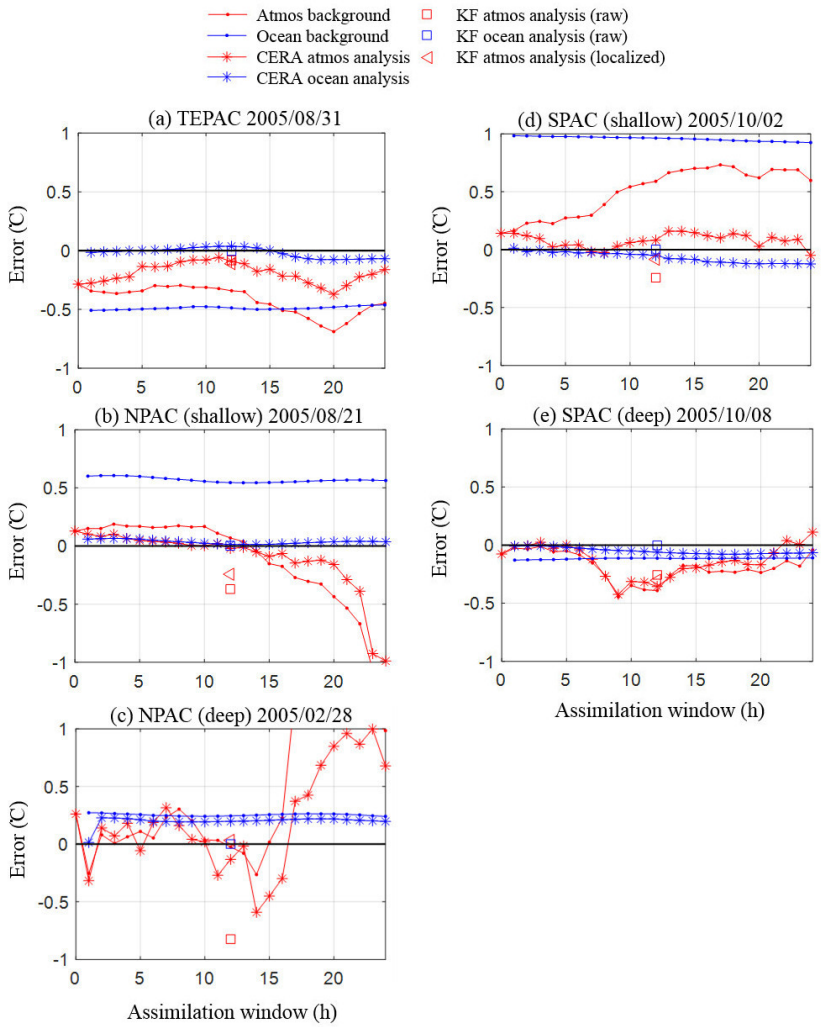

Figure 8: Background and analysis errors for the ocean and atmospheric surface temperature over the assimilation window. The ocean and atmospheric background errors are represented by the blue and red dotted lines respectively. The CERA analysis trajectory is plotted using star markers, the Kalman filter analysis valid at the observation time is plotted with a square markers (raw correlations) and a triangle marker (localised correlations).

optimal atmospheric response should be less than 0.1 C. Using localisation attenuates the overcorrection of the atmospheric errors due to the possible overestimate of the coupled correlations. The red triangle markers in panel (b) and (d) show that localised estimate have a smaller atmospheric error.

While the Kalman filter analyses fit closely the SST observations in all the experiments, the CERA system produces an estimate with a systematic larger error at the surface of the ocean when the mixed layer is deep. This situation is illustrated on the panel (c) of Figure 8 where the CERA ocean analysis error is 0.2 at the observation time. The reason of this behaviour is linked to the specification of the static ocean background covariance matrix used in the incremental variational formulation presented at line 5 of Algorithm 1 (Mogensen, Balmaseda & Weaver 2012). The vertical correlation between the ocean levels does not depend on the mixed layer depth which means that the information from a SST observation is not transferred to all the levels in the mixed layer. As a consequence, the increment produced at the beginning of the assimilation window is confined in the first few levels and will vanish in the first few time steps of the model integration. This undesirable effect is not due to the coupling strategy in CERA,

but to some flaws in the ocean background error specification. Work is ongoing to implement an ocean vertical correlation which depends on the mixed layer depth.

