ABSTRACT

Title of dissertation: STRONGLY COUPLED OCEAN-ATMOSPHERE DATA ASSIMILATION WITH THE LOCAL ENSEMBLE TRANSFORM KALMAN FILTER

Travis C. Sluka, Doctor of Philosophy, 2018

Dissertation directed by: Professor Eugenia Kalnay
Dept. of Atmospheric and Oceanic Science

Current state-of-the-art coupled data assimilation systems handle the ocean and atmosphere separately when generating an analysis, even though ocean atmosphere models are subsequently run as a coupled system for forecasting. Previous research using simple 1-dimensional coupled models has shown that strongly coupled data assimilation (SCDA), whereby a coupled system is treated as a single entity when creating the analysis, reduces errors for both domains when using an ensemble Kalman filter. A prototype method for SCDA is developed with the local ensemble transform Kalman filter (LETKF). This system is able to use the cross-domain background error covariance from the coupled model ensemble to enable assimilation of atmospheric observations directly into the ocean. This system is tested first with the intermediate complexity SPEEDYNEMO model in an observing system simulation experiment (OSSE), and then with real observations and an operational coupled model, the Climate Forecasting System v2 (CFSv2). Finally, the development of a major upgrade to ocean data assimilation used at NCEP (the Hybrid-GODAS) is

presented, and shown how this new system could help present a path forward to operational strongly coupled DA.

STRONGLY COUPLED OCEAN-ATMOSPHERE DATA ASSIMILATION WITH THE LOCAL ENSEMBLE TRANSFORM KALMAN FILTER

by

Travis C. Sluka

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2018

Advisory Committee:
Prof. Eugenia Kalnay, Chair/Advisor
Prof. James Carton
Prof. Brian Hunt
Prof. Daryl Kleist
Prof. Stephen Penny

Dedication

My wife: *"You should plant pansies outside."*

Me: *"We should be getting a hard freeze soon, it's pointless!"*

Me two months later in late December: *"It's been 60 degrees outside for weeks, I should have planted pansies a long time ago!"*

...may my research someday help improve short-term climate forecasts, so I'll know when to plant the stupid pansies.

Acknowledgments

Fred Kucharski of the Adbus Salam International Centre for Theoretical Physics (ICTP) in Trieste, Italy, very kindly offered the source code for his SPEEDYNEMO model, as well as debugging support for the model. Computational resources were provided by the University of Maryland through the Deepthought 2 (DT2) supercomputer. Assistance in generating plots was provided by Alfredo Ruiz-Barradas for the climatological analysis of the SPEEDYNEMO nature runs, and Yan Xue for the comparison plots of GODAS/Hybrid-GODAS/EN4/CMC. Hyun-Chul Lee and Rahul Mahajan of NOAA/EMC, Yan Xue of NOAA/CPC, and Steve Penny and Jim Carton of UMD, and others, have all helped with the development of the Hybrid-GODAS system. Numerous other faculty and students of the UMD Weather and Chaos group were instrumental in providing feedback and guidance throughout the course of this research.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

AOGCM	atmosphere-ocean global circulation model
CFSv2	Climate Forecasting System version 2
CPC	Climate Prediction Center
DA	data assimilation
EMC	Environmental Modeling Center
EMIC	Earth system model of intermediate complexity
EnKF	ensemble Kalman filter
GCM	global circulation model
GFDL	Geophysical Fluid Dynamics Laboratory
GODAS	global ocean data assimilation system
JCSDA	Joint Center for Satellite Data Assimilation
JEDI	Joint Effort for Data Assimilation Integration
LACC	lagged-average coupled covariance
LETKF	local ensemble transform Kalman filter
MJO	MaddenJulian oscillation
MOM	modular ocean model
NCEP	National Centers for Environmental Prediction
NWP	numerical weather prediction
OISST	optimum interpolation sea surface temperature
OSSE	observing system simulation experiment
RMSD	root mean square deviation
RMSE	root mean square error
RTPP	relaxation to prior perturbations
RTPS	relaxation to prior spread
SCDA	strongly coupled data assimilation
S2S	sub-seasonal to seasonal
SSS	sea surface salinity
SST	sea surface temperature
UMD-LETKF	universal multi-domain LETKF
WCDA	weakly coupled data assimilation

Chapter 1: Introduction

1.1 Data assimilation in coupled ocean atmosphere models

Coupled ocean-atmosphere models have become increasingly important to improving operational numerical weather prediction for a wide range of phenomena. Such phenomena on seasonal climate timescales include the forecasting of the El Niño Southern Oscillation (ENSO) and the Madden Julian Oscillation (Madden and Julian, 1972; Zebiak, 1989). It has been shown that improvements on shorter time-scales can be made with coupled ocean-atmosphere models as well for tropical cyclones (Fu et al., 2007) and annual monsoons (Klingaman et al., 2008). The importance of such systems can be noted by the choice of the European Centre for Medium-Range Weather Forecasts (ECMWF) to transition to fully coupled models for all of their operational numerical weather prediction.

Climate models typically consist of separate domains, modeled independently, which are then coupled together when running the forecast. For example, an earth system model might be composed of individual ocean, atmosphere, land, ice, and wave, models. This work will focus primarily on data assimilation in a coupled ocean-atmosphere model context, though the concepts can be applied to all components of a coupled Earth system model that would normally be considered as separate

domains.

In order to initialize these coupled models, an accurate estimate of the initial state of the atmosphere and ocean is required. Over the past several decades numerical weather prediction has advanced in part due to progress made in producing these initial conditions more accurately. The process of generating this objective analysis is known as data assimilation (DA), and can be described most succinctly as combining current observations with past information (from a previous model forecast) to produce an analysis, the best estimate of the system's current state. DA methods have evolved from simple nudging, optimal interpolation (OI), and three dimensional variational (3DVAR) methods, to the more advanced ensemble Kalman filter (EnKF) and four-dimensional variational method (4DVAR) (Kalnay, 2003). Today, further advances are made in the area of data assimilation by combining the most advanced of these existing methods into hybrid EnKF/Var systems (Bannister, 2017).

The data assimilation methods used for coupled models can be broadly divided into three categories summarized by fig 1.1:

1. **uncoupled DA** - A separate ocean background forecast and atmosphere background forecast are used with separate data assimilation.

2. **weakly coupled DA (WCDA)** - A single coupled forecast is used for the background, but DA is still performed separately on the ocean and atmosphere.

3. **strongly coupled DA (SCDA)** - A single coupled forecast is used for the background, and a single DA system is used for generating the analysis.

It is important to remember that here the word "coupled" when describing data assimilation refers to how the domains interact with each other for the generation of the analysis, and not with subsequent longer forecasts using those initial conditions. Uncoupled DA (fig 1.1a) uses the background forecast from separate ocean and atmosphere model runs, and also performs data assimilation on each domain completely independently (e.g. Saha et al. (2006); Maclachlan et al. (2015)). As an example, the original Climate Forecasting System (CFS) used at the National Centers for Environmental Prediction (NCEP) starting in 2004 (Saha et al., 2006), was a coupled ocean-atmosphere-land model that was used for climate forecasts. However, the initial conditions for the background were obtained from data assimilation cycles that used different stand-alone models. The atmospheric initial conditions were obtained from the atmosphere-only NCEP Reanalysis-2 (R2) (Kanamitsu et al., 2002), and the ocean initial conditions were obtained from the ocean-only Global Ocean Data Assimilation System (GODAS) (Behringer and Xue, 2004). R2 is prescribed sea surface temperatures from an independent SST product, and the GODAS ocean is given atmospheric fluxes from the independent R2 run. So, although CFS was a coupled model, it utilized uncoupled data assimilation to initialize the model.

With weakly coupled DA (fig 1.1b), the same coupled model that is used for subsequent forecasts is also used to generate the background for the data assimilation cycle. The actual data assimilation is, however, still then done separately. Information can be transferred between the two domains only through the integration of the background forecast. Therefore assimilation of observations in one domain

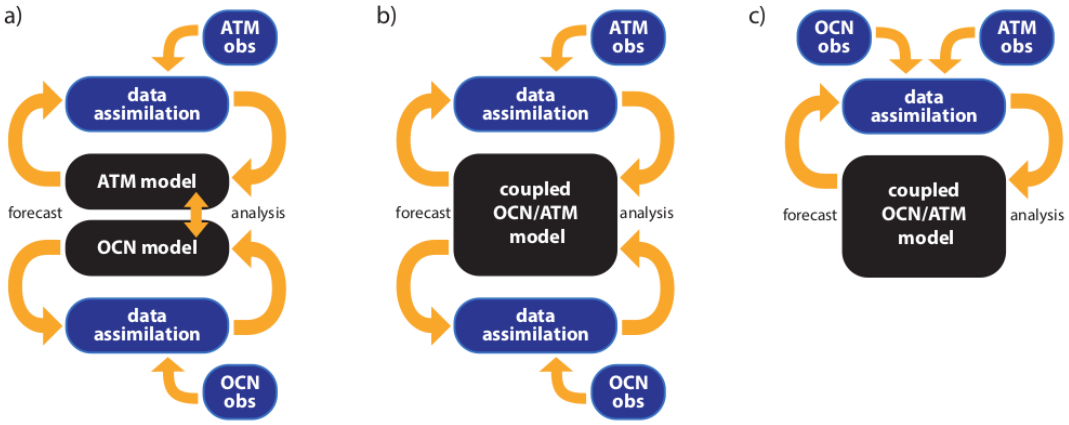

Figure 1.1: Schematic of ocean-atmosphere data assimilation coupling types: uncoupled (a), weakly coupled (b), and strongly coupled (c).

take time to become beneficial, if ever, to the other domain. WCDA has been performed successfully in at various centers including the UK Met Office (Lea et al., 2015), GFDL (Zhang et al., 2007), and NCEP (Saha et al., 2014), and has been shown to possibly produce better forecasts due to better initial conditions from the weakly coupled DA. The CFSv2, for example, uses the Climate Data Assimilation System (CDAS), in which the same coupled model is used for the forecasts and for the generation of the background for the data assimilation.

When initializing any model there is a potential for "initialization shocks" whereby some type of imbalance, artificially generated as an artifact from the data assimilation cycle, is present and can degrade the initial forecast performance. This is true of any model, but coupled models present additional opportunities for generating initialization shocks. As described by Mulholland et al. (2015), these shocks can occur when:

1. using an uncoupled model to produce the initial conditions for the coupled

forecast model

2. using a different forecast model than the model used for the data assimilation's background

3. using the same model, but with different or removed bias correction schemes

All of the above instances produce initial conditions that may be initially incompatible with the coupled model and so require a spin-up time while the flux balances adjust. Using the same coupled ocean atmosphere model for the forecasts as is used for the data assimilation should therefore reduce the initialization shocks that would otherwise be produced. Reducing these initialization shocks is thought to be important for improving seasonal forecasts (Balmaseda and Anderson, 2009) and the Madden-Julian oscillation (Marshall et al., 2011).

There are several reasons why uncoupled DA had been, and still is in some situations, preferred over WCDA. Drifts in coupled ocean-atmosphere models are a common problem, and unless they are properly accounted for these drifts can significantly impact the background used in the data assimilation. For example, it is a currently reoccurring problem for NCEP's weakly coupled CDAS used in the CFSv2 that the ocean in the tropical Atlantic develops a large cold bias and drifts away from observations. However, the standalone ocean-only GODAS, for which the data assimilation system is nearly identical, does not drift.

Strongly coupled DA (SCDA) (fig 1.1c) uses a coupled forecast for the background, similar to WCDA, but then also performs the data assimilation as a single system. In this way, observations in one domain are able to immediately impact

other domains during the analysis step. SCDA is able to transmit corrections to the state of one domain both through the integration of the model background forecast, and also by utilizing the cross-domain background-error covariance during the analysis step. Such coupling allows observations in one domain to instantaneously impact the state variable in the other domain. Strongly coupled DA should be able to extract more information from the same observations, given that there are meaningful ocean-atmosphere correlations to be used by the data assimilation system, and should retain a better balance between the two domains.

It is being increasingly realized by various centers in the United States (National Academies of Sciences Engineering and Medicine, 2016) and Europe (ECMWF, 2016) that moving from weakly to strongly coupled DA has the potential to be beneficial for sub-seasonal to seasonal (S2S) prediction and may be crucial for improving long range forecasts. There are currently no operational systems utilizing SCDA for Earth system models, and until now, the majority of research into SCDA has been performed using simplistic models. Efforts are now being made to accelerate the development of SCDA research (Lawless, 2012; Penny and Hamill, 2017).

There are several reasons why ocean-atmosphere models have used WCDA, initializing the two domains independently. There is often a much higher observational coverage in the atmosphere compared to the ocean, and additionally ocean observations had traditionally been slower to be processed, lagging behind the sub-daily times common for synoptic atmospheric observations. Advancements in data assimilation have usually been focused on improving the atmosphere first, and so ocean DA has generally lagged behind. For example both NCEP (Kleist and Ide, 2015)

and ECMWF (Bonavita et al., 2012) have operational hybrid ensemble/variational DA systems for the atmosphere, though are still using traditional 3-D variational methods for the ocean (Saha et al., 2014; Mogensen et al., 2012). There have been no significant updates to the global ocean DA system used by NCEP since about 2003 (Behringer and Xue, 2004) relying on the same 3d variational system. Different temporal and spatial scales and different grid types by each model adds complications as well.