7 Timescales of synchronisation

The case studies presented in Figure 8 show that it takes several hours for CERA to synchronise the air temperature from a SST correction made at the beginning of the assimilation window. This can be seen by comparing the evolution of the atmospheric background (dotted red line) with the atmospheric analysis (red line with star markers). This response time of the atmospheric surface temperature increment to the assimilation of the SST observation has been illustrated in Figure 9 where the evolution of the size of the atmospheric increment normalised by the size of the SST increment at the same valid time is plotted. To characterise the timescale, we grouped TEPAC, shallow mid-latitudes, and deep mid-latitude cases together.

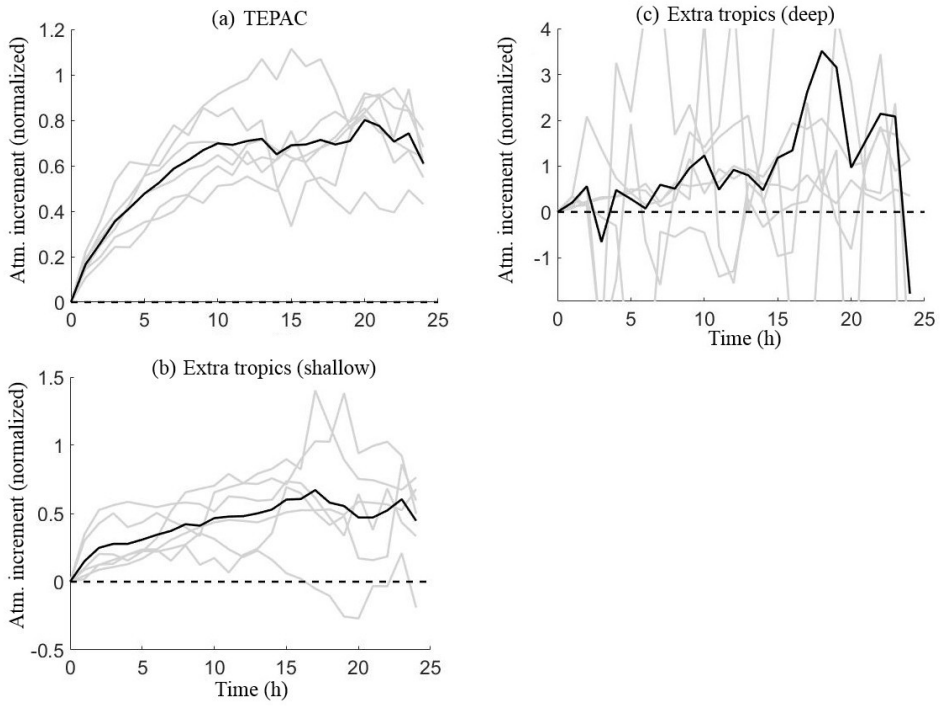

Figure 9: Time series of atmospheric surface temperature increment normalized by SST increment at the same valid time in the single observation experiments(Tropical East Pacific (a), shallow mixed layer cases for NPAC and SPAC (b), and deep mixed layer cases for NPAC and SPAC (c)). Plotted is the magnitude of the atmospheric increment in the CERA outer loop normalised by the magnitude of the SST perturbation (grey lines are the single observation experiments and black lines are the means).

When the cross-correlations are large (TEPAC and shallow mixed layer, panels (a) and (b) in Figure 9), the coupled model takes between 6 to 12 hours to transfer the information from the ocean increment to the atmosphere. At that point, the normalised atmospheric increment reaches a plateau which shows that the atmosphere and the ocean has reached some kind of consistency and balance. This transfer of information is faster in the TEPAC region (an average of 6 hours) and longer in mid-latitudes (an average of 12 hours). When the cross-correlations are small (deep mixed layer, panel (c) in Figure 9), the atmospheric response is not as clear. One reason is the ocean increment confined in the first few

levels when the mixed layer is deep. As the whole mixed layer column has not been corrected, the ocean increment will vanish after few time steps of the coupled model and the atmosphere does not have the time to adjust to it. This should be corrected by the implementation of ocean vertical correlation which depends on the mixed layer depth.

8 Summary and conclusions

This paper has studied the implicit cross-correlations generated by the CERA atmosphere-ocean coupled assimilation system developed at ECMWF. They have been compared to explicit cross-correlations computed from an ensemble of coupled forecast. The series of single observation experiments suggest that oceanic and atmospheric increments can be effectively synchronised by either the implicit coupled analysis based on CERA or by the explicit coupled analysis based on coupled ensemble error covariance. Despite the similarities measured in the middle of the 24 hour assimilation window where the SST single observation is valid, some differences exist.

We found that it took between 6 and 12 hours for the outer loop coupling to synchronise the coupled increments. In contrast, coupled error covariances can be specified at the beginning of the window, hence resulting in the synchronised initial conditions. This finding suggests that a long assimilation window (at least 12 hours) is necessary for CERA to be an effective strategy for coupled data assimilation. This finding is consistent with an interpretation that a long-window 4DVAR shifts the burden of specifying initial time covariances to the burden of providing accurate models of forward and backwards dynamics. For shorter 6-hour assimilation windows, the benefit of the outer loop strategy might be smaller and the use of a coupled initial time covariance more beneficial.

Based on our findings, we propose three strategies to alleviate the limited positive impact of the outer loop coupling in forecast systems that are embedded in a short (6 hour) operational analysis cycle. The first strategy takes place when the information from one assimilation cycle is transferred to the next cycle. In most operational centres, the analysis close to the end of the window is used as the background for the next assimilation cycle and the model is integrated from it to produce the so-called first-guess trajectory. To increase the length of the model integration and enhance the synchronisation of the increments, the analysis at the beginning of the previous assimilation window could be used as the initial condition to produce the first-guess trajectory. This will increase the integration time and should provide a more consistent and balanced model trajectory. This longer warm-up strategy can be further aided using the incremental analysis update (IAU) initialisation as currently used by the UK Met Office to balance the initial state of the atmospheric model before the next assimilation cycle (Lorenc et al. 2015).

Some operational centres use a short assimilation window as they need a new analysis every 6 hours to initialise a new forecast. The second strategy is to use longer overlapping assimilation windows in a way that a new analysis is still produced at the required times (Fisher et al. 2011). One could imagine an assimilation system with a 24-hour assimilation window shifting every 6 hours by 6 hours. Such system will produce the required initial conditions to produce forecasts every 6 hours while allowing enough time to synchronise the two components. Such assimilation systems with overlapping analysis have been fraught upon by the community because they violate some assumptions by assimilating observations twice and hence correlating forecast and observational errors. However, several mathematically rigorous ideas have emerged recently that pave path for assimilation with overlapping windows. Data could be thinned in a way that overlapping windows do not use the same data twice (Bonavita et al. 2017), the first guess trajectory and the linearisation trajectory could be decoupled (Elias Holm, personal communication, 2017), or the observation errors could be inflated when the same data is assimilated multiple

times (Bocquet & Sakov 2014).

A third strategy to mitigate the effects of short operational window, is to use different assimilation windows in the ocean and atmosphere. For example, the U.S. Naval Research Laboratory is currently developing a system where atmospheric 4D-Var system operates on a traditional 6-hour assimilation cycle but the ocean 3D-Var system operates on a longer 24-hour assimilation cycle. That is 4 atmospheric DA cycles are performed for one ocean cycle, with all short forecasts performed using a fully coupled model. In such a system, it is reasonable to expect that the atmospheric state can be brought into balance with the ocean update within the 24-hour ocean assimilation cycle.