Much work goes into specifying the background error covariance for either an atmospheric or oceanic variational data assimilation system, and needing to specify cross domain covariance would greatly add to the difficulty. The ocean-atmosphere boundary layer between two domains, is often insufficiently modeled. The first level of NCEP's current global ocean system is 10 meters thick, far too coarse to properly resolve a diurnal cycle in the ocean. It is recognized that the vertical resolution of the models at the interface, and parameterizations used for representing the boundary layer need to be improved (Grissom et al., 2017; Cravatte et al., 2015).

1.2 Strongly coupled DA and recent research

The possible reduction of initialization shocks is one motivation for moving from uncoupled data assimilation to weakly coupled DA (Mulholland et al., 2015). Even when the same model is used for the data assimilation background and the subsequent forecast, initialization shocks can still occur due to the way in which the data assimilation is performed. For example, ensemble data assimilation often

specifies a localization radius, eliminating the impact of an observation from model grid points that are too far away to account for spurious correlations caused by insufficient ensemble size (Greybush et al., 2011). If the localization radius is too small, such as less than the Rossby radius of deformation, the analysis is left in an imbalanced state, and gravity waves are generated at the first time-steps of the forecast, and therefore represent an initialization shock. This has been observed within a single model domain, but it is also likely that weakly coupled data assimilation (which is in effect using domain localization) is generating cross domain shocks, as the surface states and fluxes might no longer be in a physically realistic state.

No operational strongly coupled ocean-atmosphere data assimilation system currently exists but the closest to SCDA is the coupled ECMWF reanalysis (CERA) (Laloyaux et al., 2015). With this system, the atmosphere data assimilation is performed with a 4DVAR, and the ocean with a separate 3DVAR. The analysis step is therefore inherently weakly coupled DA, and observations do not directly impact across the domain. However, the 4DVar uses an outer loop, whereby a minimization is performed with the tangent linear model, and then the full non-linear coupled model is run again before repeating the inner loop minimization. By using the full coupled model in the outer loop, the observations are able to indirectly impact the opposite domains as can be seen in single observation experiments (fig 1.2). This method lies somewhere between SCDA and WCDA (Quasi-SCDA), since the variational minimization does not work across the domain, but the analysis is impacted due to the model integration of the outer loop.

Most research with truly strongly coupled data assimilation has been carried

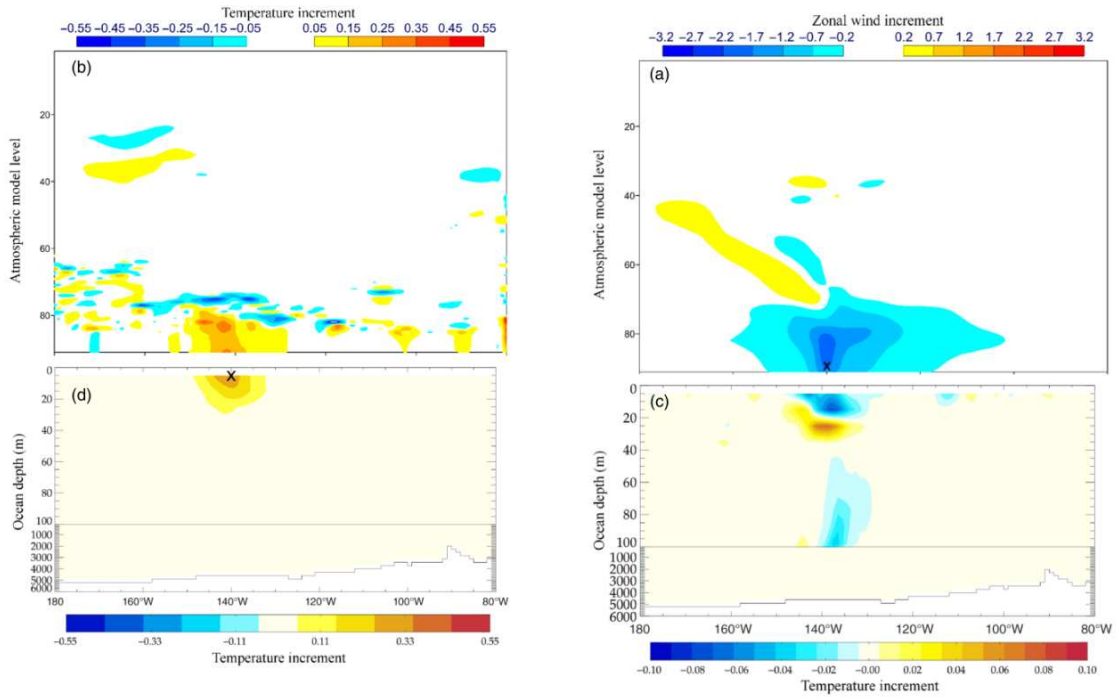

Figure 1.2: Single observation tests with the ECMWF CERA climate reanalysis showing the cross domain impacts due to it's 'quasi-strongly coupled' implementation. A single ocean temperature observation at 5m depth (left). A single atmospheric zonal wind observation at the lowest atmospheric level (right). (Laloyaux et al., 2015)

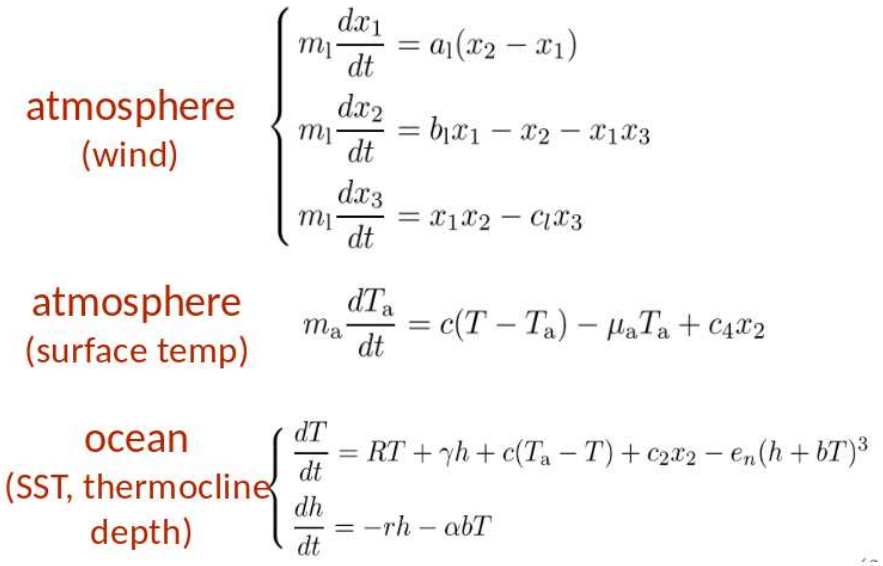

Figure 1.3: A simplified coupled ocean-atmosphere single column model from Liu et al. (2013).

out using simple one-dimensional toy models, such as coupled Lorenz models (fig 1.3). Most have used an EnKF (Lu et al., 2015a; Liu et al., 2013; Han et al., 2013; Luo and Hoteit, 2014; Tardif et al., 2014) although some work has been done with 4DVAR data assimilation as well (Smith et al., 2015). Liu et al. (2013) found that strongly coupled DA provides substantial improvements, with the greatest impacts seen by assimilating atmospheric observations into the oceans in the extra-tropics. Tardif et al. (2014) found that the Atlantic meridional overturning circulation (AMOC) in the ocean can is not be correctly initialized in their low dimensional model even when only time averaged atmospheric observations are assimilated.

One study to date with a more realistic coupled ocean atmosphere model has explored SCDA (Lu et al., 2015b). Lu et al. (2015b) uses a low resolution ocean

atmosphere model to assimilate averaged atmospheric observations into the ocean. They found that weakly coupled DA is detrimental with regard to maintaining a proper correlation at the ocean atmosphere interface fig 1.4. Also, correlations between the two domains are stronger when the atmospheric observations are averaged over at least a week. A method called the lagged average coupled covariance (LACC) is developed. Using this method, they find that they, similar to the simplified 1-dimensional models, are able to improve the extra-tropics by assimilating atmospheric observations into them (fig 1.5). The LACC experiments, though, were driven with monthly averaged SSTs, at shorter timescales the signals would have been damped.

For strongly coupled DA systems to be practical for operational numerical weather prediction (NWP), the ocean and atmosphere DA cycles should use similar observational windows. Ensemble Kalman filters (EnKFs), in contrast to 4D-Variational methods, perform best with short assimilation windows (Kalnay et al., 2007). Thus, the EnKF allows for both systems to perform assimilation at the shorter window length of the atmosphere (Singleton, 2011). With a variational system, the cross domain covariance needs to be explicitly defined in some manner. There is a significant amount of work, even for a single domain in uncoupled DA, that goes into generating these background error covariances. It would be expected that things would be further complicated by having to generate the covariances for a coupled system, though some groups are pursuing this route (Frolov et al., 2016; Smith et al., 2018). A benefit of the EnKF is that these background-error covariances are automatically generated by the ensemble, assuming that such correlations

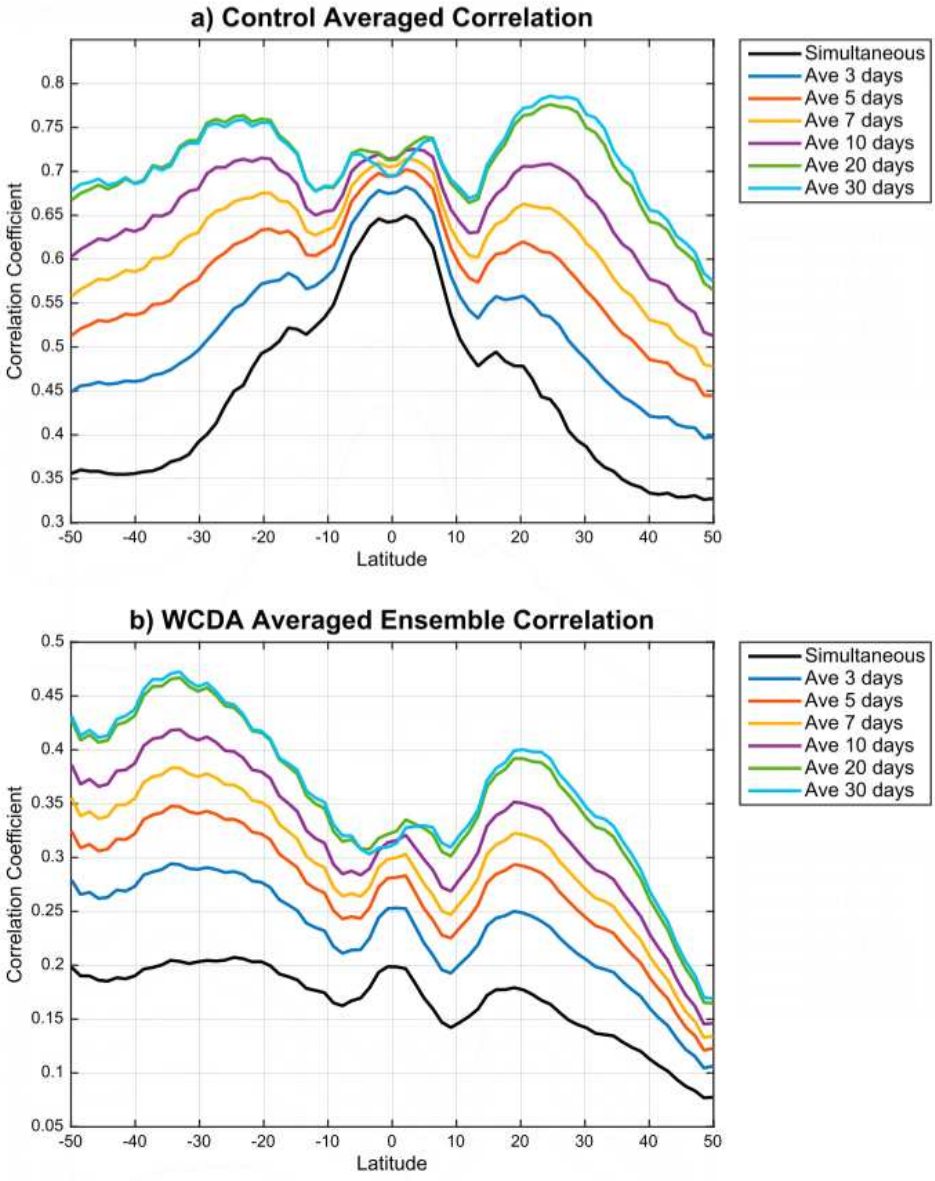

Figure 1.4: Impact of coupled DA on cross domain correlations, from Lu et al. (2015b).

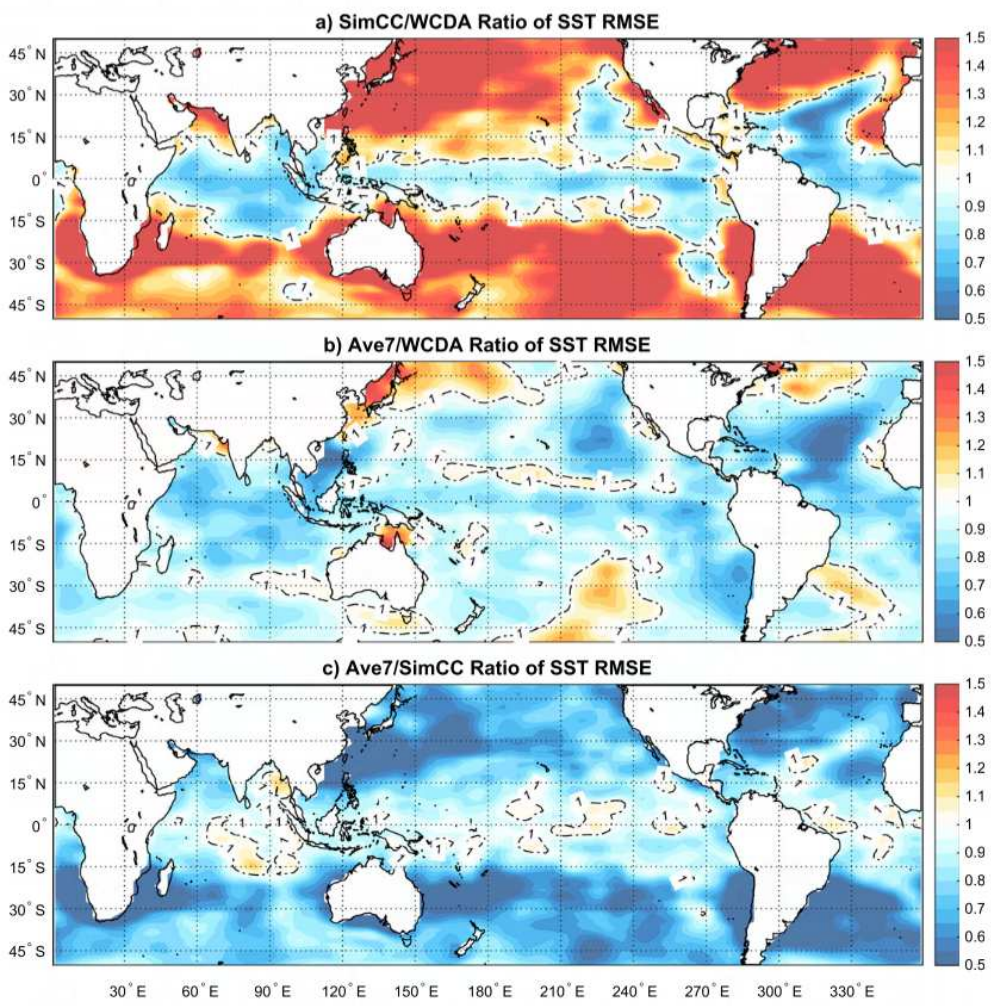

Figure 1.5: Impact of strongly coupled DA on intermediate complexity model from Lu et al. (2015b). Simple simultaneous SCDA worsens the analysis (top) but assimilating averaged atmospheric observations improves the analysis compared with WCDA (middle) and simultaneous SCDA (bottom).

are indeed present in the model. Special care does need to be taken though to prune spurious correlations that appear due to rank deficiency of the ensemble (Yoshida and Kalnay, 2018). Given the EnKFs preference for short assimilation windows, allowing us to perform all DA at the atmospheric timescale (6 hrs to 1 day), as well as not needing to manually specify cross-domain background-error covariances, the EnKF was chosen here to pursue SCDA experiments with a realistic model.

Given that multiple studies with low dimensional ocean atmosphere models (Liu et al., 2013; Tardif et al., 2014) demonstrate that SCDA has a more beneficial impact when assimilating the atmospheric observations into the ocean, the initial experiments for this dissertation will focus on the same method of assimilating atmospheric observations into the ocean with the LETKF.

1.3 Local ensemble transform Kalman filter

A brief description of how the LETKF operates, and how its formulation can benefit from strongly coupled DA follows. The LETKF (Hunt et al., 2007) is a type of ensemble Kalman filter (Evensen, 1994), using an ensemble of forecasts $\left\{\mathbf{x}^{b(i)} : i = 1, 2, ..., k\right\}$ to determine the statistics of the background error covariance. This information is combined with new observations $\mathbf{y}^o$, to generate an analysis mean, $\bar{\mathbf{x}}^a$, and a set of new ensemble members, $\mathbf{x}^{a(i)}$. First, the model state is mapped to observation space by applying a non-linear observation operator H to each background ensemble member $\mathbf{y}^{b(i)} = H\mathbf{x}^{b(i)}$. If the observed and modeled variables are the same, H is simply an interpolation of the model state to the

observation locations. The weights $\bar{\mathbf{w}}^a$ are calculated to find the analysis mean $\bar{\mathbf{x}}^a$

$$\tilde{\mathbf{P}}^a = \left[(k-1)\,\mathbf{I} + \left(\mathbf{Y}^b\right)^T \mathbf{R}^{-1}\mathbf{Y}^b\right]^{-1} \tag{1.1}$$

$$\bar{\mathbf{w}}^a = \tilde{\mathbf{P}}^a \left(\mathbf{Y}^b\right)^T \mathbf{R}^{-1} \left(\mathbf{y}^o - \bar{\mathbf{y}}^b\right) \tag{1.2}$$

$$\bar{\mathbf{x}}^a = \bar{\mathbf{x}}^b + \mathbf{X}^b\bar{\mathbf{w}}^a \tag{1.3}$$

where $\bar{\mathbf{x}}^b$ and $\bar{\mathbf{y}}^b$ are the ensemble mean of the background in model space and observation space, respectively, $\mathbf{X}^b$ and $\mathbf{Y}^b$ are the matrices whose columns represent the ensemble perturbations from those means, and $\mathbf{R}$ is the observation error covariance matrix. Last, the set of weights $\mathbf{W}^a$ are calculated to find the perturbations in the model space for the analysis ensemble by

$$\mathbf{W}^a = \left[(k-1)\,\tilde{\mathbf{P}}^a\right]^{1/2} \tag{1.4}$$

$$\mathbf{X}^a = \mathbf{X}^b\mathbf{W}^a \tag{1.5}$$

In practice, the LETKF is able to calculate the above equations in parallel for each grid point j using the subset of observations, $\mathbf{y}_j^o$, within its localization radius. This makes the LETKF computationally efficient and highly scalable. For weakly coupled DA, $\mathbf{y}_j^o$ contains only observations from the same domain as the grid point

being considered, whereas $\mathbf{y}_j^o$ can contain both atmospheric ($\mathbf{y}_{atm}^o$) and ocean ($\mathbf{y}_{ocn}^o$) observations in strongly coupled DA.

The LETKF benefits from strongly coupled DA in two key ways. First, the calculation of $\bar{\mathbf{x}}^a$ by equations (1.1) − (1.3) uses the cross domain error covariance to allow observations in one domain to directly inform the analysis mean calculated at grid points in the other domain. This can be especially beneficial to the ocean, where observations are often sparse compared with the observation densities of the atmosphere.

Second, the creation of the analysis ensemble by equations (1.4) and (1.5) maintains balance between the two domains within each ensemble member. Neighboring grid points use overlapping sets of observations, and since $\mathbf{y}_j^o$ will be nearly identical for adjacent grid points, $\mathbf{W}_j^a$ will be similar as well (Yang et al., 2009). Similar weights for neighboring grid points, both vertically and horizontally, ensures the ensemble perturbations are kept "matched together" at the domain interface. Weakly coupled DA is not able to retain this ocean-atmosphere surface balance within the ensemble members.

1.4 Outline of this research

The goal of this research is produce a prototype strongly coupled ocean-atmosphere data assimilation system that would be suitable for operational numerical weather prediction. Previous studies have not been conducted with operational quality models, and those that have used realistic models are oriented towards differ-

ent timescales, such as Lu et al. (2015b) focusing on monthly and weekly timescales. These timescales, and associated methods, are not practical for real-time operations. The methods developed here will focus on daily and sub-daily SCDA cycles.

The research will be conducted in three steps. The first step toward SCDA in an operational cycle, described in Chapter 2, will be the construction of a strongly coupled DA system using the local ensemble transform Kalman filter (LETKF). This system will be tested with an observing system simulation experiment (OSSE) and a simplified climate model (SPEEDYNEMO). This model is still able to produce realistic phenomena, though is simple enough to be run extremely fast.

Next, a realistic, operational quality model, the Climate Forecasting System version 2 (CFSv2) will be used in Chapter 3, this time under an OSE experiment using real observations. The model is identical to that which is being used operationally by NCEP, though with a lower horizontal resolution in the atmosphere due to computational constraints. This CFSv2-LETKF strongly coupled system will be tested with real observations. Assimilation of real observations presents a slew of additional difficulties such as observation and model biases, and so these experiments will only be carried out with a limited subset of insitu observations to demonstrate potential improvements and difficulties with a full system.

Finally, the development of a next generation ocean data assimilation system (Hybrid-GODAS) at NCEP will be discussed in Chapter 4. While Hybrid-GODAS does not directly relate to strongly coupled data assimilation at the moment, this system uses the LETKF as its backbone, and may form the foundation of a near-future coupled ocean/ice/wave data assimilation at NCEP, a possible gateway toward a

full earth system model SCDA system.

Chapter 2: Intermediate Model OSSE: SPEEDYNEMO-LETKF

Before beginning strongly coupled data assimilation experiments with a full ocean-atmosphere general circulation model, it is useful to apply the method first to a more simple, yet still realistic, model. There exists a general hierarchy of coupled ocean-atmosphere climate models, ranging from simplified analytical models of ENSO variability (e.g. Battisti and Hirst (1989); Zebiak (1989)) and the MJO (Madden and Julian, 1972) to full scale state-of-the-art models used for operational climate prediction such as the Climate Forecasting System (CFSv2) (Saha et al., 2014), Community Earth System Model (CESM) (Hurrell et al., 2013), or Goddard Earth Observing System Model (Molod et al., 2012), just to name a few. Most studies of strongly coupled data assimilation to date have been performed using one extreme of this hierarchy: one-dimensional simplistic models (e.g. Singleton (2011); Han et al. (2013); Luo and Hoteit (2014)). With the ultimate goal of providing improvements to the operational coupled data assimilation systems used for numerical weather forecasting, it is therefore logical to move toward studying SCDA with intermediate complexity models, also known as Earth system model of intermediate complexity (EMIC), (Kucharski et al., 2013).

EMICs are a big step up in terms of complexity compared with simple one-

dimensional analytical models and share many similarities with ocean-atmosphere GCMs. Where they differ from full scale GCMs is usually with regard to lower resolution, both in the horizontal and the vertical direction, which allows for longer step sizes and therefore lowers the costs of the dynamical core. Also, parameterizations usually have been simplified as much as possible, such as by not including a diurnal cycle, which greatly speeds up the radiative transfer parameterizations. These types of simplifications help reduce the costs of the model physics. Despite the simplifications present with an EMIC, they are still three dimensional models with orography in the general shape of the continents and oceans, and are able to often represent features such as seasonal cycles, Hadley circulations, jet streams, inter-tropical convergence zones, and even ENSO-like phenomena if system is tuned correctly. This balance between model speed and the fidelity required to still exhibit somewhat realistic phenomena is extremely beneficial to ensemble data assimilation. Multiple iterative experiments utilizing many ensemble members are often required to test and tune a system. However, this requirement would likely be overly costly to do first with a full atmosphere-ocean general circulation model (AOGCM).

The SCDA system developed here is designed with an operational data assimilation cycle in mind, namely a 6 hour cycle similar to that used by the National Center for Environmental Prediction's (NCEP) Climate Forecasting System version 2 (CFSv2). Since the SPEEDY-NEMO model is very fast, testing can be done in short time, and the lessons learned in the process and the resulting system developed can then be ported to use an operational-quality model such as the CFSv2.

Two unique experiments are carried out with the SPEEDY-NEMO, each com-

paring the performance of weakly coupled versus strongly coupled DA:

1. atmospheric observations only

2. ocean observations only

Experiments with both ocean and atmosphere observations being simultaneously assimilated into both domains were attempted as well, but were not successful, as will be described in the concluding section.

2.1 Method

The following experiments of SCDA with an intermediate model are performed as a series of observing system simulation experiments (OSSEs). In an OSSE, the model is first integrated for a long period of time and the output is saved as what is called a "nature run". This nature run is the truth from which synthetic observations are generated and to which the subsequent data assimilation experiment analyses are compared against to evaluate performance. By utilizing an OSSE with this nature run, the experiments and performance metrics are simplified because the truth is known. This would not be the case if initial experiments were to use actual observations because the truth from which real observations were obtained is not known. Also, by using the same model in the generation of the nature run as is used in the data assimilation experiments (known as a "perfect twin" experiment) we can ignore the effects of model and observation bias that must be handled with actual observations. A drawback to this approach that must be kept in mind is that since

model and observation biases are not introduced it is possible that results from such a system are not actually relevant to an operational system using real observations.