As mentioned earlier, the outer loop coupling provides only a first approximation to the fully coupled 4DVAR. Ideally, the outer loop coupling will be augmented by both the coupled adjoint and the specification of the initial time coupled covariances. Such systems could mitigate for the deficiencies of the ensemble based coupled covariance specification and for the cases where the coupled model will not be able to synchronise the unbalanced increments. For example, when ice velocities are assimilated (Barth et al. 2015, Meier et al. 2000), the analysed ice velocities do not persist unless the atmospheric analysis has a corresponding change in surface winds. Another example is a slow coupling of the SST variation and convection over the warm pool area (Fujii et al. 2009, Saha et al. 2010). To properly resolve such slow coupling, the CERA outerloop will need to be increased far beyond the current 24 hours which is impractical with a high-resolution atmospheric model. Instead it will be possible to specify coupled, ensemble-informed covariances at the initial time. From an ensemble perspective, more work is needed to develop effective cross-fluid localisation strategies that can combine the adaptive localisation of (Ménétrier et al. 2015*a*,*b*) that is noisy for small ensemble sizes and a more predictable but suboptimal performance of a static coupled localisation, such as the one used in Frolov et al. (2016).

As forecasts progress towards coupled modelling in the different operational NWP centres, interactions between the different Earth system components need to be fully taken into account, not only during the forecast but also for the definition of the initial conditions of the forecasts. The outer loop coupling strategy developed initially at ECMWF for the atmosphere and the ocean has been extended to the assimilation of sea-ice concentration. Work is ongoing to investigate how the coupling with the land and the wave assimilation systems could be improved in the future.

9 Acknowledgements

The work described in this article was supported by the Office of Naval Research award number N0001412WX20323 and the ERA-CLIM2 project, funded by the European Unions Seventh Framework Programme under grant agreement no. 607029.

References

Bannister, R. (2008), 'A review of forecast error covariance statistics in atmospheric variational data assimilation. I: Characteristics and measurements of forecast error covariances', *Quarterly Journal of the Royal Meteorological Society* **134**(637), 1951–1970.

Barth, A., Canter, M., Schaeybroeck, B. V., Vannitsem, S., Massonnet, F., Zunz, V., Mathiot, P., Alvera-Azcrate, A. & Beckers, J.-M. (2015), 'Assimilation of sea surface temperature, sea ice concentration and sea ice drift in a model of the southern ocean', *Ocean Modelling* **93**, 22 – 39.

Bocquet, M. & Sakov, P. (2014), 'An iterative ensemble kalman smoother', *Quarterly Journal of the Royal Meteorological Society* **140**(682), 1521–1535.

Bonavita, M., Holm, E., Isaksen, L. & Fisher, M. (2016), 'The evolution of the ECMWF hybrid data assimilation system', *Quarterly Journal of the Royal Meteorological Society* **142**(694), 287–303.

Bonavita, M., Tremolet, Y., Holm, E., Lang, S., Chrust, M., Janiskova, M., Lopez, P., Laloyaux, P., De Rosnay, P., Fisher, M., Hamrud, M. & English, S. (2017), A strategy for data assimilation, Technical Memorenda 800, ECMWF.

Buizza, R., Milleer, M. & Palmer, T. (1999), 'Stochastic representation of model uncertainties in the ecmwf ensemble prediction system', *Quarterly Journal of the Royal Meteorological Society* **125**(560), 2887–2908.

Courtier, P., Thépaut, J.-N. & Hollingsworth, A. (1994), 'A strategy for operational implementation of 4D-Var, using an incremental approach', *Quarterly Journal of the Royal Meteorological Society* **120**, 1367–1388.

Daley, R. (1991), *Atmospheric data analysis*, Cambridge University Press.

Dee, D. P. (2005), 'Bias and data assimilation', *Quarterly Journal of the Royal Meteorological Society* **131**(613), 3323–3343.

ECMWF (2016), IFS cycle 41R2, Official documentation, ECMWF.
URL: *https://www.ecmwf.int/en/forecasts/documentation-and-support/changes-ecmwf-model/cy41r2-summary-changes*

Feng, X., Haines, K. & de Boissson, E. (2018), 'Coupling of surface air and sea surface temperatures in the cera-20c reanalysis', *Quarterly Journal of the Royal Meteorological Society* pp. n/a–n/a.
URL: *http://dx.doi.org/10.1002/qj.3194*

Fisher, M., Tremolet, Y., Auvinen, H., Tan, D. & Poli, P. (2011), Weak-constraint and long window 4DVAR, Technical Memorenda 655, ECMWF.