2.1.1 Data assimilation

Data assimilation will be accomplished by using the Local Ensemble Transform Kalman Filter (LETKF). Two separate systems will be used, one for the atmosphere and one for the ocean. For the atmosphere, the SPEEDY-LETKF (Miyoshi, 2005) has widely been used for past data assimilation experiments and will be used here with only minor modifications. For the ocean, the NEMO-LETKF has been developed, but it is based mostly on the ocean LETKF developed by Penny et al. (2013). The general structure of how these computer programs will be connected to form a strongly coupled data assimilation system are shown in Fig 2.1, with mathematical justification given by Eq 1.1-1.5. The code for each domain is actually two executables, OBSOP and LETKF. The OBSOP program first performs the observation operator on the ensemble model state, $\mathbf{y}^b = H\mathbf{x}^b$, transforming the state to observation space. In this case OBSOP is specific to the domain. The SPEEDY-LETKF OBSOP will only process atmospheric observations, and the NEMO-LEKTF OBSOP will only process the ocean observations. Once the observation departures have been calculated, the next executable, LETKF, can use the departures regardless of which domain it comes from. The NEMO-LETKF can use the output from SPEEDY-LEKTF OBSOP and vice-versa. The only modification to the SPEEDY-LETKF and NEMO-LETKF LETKF codes required is to inform the solver how the

vertical localization should be performed for observations from the opposite domain.

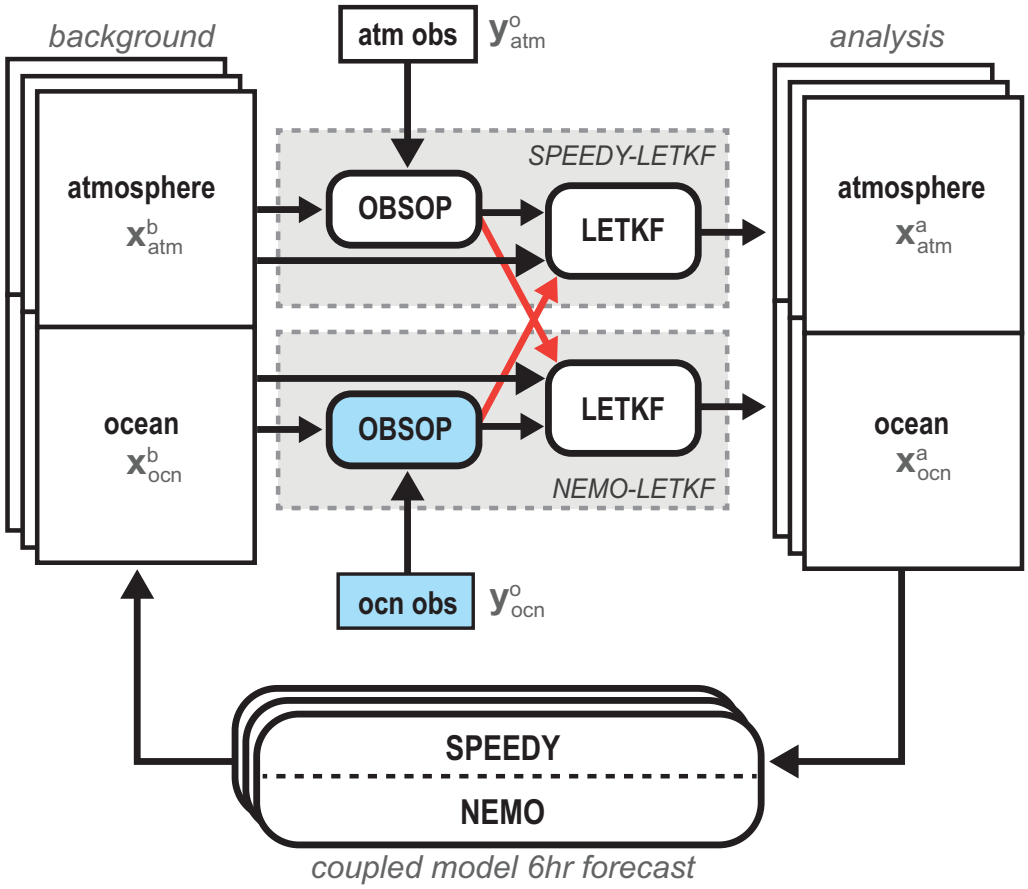

Figure 2.1: Schematic of the LETKF configuration for the coupled SPEEDYNEMO system. Shared observational departures (red arrows) between the separate LETKF systems enable them to effectively perform as a single strongly coupled DA system.

It is possible to alternatively have a single strongly coupled SPEEDYNEMO-LETKF that handles the entire system in one executable. However, it is preferable to use the approach given here for several reasons. By keeping the LETKF code separate for the two domains, the code is simpler and easier to follow. Since the atmosphere and ocean often are run with different grid resolutions, the code is kept

cleaner by keeping the LETKFs separate. The benefits might seem small for for just a two domain ocean-atmosphere SCDA system, but the complexity would quickly add up for a full Earth system model. The land, atmosphere, ocean, ice, wave, etc, can all be coded as separate LETKFs that are then made into a SCDA system just by sharing the observation operator output.

It was found that the respective SPEEDY-LETKF and NEMO-LETKF codebases contained an unnecessarily large number of places with hard-coded variable names, which slowed down transition from WCDA to SCDA experiments. Because of this, a prototype for a universal LETKF that should work with any geophysical system was created. This publicly available software is a complete re-factorization of the original Miyoshi code, termed the universal multi-domain LETKF (UMD-LETKF, Sluka (2018b)) and should allow a future researcher to experiment with SCDA without having to change any of the LETKF code. All state variables, domains, parameters, etc., are controlled by a set of configuration files. The UMD-LETKF software is described in more details in Chapter 5, and its development continues as it is made ready for adoption by the marine modeling group at NCEP.

Inflation methods

To account for errors in the background estimate, covariance inflation is typically required to keep the ensemble spread from becoming too small or even collapsing. There are several choices for covariance inflation including a constant multiplicative factor (Anderson, 2001), adaptive multiplicative (Miyoshi, 2011), addi-

tive (Houtekamer and Mitchell, 2005), relaxation to prior background (Zhang et al., 2004), and relaxation to prior spread (Whitaker and Hamill, 2012). Also, during the model integration stochastic parameterizations (Shutts, 2005; Berner et al., 2009) can be used to increase the ensemble spread. Additive inflation is typically more important for dealing with model error (Whitaker and Hamill, 2012). Since we are using a perfect model OSSE, a form multiplicative inflation will instead be used.

Several of these inflation methods have been previously incorporated into the LETKF code, but the ones already provided create problems with the ocean. Constant multiplicative inflation is the simplest, but this method causes problems with the ocean-LETKF. With a constant multiplicative factor, unobserved regions of the ocean (such as the southern hemisphere before the satellite era, or the deeper ocean) will have their spread continue to grow unbounded and will eventually blow up unless artificial bounds are placed on the spread at various locations and depths.

Adaptive inflation (Miyoshi, 2011) is often used for data assimilation experiments with the SPEEDY-LETKF. This method estimates an evolving spatially varying multiplicative covariance inflation factor, and often works well assuming the observation network does not change in time and that model grid-point locations with an observation have **all** of its state variables observed. Since the same inflation factor is applied to all state variables, the unobserved state variables can "blow-up", unless variable localization is used in the data assimilation. Temperature and salinity are typically the only widely sampled variables in the ocean, and adaptive inflation therefore has a tendency to cause the spread of the ocean currents to grow unbounded and fails to work.

It is for these reasons that relaxation to prior spread (RTPS) method of Whitaker and Hamill (2012) was added to the LETKF. This method, shown by eq 2.1, inflates the spread of the analysis, σ^a, some percentage, α, back toward the original spread of the background, σ^b, where $\mathbf{x}_i^{'a}$ are the analysis ensemble perturbations for each ensemble member i.

$$\mathbf{x}_i^{'a} \leftarrow \mathbf{x}_i^{'a} \left(\alpha \frac{\sigma^b - \sigma^a}{\sigma^b} + 1 \right) \tag{2.1}$$

This method is simple to implement and has the benefit of inflating the analysis more where observation density is higher (similar to what adaptive inflation will do), but without over-inflating as adaptive can do if there are no directly observed variables. RPTS was chosen over the similar relaxation to prior perturbations (RTPP) (Zhang et al., 2004) in part because it was shown by Whitaker and Hamill (2012) that RTPS is less sensitive to the choice of the α value. RTPP was shown to blow-up if α is chosen from outside a narrower range of acceptable values, however, this result was found with a different flavor of EnKF and might not hold true with the LETKF.

2.1.2 SPEEDYNEMO model

The SPEEDYNEMO model is a simplified ocean-atmosphere coupled model that was developed by Kucharski et al. (2015) to investigate the role of the Atlantic Multidecadal Oscillation (AMO) on the tropical Pacific. The primary benefit of this model is that the configuration and parameterizations are comprehensive enough to create realistic atmospheric phenomena while at the same time is computationally cheap. For example, the atmospheric component by itself is able to perform a one year simulation in a mere 6 minutes on a single core of a standard desktop computer (Kucharski et al., 2013). This intermediate-complexity model enables ensemble experiments with a fast turn around time and is therefore chosen for the strongly-coupled OSSEs.

The atmospheric component consists of the Simplified Parameterization, primitivE Equation DYnamics (SPEEDY) model, version 41 (Molteni, 2003; Kucharski et al., 2006). SPEEDY is a hydrostatic, eight-level sigma coordinate spectral model with T30 resolution and is capable of producing fairly realistic phenomena despite the simplified parameterizations. This model is a small upgrade from the SPEEDY model used for the initial LETKF experiments (Miyoshi, 2005), primarily using an increase in the number of vertical levels from 7 to 8.

The ocean component consists of the Nucleus for European Modeling of the Ocean (NEMO) (Madec, 2008). NEMO is configured with the ORCA2 grid, a 30 level vertical z-coordinate grid with a 2° horizontal tripolar grid that tapers to 0.25° at the equator to capture equatorial wave dynamics.

The SPEEDY and NEMO models are coupled by exchanging SST from the ocean to the atmosphere, and total heat flux, shortwave radiation, wind stress, and evaporation minus precipitation from the atmosphere to the ocean. This exchange was originally performed by the OASIS coupler, though the coupler was removed for simplicity due to the fact that the coupling period (6 hours) would be the same as the data assimilation period and so the data assimilation code is used as the coupler. If a period other than this is required the OASIS coupler will have to be reinserted. The original SPEEDYNEMO has been modified to produce instantaneous model output every 6 hours. The LIM ice model that comes with NEMO is turned off and sea ice distribution is prescribed by using observed monthly climatology from ERA-15 (Gibson et al., 1999)

In order for the model to produce ENSO-like patterns when run as a freely running nature run, a flux correction has to be applied. Following Kröger and Kucharski (2011) a one-way anomaly coupling is applied from the ocean to the atmosphere. This corrects a cold bias in the East Pacific. The reduced East-West SST gradient allows for El Niño-Southern Oscillation-type variability to occur (fig 2.2). However, the model has an ENSO pattern with a lower frequency and more persistent signal than what should be observed (fig 2.3). Also, the SST anomalies are slightly too far to the west and do not move northward along the coast.

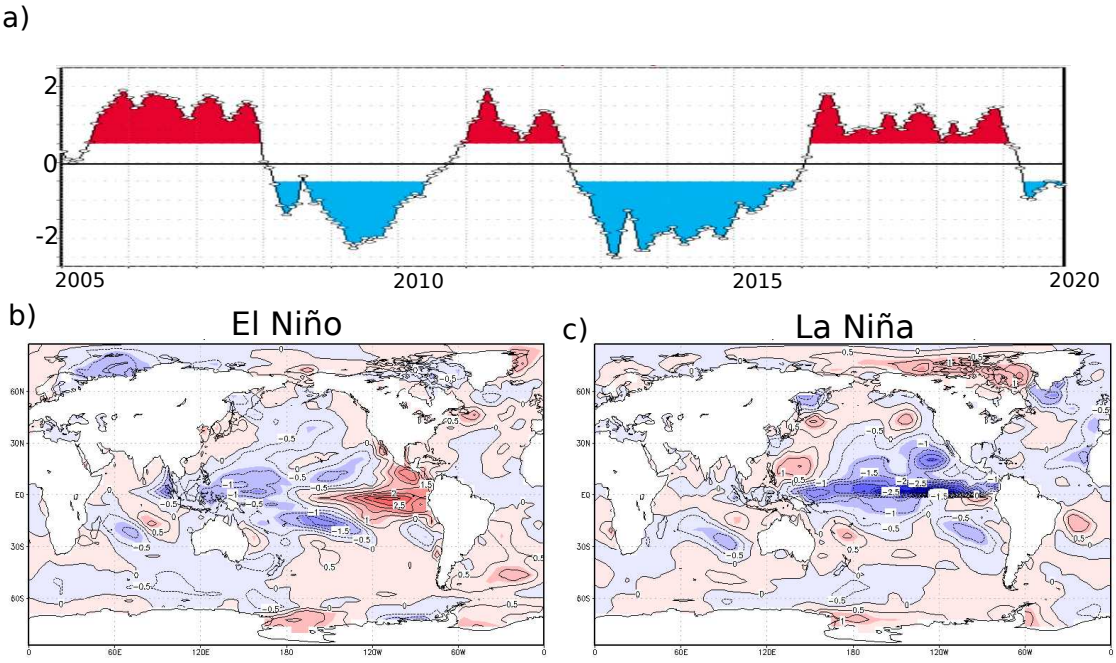

Figure 2.2: The spontaneous ENSO-like pattern seen in the SPEEDYNEMO during a long nature run. Shown are the SST anomaly over the Niño 3.4 region (a) and an example of a typical El Niño phase during 2007 (b) and La Niña phase during 2008 (c)

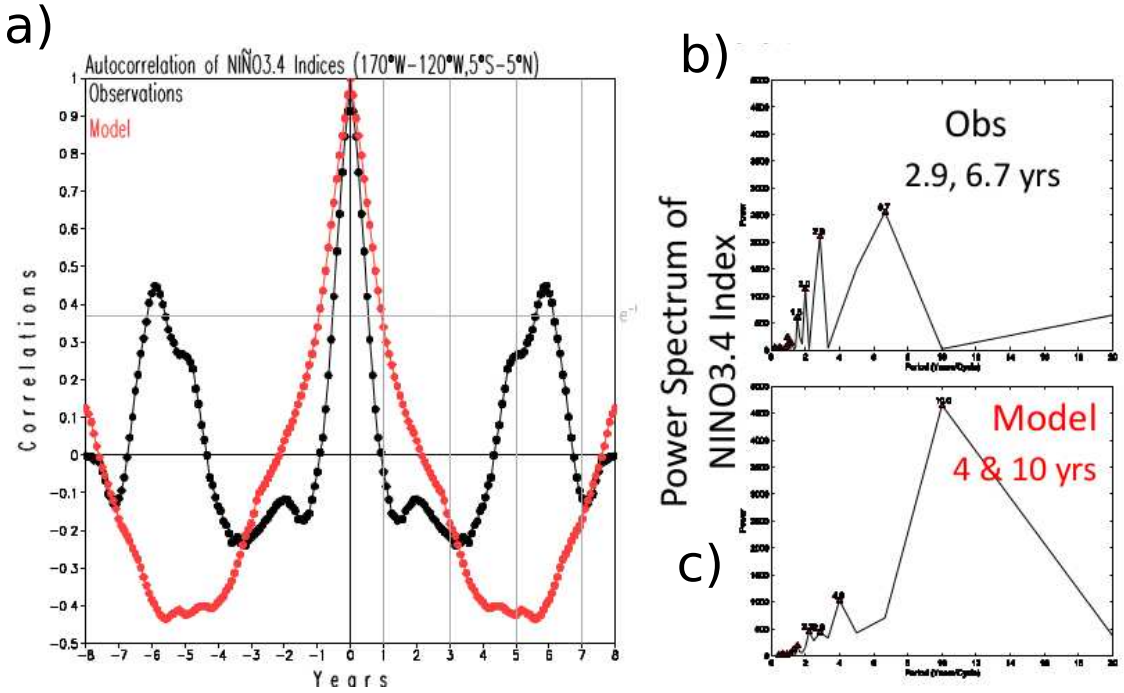

Figure 2.3: Auto-correlation of NINO3.4 indices (a) for observed conditions (black) and the SPEEDYNEMO model nature run (red). Power spectrum of observations (b) and the SPEEDYNEMO model (c). SPEEDYNEMO is shown to produce ENSO phases that last too long and at a period longer than observed in nature. Image courtesy of Alfredo Ruiz-Barradas.

2.2 One-way SCDA with atmospheric observations

Perfect model observation system simulation experiments (OSSEs) are conducted using synthetic atmospheric observations. First, a long free run of the model is performed and used as the truth for the remainder of the experiments. To generate this nature run, SPEEDY-NEMO is first initialized with climatological ocean temperature and salinity, an atmosphere at rest, and run freely for 20 years to spin-up. The subsequent 6 years are then saved as the nature run.

From this run, synthetic rawinsonde observations and satellite retrievals are generated every 6 hours at the locations shown in fig 2.4. This provides observations of surface pressure (Ps) and vertical profiles of temperature (T), humidity (q), and wind(U,V). For each observation the values at the appropriate times and positions of the nature run are used and independent Gaussian errors are added with zero mean and unit standard deviation (1 hPA, 1° C, 1 g/kg, and 1 m/s). It should be emphasized that for this first OSSE no ocean observations are generated or assimilated. For simplicity, observations are only generated at the analysis times, though a 4D-LETKF that uses observations throughout a window would be expected to perform similarly. Also, these observations are generated at the exact grid-points, and so no interpolation is needed from the observation operator.

Two runs of the data assimilation system are performed, one with weakly coupled DA (WEAK) and a second with strongly coupled DA (STRONG). For WEAK the atmospheric observations are assimilated only into the atmosphere by the SPEEDY-LETKF. NEMO-LETKF is not run and the ocean is updated every 6

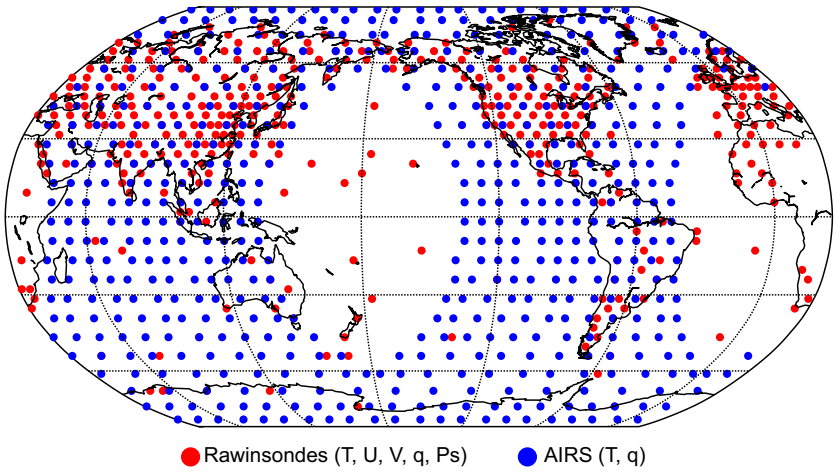

Figure 2.4: Locations of atmospheric observations for SPEEDYNEMO OSSE over a single 6 hour period. The locations of the AIRS-like T and q satellite observations changes with each 6 hour time period to provide global coverage.

hours only through the normal flux exchanges of heat, momentum, and evaporation minus precipitation. This method is the standard way a coupled ocean-atmosphere system would be run given only atmospheric observations.

For the second data assimilation run with strongly coupled DA (STRONG), both SPEEDY-LETKF and NEMO-LETKF are given the atmospheric observation departures, thereby allowing the atmospheric observations to be assimilated into the ocean. In this case the ocean state is corrected both by the fluxes from the atmosphere during the model integration and by the data assimilation that is performed every 6 hours.

It should be emphasized that in both cases the atmosphere ($\mathbf{x}^a_{atm}$) and the ocean ($\mathbf{x}^a_{ocn}$) analyses are generated separately by the respective SPEEDY-LETKF and NEMO-LETKF codes, though for STRONG this is mathematically identical to having a single LETKF handling the entire state ($\mathbf{x}^a$) due to the fact that cross-

horizontal localization	1000 km
vertical localization (atm)	$0.1 \ln(\Delta P)$
vertical localization (ocn)	none
ensemble size	40
inflation	RTPS ($\alpha_{ocn} = 90\%$ $\alpha_{atm} = 60\%$)

Table 2.1: Data assimilation parameters used by the SPEEDYNEMO OSSEs

domain observational departures are shared with the two systems.

Starting with an arbitrarily labeled date of January 1, 2005, both experiments are initialized with identical ensemble members that are randomly chosen from subsequent years of the nature run. This gives the initial conditions sufficient error and spread from which to start the experiments. STRONG and WEAK experiments are then run for 6 years using the data assimilation parameters summarized by table 2.1. A horizontal localization radius of 1000km is used in both the atmosphere and the ocean. Vertical localization in the atmosphere is carried out by each model level so that observations at one level only have minimal impact on the levels above and below it.

No vertical localization is used in the ocean, the entire water column is therefore able to be updated by the atmospheric observations at the lowest level of the atmosphere in STRONG. It has been shown that while vertical localization is very important for atmospheric data assimilation, the ocean performs better without any vertical localization (Penny et al., 2015). In fact, not using any vertical localization in the ocean is computationally more efficient, since the analysis needs to only be calculated once for the entire column instead of separately at each level, and is more able to produce an analysis where the water column remains in hydrostatic balance.

The relaxation to prior spread method, or RTPS (Whitaker and Hamill, 2012), is used and has been tuned with different values for the atmosphere ($\alpha_{atm} = 0.6$) and the ocean ($\alpha_{ocn} = 0.9$) so that in the STRONG run an ensemble spread of similar magnitude to the root-mean-square-error is maintained. The use of different RTPS values for the two domains is not ideal, however. A major benefit of SCDA is that it should improve balance at the domain interface, but this balance may be disturbed by not keeping the ensemble perturbations matched up perfectly. Preferably, other methods of increasing ensemble spread for the ocean, such as using a higher eddy-permitting resolution or stochastic perturbations, should be used so that an identical RTPS value can be used for both domains. These changes are not practical for the following experiments, though.

Results

The root mean square error (RMSE), as given by eq 2.2, is used as the predominant verification method for the performance of the experiments. The RMSE is calculated for the ocean temperature, salinity, and sea surface height, as well as the atmospheric variables.

$$RMSE = \sqrt{\frac{1}{n}\sum_{j=1}^{n}(x_j - \hat{x}_j)^2} \tag{2.2}$$

The difference in analysis RMSE as compared to the nature run truth for STRONG minus WEAK (fig 2.6) shows that the ocean is significantly improved when strongly coupled DA is used to assimilate atmospheric observations into both

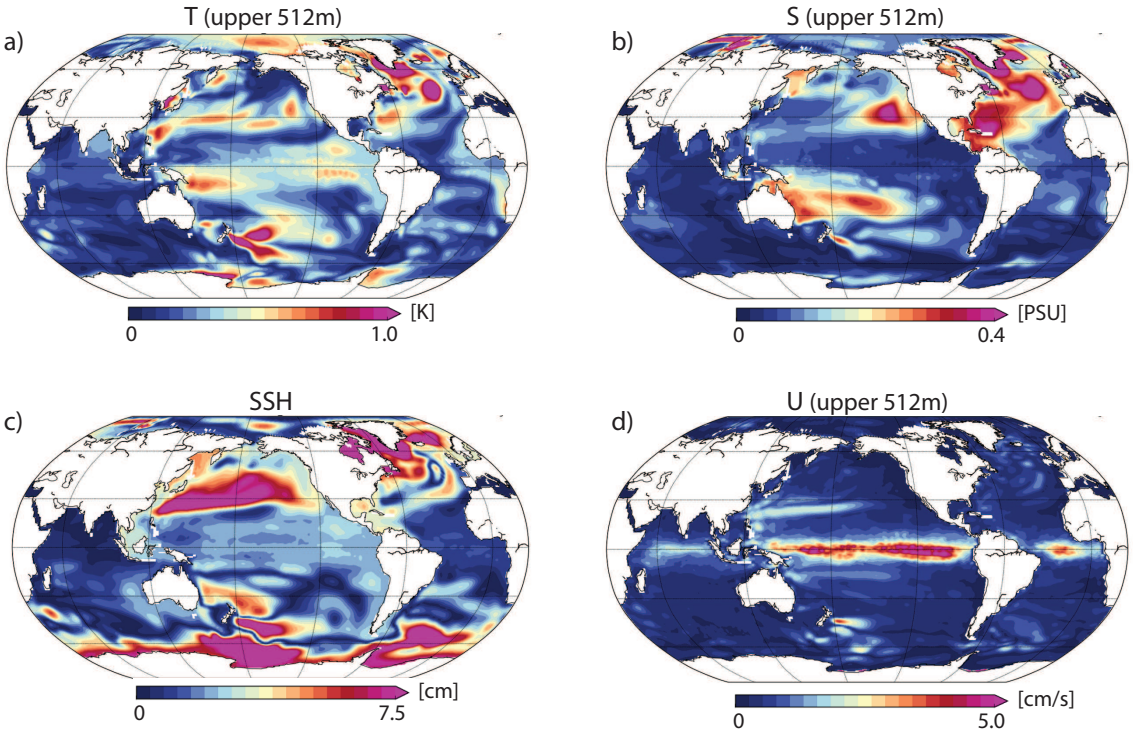

Figure 2.5: Temporally averaged ocean analysis RMSE for the WEAK run. Shown are upper temperature (a), salinity (b), sea surface height (c), and zonal currents (d).

the ocean and atmosphere at the same time. The near-surface temperatures and SSH RMSE are reduced compared with WEAK results by about 50% after an initial spin-up period of just a couple of weeks. The Northern Hemisphere (NH) and tropics, which have the largest initial errors in the WEAK run (fig 2.5), also improve the most in STRONG.