Frolov, S., Bishop, C. H., Holt, T., Cummings, J. & Kuhl, D. (2016), 'Facilitating strongly coupled oceanatmosphere data assimilation with an interface solver', *Monthly Weather Review* **144**(1), 3–20.

Fujii, Y., Nakaegawa, T., Matsumoto, S., Yasuda, T., Yamanaka, G. & Kamachi, M. (2009), 'Coupled climate simulation by constraining ocean fields in a coupled model with ocean data', *Journal of Climate* **22**, 5541–5557.

Gratton, S., Laloyaux, P., Sartenaer, A. & Tshimanga, J. (2011), 'A reduced and limited-memory preconditioned approach for the 4d-var data-assimilation problem', *Quarterly Journal of the Royal Meteorological Society* **137**(655), 452–466.

Kalnay, E. (2003), *Atmospheric modelling, data assimilation and predictability*, Cambridge University Press.

Laloyaux, P., Balmaseda, M., Dee, D., Mogensen, K. & Janssen, P. (2016), 'A coupled data assimilation system for climate reanalysis', *Quarterly Journal of the Royal Meteorological Society* **142**(694), 65–78.

Laloyaux, P., de Boisseson, E., Balmaseda, M., Bidlot, J., Broennimann, S., Buizza, R., Dalhgren, P., Dee, D., Haimberger, L., Hersbach, H., Kosaka, Y., Martin, M., Poli, P., Rayner, N., Rustemeier, E. & Schepers, D. (2018), 'CERA-20C: A coupled reanalysis of the twentieth century', *Journal of Advances in Modeling Earth Systems* .

Laloyaux, P., Thépaut, J.-N. & Dee, D. (2016), 'Impact of scatterometer surface wind data in the ECMWF coupled assimilation system', *Monthly Weather Review* **144**(3), 1203–1217.

Lea, D. J., Mirouze, I., Martin, M. J., King, R. R., Hines, A., Walters, D. & Thurlow, M. (2015), 'Assessing a new coupled data assimilation system based on the met office coupled atmospherelandoceansea ice model', *Monthly Weather Review* **143**(11), 4678–4694.

Lorenc, A. C., Bowler, N. E., Clayton, A. M., Pring, S. R. & Fairbairn, D. (2015), 'Comparison of hybrid-4DEnVar and hybrid-4DVar data assimilation methods for global nwp', *Monthly Weather Review* **143**(1), 212–229.

Madec, G. (2008), NEMO ocean engine, Note du Pole de modélisation 27 ISSN No 1288-1619, Institut Pierre-Simon Laplace.

Meier, W. N., Maslanik, J. A. & Fowler, C. W. (2000), 'Error analysis and assimilation of remotely sensed ice motion within an arctic sea ice model', *Journal of Geophysical Research: Oceans* **105**(C2), 3339–3356.

Ménétrier, B. & Auligne, T. (2015), 'Optimized localization and hybridization to filter ensemble-based covariances', *Monthly Weather Review* **143**(10), 3931–3947.

Ménétrier, B., Montmerle, T., Michel, Y. & Berre, L. (2015*a*), 'Linear filtering of sample covariances for ensemble-based data assimilation. part i: Optimality criteria and application to variance filtering and covariance localization', *Monthly Weather Review* **143**(5), 1622–1643.

Ménétrier, B., Montmerle, T., Michel, Y. & Berre, L. (2015*b*), 'Linear filtering of sample covariances for ensemble-based data assimilation. part ii: Application to a convective-scale nwp model', *Monthly Weather Review* **143**(5), 1644–1664.

Mogensen, K., Balmaseda, M. & Weaver, A. (2012), The NEMOVAR ocean data assimilation system as implemented in the ECMWF ocean analysis for System 4, Technical Memorenda 668, ECMWF.

Mogensen, K., Keeley, S. & Towers, P. (2012), Coupling of the NEMO and IFS models in a single executable, Technical Memorenda 673, ECMWF.

Mulholland, D. P., Laloyaux, P., Haines, K. & Balmaseda, M. A. (2015), 'Origin and impact of initialization shocks in coupled atmosphereocean forecasts', *Monthly Weather Review* **143**(11), 4631–4644.