Ocean salinity errors are reduced more slowly than temperature, but this reduction continues for several years throughout the duration of the experiment. Globally, the strongly coupled DA reduces errors in salinity and temperature an average of 46% for the upper ocean over the last 5 years of the experiment.

Annual variations in the RMSE reduction by STRONG can be seen at the ocean surface (fig 2.6a). The northern hemisphere mid-latitudes experiences the greatest improvement in SST during the winter/spring months, averaging 52% over the last 5 years while only 37% over the summer months. The same is true for the southern hemisphere, though with a smaller amplitude. These results could be expected due to stronger mid-latitude atmospheric dynamics driving the ocean during the winter and spring months, as well as well as a deeper mixed layer depth in the winter.

Figure 2.7 shows the spatial patterns of analysis RMSE reduction between the two cases. The ocean state is improved most in the NH midlatitudes where the greatest density of atmospheric observations are and where the ocean is generally considered to be driven by weather anomalies. The NEMO-LETKF is configured to use no vertical localization in the ocean, which enables observations above the ocean to impact the entire water column, accelerating the improvement of the barotropic mode of the ocean. The strongest improvements in the northern Atlantic extend down below 2.5 km. Although SST errors are not reduced significantly in the tropical Pacific, RMSE errors of the subsurface waters in the upper 250m are reduced by about 1 $^\circ$ C.

Assimilating atmospheric observations into the ocean corrects the sea surface

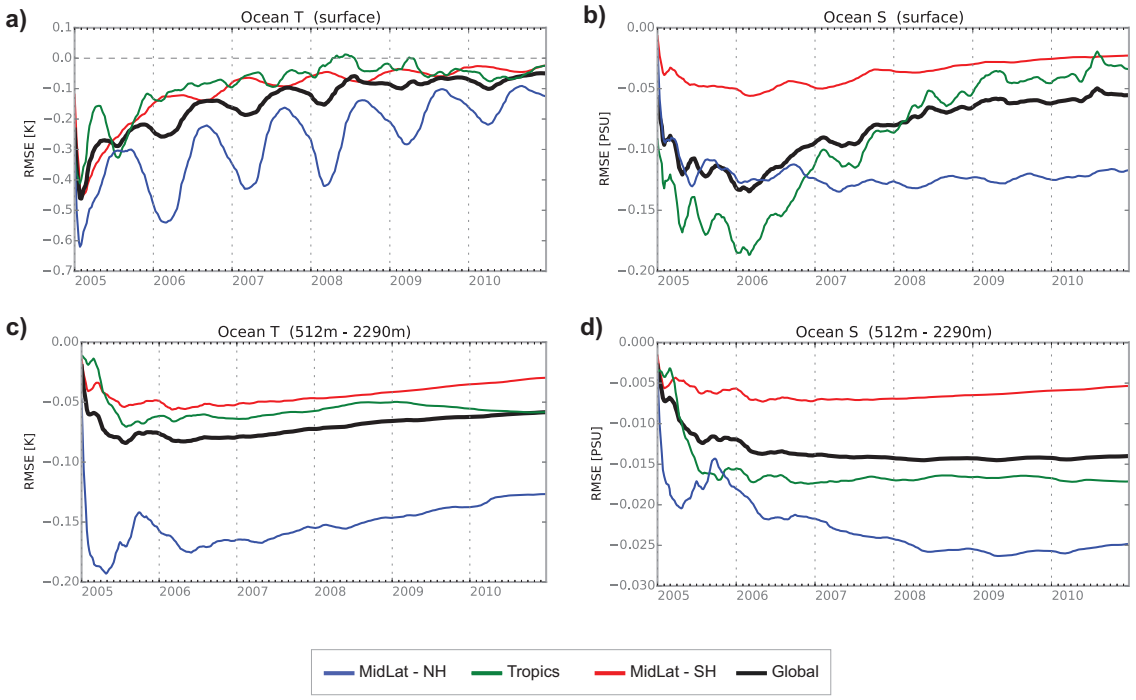

Figure 2.6: Spatially averaged difference of analysis RMSE for STRONG-WEAK using only atmospheric observations. Negative values indicate improvements by strongly coupled DA. Shown are results averaged over the Northern Hemisphere mid-latitudes (blue), tropics (green), and Southern Hemisphere mid-latitudes (red). Shown are temperatures (a and c) and salinity (b and d) at the surface (a and b) and deep ocean (c and d). Sluka et al. (2016)

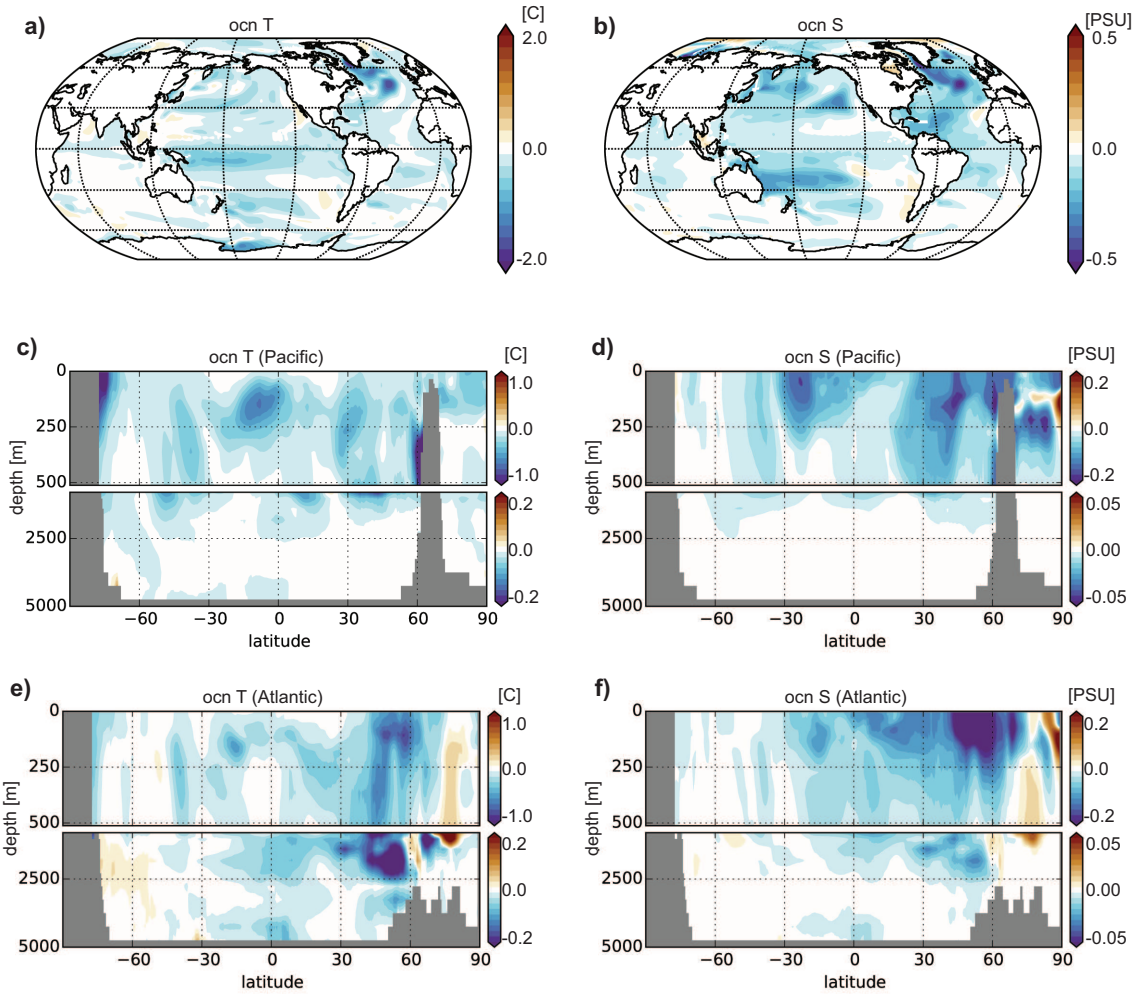

Figure 2.7: Time mean difference of analysis RMSE for STRONG-WEAK using only atmospheric observations. Negative values (blue) indicate improvements by strongly coupled DA. Shown are temperatures (a,c, and e) and salinity (b,d,and f) over the upper 500 meters (a and b) and cross sections of the Pacific (c and d) and Atlantic (e and f) basins. Sluka et al. (2016)

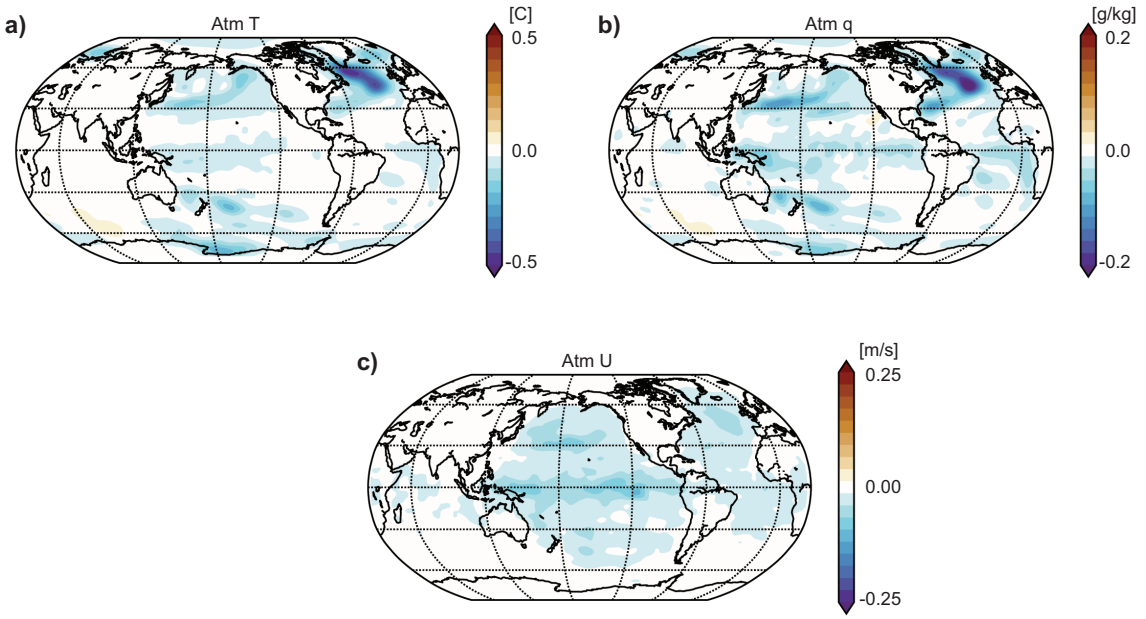

Figure 2.8: Time mean difference of analysis RMSE for STRONG-WEAK using only atmospheric observations. Negative values (blue) indicate improvements by strongly coupled DA. Shown are temperature (a) and humidity (b) at the lowest model levels, and zonal wind speed through the troposphere (c). Sluka et al. (2016)

temperatures in STRONG which then reflect back on the atmosphere, resulting in a reduction in atmospheric RMSE in STRONG, as shown in figure 2.8. Improvements in atmospheric temperature and humidity at the lowest model levels overlap the same areas of the ocean (fig 2.7) experiencing corrected SSTs. Precipitation and other fluxes are all improved in these areas as well. Zonal winds are improved throughout the troposphere of the tropical Pacific, presumably from an improved Walker circulation, as well as over the oceanic NH mid-latitudes.

In addition to the STRONG and WEAK cases using all atmospheric observations, a similar experiment was performed using only rawinsonde observations (fig 2.9). Although extremely few observations are directly over the oceans, the strongly coupled data assimilation was still able to provide similar improvements in

most regions, except for the Southern Hemisphere where there are too few rawinsondes.