Nozomi, S., Toshiyuki, A., Shuhei, M., Takashi, M., Takahiro, T., Toru, M., Hiromichi, I. & Yoichi, I. (2008), 'Development of a fourdimensional variational coupled data assimilation system for enhanced analysis and prediction of seasonal to interannual climate variations', *Journal of Geophysical Research: Oceans* **113**(C10), 1–21.

Palmer, T., R. Buizza, F. D.-R., T. Jung, M. L., G. Shutts, M. S. & Weisheimer, A. (2009), Stochastic parametrization and model uncertainty, Technical Memorenda 598, ECMWF.

Rabier, F. (2005), 'Overview of global data assimilation developments in numerical weather-prediction centres', *Quarterly Journal of the Royal Meteorological Society* **131**(613), 3215–3233.

Saha, S., Moorthi, S., Pan, H., Wu, X., Wang, J., Nadiga, S., Tripp, P., Kistler, R., Woollen, J., Behringer, D., Liu, H., Stokes, D., Grumbine, R., Gayno, G., Wang, J., Hou, Y., Chuang, H., Juang, H., Sela, J., Iredell, M., Treadon, R., Kleist, D., Delst, P. V., Keyser, D., Derber, J., Ek, M., Meng, J., Wei, H., Yang, R., Lord, S., Dool, H. V. D., Kumar, A., Wang, W., Long, C., Chelliah, M., Xue, Y., Huang, B., Schemm, J., Ebisuzaki, W., Lin, R., Xie, P., Chen, M., Zhou, S., Higgins, W., Zou, C., Liu, Q., Chen, Y., Han, Y., Cucurull, L., Reynolds, R., Rutledge, G. & Goldberg, M. (2010), 'The NCEP climate forecast system reanalysis', *Bulletin of the American Meteorological Society* **91**, 1015–1057.

Schepers, D. (2018), CERA-SAT, ECMWF Newsletter 155, ECMWF.

Sluka, T. C., Penny, S. G., Kalnay, E. & Miyoshi, T. (2016), 'Assimilating atmospheric observations into the ocean using strongly coupled ensemble data assimilation', *Geophysical Research Letters* **43**(2), 752–759.

Smith, P. J., Fowler, A. M. & Lawless, A. S. (2015), 'Exploring strategies for coupled 4D-Var data assimilation using an idealised atmosphereocean model', *Tellus A: Dynamic Meteorology and Oceanography* **67**(1), 27025.

Smith, P. J., Lawless, A. S. & Nichols, N. K. (2017), 'Estimating forecast error covariances for strongly coupled atmosphereocean 4d-var data assimilation', *Monthly Weather Review* **145**(10), 4011–4035.

Smith, P. J., Lawless, A. S. & Nichols, N. K. (2018), 'Treating sample covariances for use in strongly coupled atmosphere-ocean data assimilation', *Geophysical Research Letters* **45**(1), 445–454. 2017GL075534.

Titchner, H. A. & Rayner, N. A. (2014), 'The met office hadley centre sea ice and sea surface temperature data set, version 2: 1. sea ice concentrations', *Journal of Geophysical Research: Atmospheres* **119**(6), 2864–2889. 2013JD020316.

Tremolet, Y. (2005), Incremental 4D-Var convergence study, Technical Memorenda 469, ECMWF.

Tshimanga, J., Gratton, S., Weaver, A. T. & Sartenaer, A. (2008), 'Limited-memory preconditioners, with application to incremental four-dimensional variational data assimilation', *Quarterly Journal of the Royal Meteorological Society* **134**(632), 751–769.

Veerse, F. & Thépaut, J.-N. (1998), 'Multiple-truncation incremental approach for four-dimensional variational data assimilation', *Quarterly Journal of the Royal Meteorological Society* **124**, 1889–1908.

Wada, A. & Kunii, M. (2017), 'The role of oceanatmosphere interaction in typhoon sinlaku (2008) using a regional coupled data assimilation system', *Journal of Geophysical Research: Oceans* **122**(5), 3675–3695.