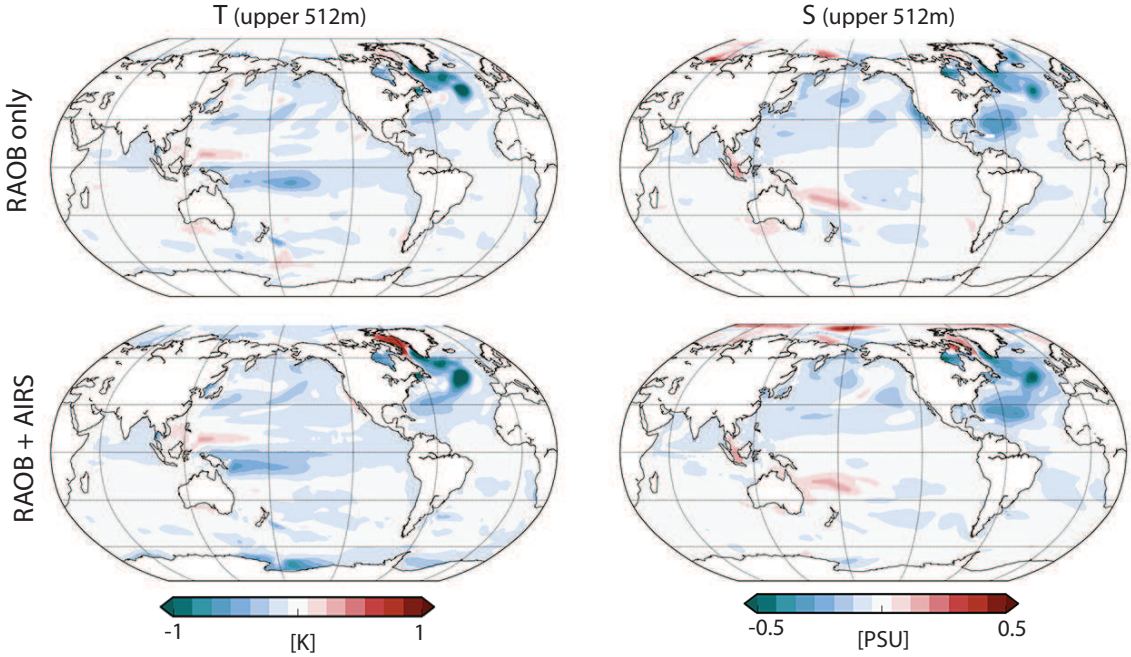

Figure 2.9: Time mean difference of analysis RMSE for STRONG-WEAK using only rawinsonde observations (top) compared with the rawinsonde and satellite observations (bottom). Shown are RMSE in upper 512m for ocean temperature (left) and salinity (right).

2.3 One-way SCDA with ocean observations

Experiments similar to the previous atmospheric-only observation experiments are repeated, although this time using only ocean observations in both weakly and strongly coupled settings. The same nature run is utilized, and the STRONG vs WEAK experiments are run over the same 2005-2010 time period. This configuration tests what impacts ocean observations have on the atmosphere when using strongly coupled data assimilation. Synthetic observations are generated to mimic temperature and salinity profiles from Argo floats (Roemmich et al., 2009) and satellite based sea surface temperature platforms. Argo floats provide temperature and salinity profiles down to about 2km with nearly global coverage, although with a sparse 300 or so profiles a day, starting in the early 2000's. With fewer observations now being assimilated into the ocean and atmosphere, less inflation was required. The RTPS value was reduced from 90% to 50% for both the atmosphere and ocean. Horizontal localization for the ocean is 700km at the equator, decreasing to 200km at the poles. Atmospheric horizontal localization of 1000km was used.

Roughly 120 temperature and salinity profiles each day are generated with random locations to simulate the Argo profiles (fig 2.10). These positions are pseudo-randomly generated, with less likelihood of observations being generated in the polar regions, or in areas where the ocean depth is less than 2km. The number generated is less than the actual number of operational Argo profiles that are currently present in reality, which consists of roughly 3000 floats that give profiles every 10 days, for a total of 300 profiles a day. This reduction is justified given the reduced resolution

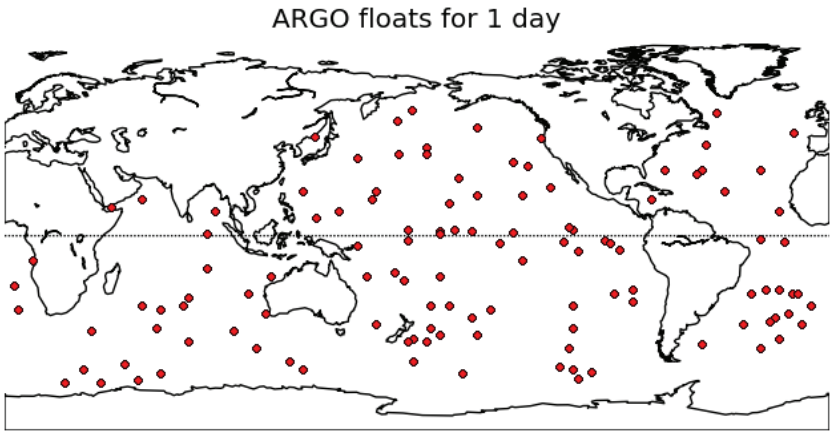

Figure 2.10: Example location of synthetic Argo T/S ocean profile observations for one day. The locations are randomly generated to cover an area similar to that of the actual Argo system with about 200 profiles per day.

and variability of the low resolution NEMO system being used. The values generated from the nature run have Gaussian noise, with mean of zero, and a depth dependent standard deviation as given by fig 2.11. This simulates increased errors near the thermocline. Synthetic satellite based SST observations are generated by randomly selecting 1/4 of the ocean grid points each day, and adding 1 degree error.

Results of the RMSE differences for STRONG vs WEAK are shown in fig 2.12. Overall, it can be seen that the error in the heat flux between the atmosphere and ocean is improved globally when ocean observations are assimilated into the atmosphere. The mid latitudes require a spin-up of several months before SCDA is able to outperform WCDA, but then after this period these fluxes are consistently better throughout the remainder of the experiment. Initially, the other atmospheric variables plotted (humidity, temperature, and wind speed, all at the lowest model level) show reduced errors with SCDA for the first 2 years, but then SCDA performs

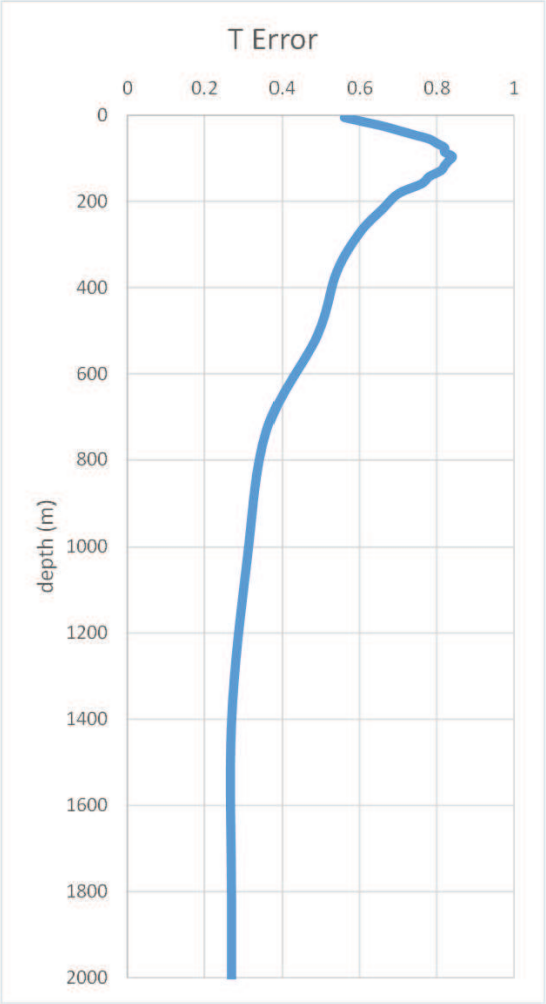

Lvl	Depth (m)	T Err	S Err
1	4.999938	0.56	0.056
2	15.00029	0.61	0.061
3	25.00176	0.66	0.066
4	35.00541	0.7	0.07
5	45.01332	0.74	0.074
6	55.0295	0.78	0.078
7	65.06181	0.8	0.08
8	75.12551	0.82	0.082
9	85.25037	0.82	0.082
10	95.49429	0.84	0.084
11	105.9699	0.83	0.083
12	116.8962	0.82	0.082
13	128.6979	0.81	0.081
14	142.1953	0.78	0.078
15	158.9606	0.76	0.076
16	181.9628	0.7	0.07
17	216.6479	0.66	0.066
18	272.4767	0.6	0.06
19	364.303	0.54	0.054
20	511.5348	0.49	0.049
21	732.2009	0.36	0.036
22	1033.217	0.31	0.031
23	1405.698	0.27	0.027
24	1830.885	0.27	0.027

Figure 2.11: Argo profile errors are randomly generated with a Gaussian distribution using depth dependent standard deviation given in the table.

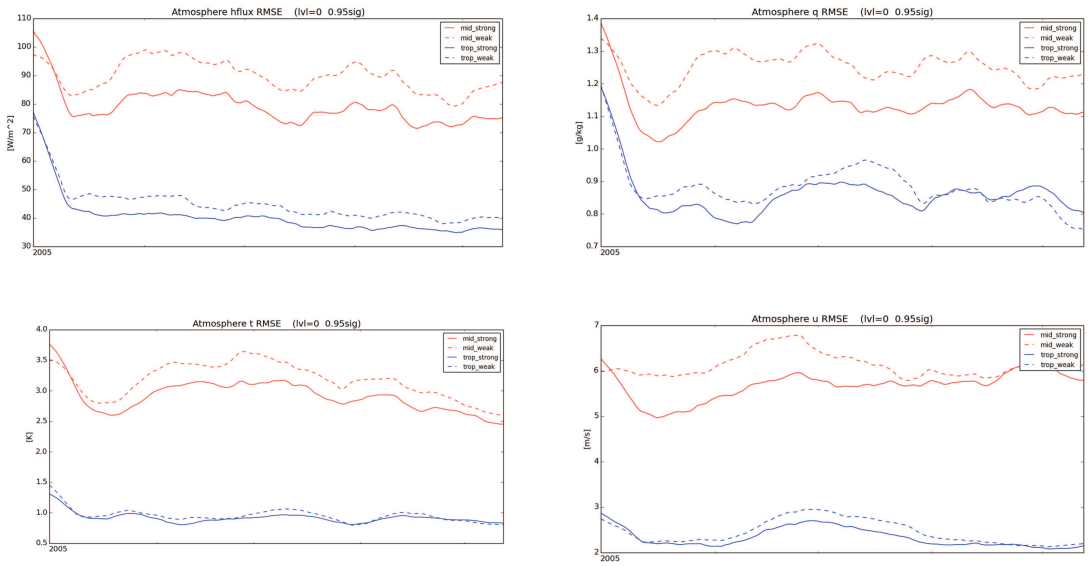

Figure 2.12: Spatially averaged analysis RMSE for ocean observations only experiments. Shown are results for STRONG (solid) and WEAK (dashed) for both the tropics (blue) and mid-latitudes (red). Atmospheric variables are heat flux (a), and humidity (b), temperature (c), and zonal wind (d), at the lowest levels.

the same as WCDA in the tropics. In the extra-tropics, temperature and humidity continue to have lower RMSE with SCDA compared with WCDA throughout the duration of the experiment.

The results given in fig 2.13 show that in these experiments the atmospheric observations have a slightly larger impact with SCDA in the tropics (especially zonal winds), whereas ocean observations have a stronger impact in the extra-tropics.

This latitudinal variation in the effects of SCDA can possibly be explained by an observed latitudinal variation of which domain typically dynamically drives the other. Several studies (e.g. Pena et al. (2003); Ruiz-Barradas et al. (2017)) have sought identify the spatial variation of which domain is the main driver of the coupled system, the atmosphere or the ocean, by examining the coupled anomalies

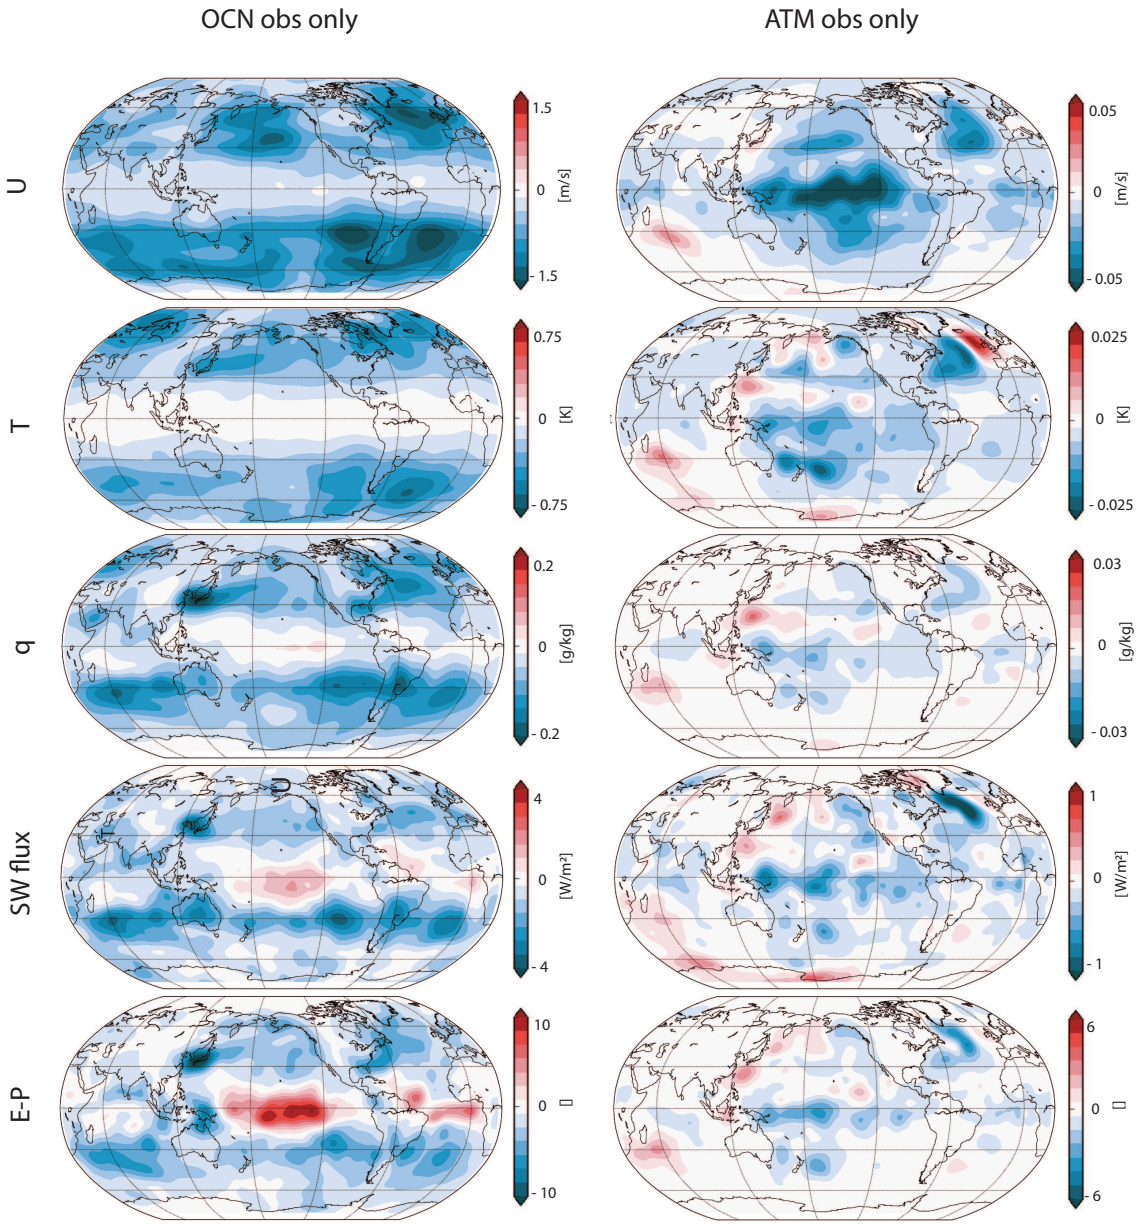

Figure 2.13: Temporally averaged difference of analysis RMSE for STRONG-WEAK for several atmospheric variables. Shown are experiments using ocean observations only (left) and atmospheric observations only (right). Variables shown are at the lowest model level for zonal wind (U), temperature (T), humidity (q), as well as downward shortwave radiation (SWFlux) and evaporation minus precipitation (E-P). Ocean observations mainly show improvement in the extra-tropics, whereas atmospheric observations show improvements in the tropics.

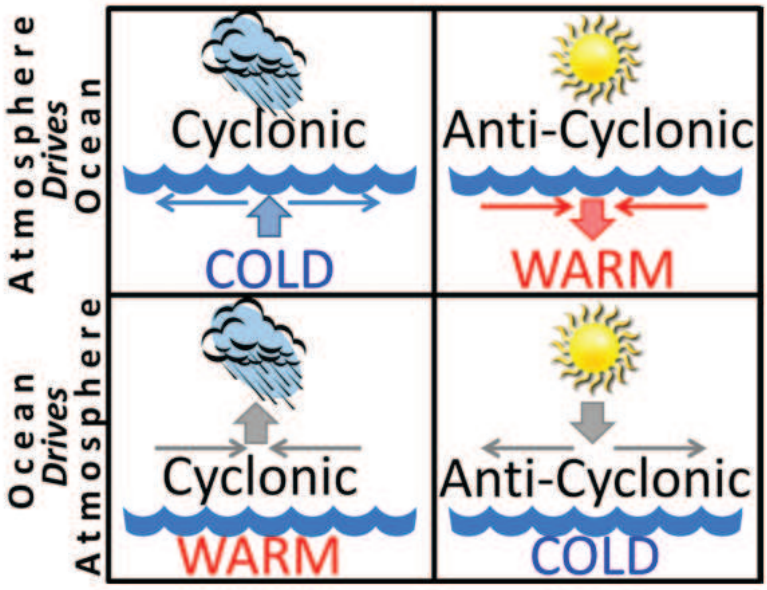

Figure 2.14: Schematic of local dynamical ocean-atmosphere coupling. Shown are instances where atmospheric anomalies drive the ocean (top) and where ocean anomalies drive the atmosphere (bottom). Ruiz-Barradas et al. (2017)

in atmospheric vorticity at 850mb and ocean SST. Figure 2.14 summarizes these coupled anomalies. The atmosphere typically drives the ocean in the mid-latitudes, where a cyclonic atmosphere is associated with storminess which drives ocean cooling through reduced shortwave radiation and Ekman upwelling in the ocean. Conversely, an anticyclonic atmosphere is associated with higher downward radiation and therefore a warmer ocean. The opposite is true in the tropics, where the ocean is typically the main dynamic driver of the coupled system. Warm SST anomalies lead to convection and subsequent cyclonic vorticity in the atmosphere, while cold SSTs lead to an anti-cyclonic atmosphere.

With data assimilation, the flow of information across the ocean-atmosphere domain interface is accomplished in two ways:

1. the previously explained dynamical forcing which is mediated by the flux exchange during model integration

2. by performing SCDA and utilizing the coupled covariance to allow observations in one domain to impact the other domain

Since the atmosphere is the main dynamical driver of the system in the mid-latitudes, the impact of atmospheric observations assimilated into only the atmosphere are inherently felt by the ocean anyway even with WCDA thanks to the atmosphere passing that information to the ocean via the fluxes, and so assimilation of atmospheric observations into the ocean under SCDA have little impact in the mid-latitudes. On the other hand, if an ocean observation is assimilated into the ocean only with WCDA, then that information has no way of improving the atmosphere, since the ocean is not dynamically driving the atmosphere in this location. The only way to have an ocean observation improve the atmosphere is through SCDA. The coupled covariance enables the ocean observation to correct the atmosphere, by choosing the atmosphere that led to that ocean state. In the tropics everything is reversed, and atmospheric observations have a bigger impact under SCDA due to the ocean being the main dynamical driver of the atmosphere.

The above results, however, are complicated by the fact that the SPEEDY atmospheric model has a tendency to easily become unstable during the course of data assimilation cycles. This problem is more likely to happen for SPEEDY in cases involving an imperfect model or imperfect boundary conditions. The SCDA and WCDA experiments performed are inherently using imperfect boundaries, since the

SST during the data assimilation run will necessarily not be the same as that of the nature run. These types of instabilities developing in the atmospheric model are the result of an accumulation of energy at the shortest wavelengths in temperature and humidity at the lowest model levels (e.g. fig 2.15) and were found by several other students, Cathy Thomas (Thomas, 2017) and Yan Zhou (Zhou, 2014). It appears that a regular grid of observations, along with an imperfect model, are enough to force the problem. Neither student was able to eliminate the spurious high energy waves completely, but it was found that performing an additional spectral truncation of the analysis, or performing the analysis at a resolution lower than the forecast were helpful methods.

Even with attempting additional spectral truncation of the atmospheric state by removing all spectral coefficients with wave-number $n >= 30$, SCDA experiments with ocean observations assimilated into the atmosphere were not able to be run for indefinite periods of time without the atmospheric model blowing up. The results in fig 2.12 and fig 2.13 are therefore not completely conclusive. The atmospheric surface pressure appears fine for the first year or two before the high frequency variations in tropical pressure become visible. However, even with surface pressure becoming unstable, the other fields plotted show little difference over the years, meaning the results might be trusted somewhat. These instabilities, and the inability by several students to overcome them, are the major reasoning for moving on to more realistic SCDA tests with the CFSv2 model as described in the subsequent chapters.

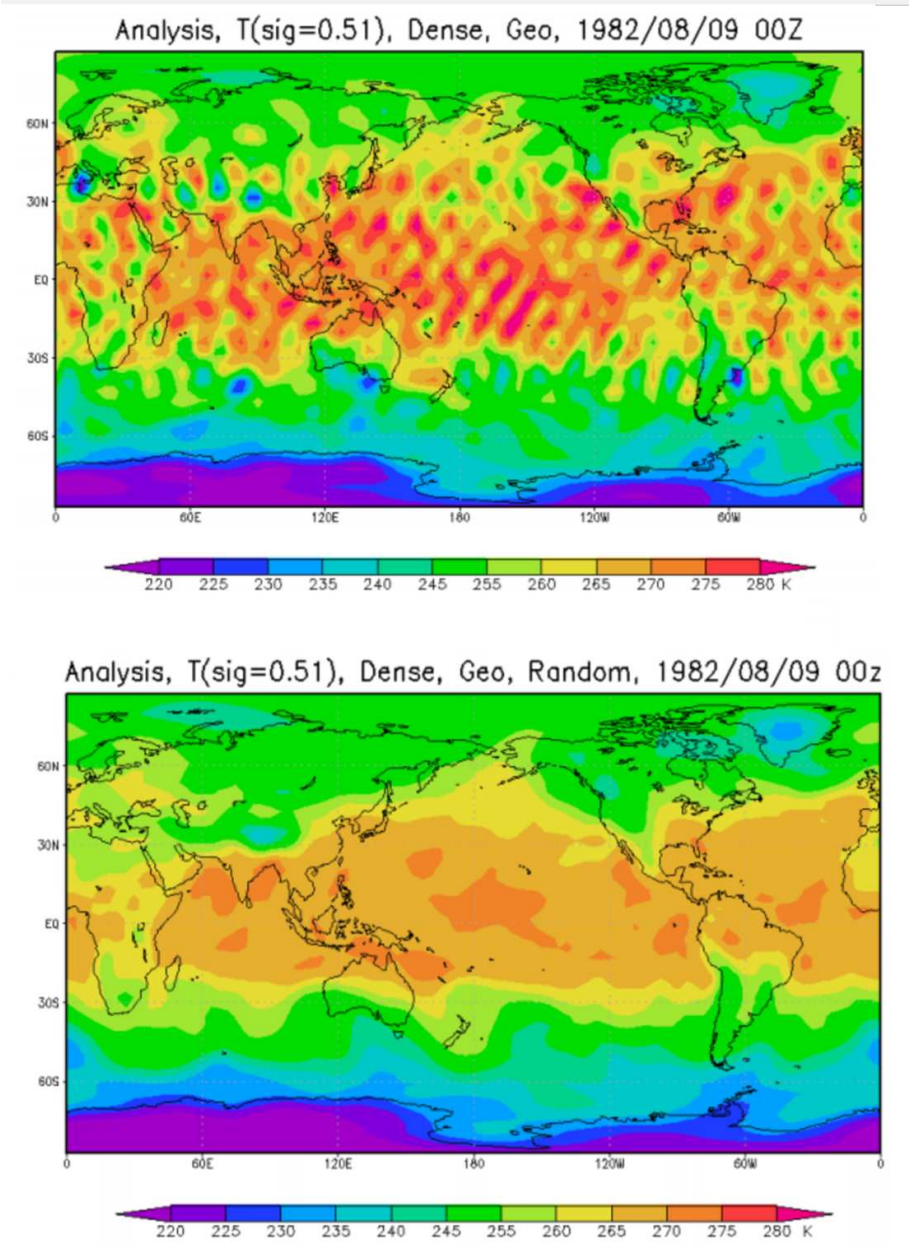

Figure 2.15: Instance of SPEEDY model exhibiting growing standing waves in the temperature field at the $\sigma = 0.5$ level. These waves eventually cause the model to blow up (top), compared with a normal run (bottom). Image courtesy of Cathy Thomas.

2.4 Summary

By performing strongly coupled DA, where the ocean-atmosphere states and observations are effectively treated as a single system, improvements can be seen in both domains compared to weakly coupled DA in specific situations. Sharing the ensemble observational departures between the separate SPEEDY-LETKF and NEMO-LETKF systems takes advantage of the cross-domain background error covariance and allows atmospheric observations to directly impact the ocean. Ocean observations can similarly improve the atmosphere, though this is more difficult for long term runs with SPEEDYNEMO given the instabilities in the atmosphere.

Several findings with the SPEEDYNEMO system are interesting and should be further explored with a more realistic (and more stable) system. It is hinted at that atmospheric observations have more of an impact with SCDA in the tropics, and ocean observations having more of an impact with SCDA in the extra-topics. This latitudinal dependence on the directionality of impact for observations can possibly be explained in terms of the latitudinal dependence on who the dynamical driver of the system is. I.e, if the ocean is the dynamical driver, atmospheric observations should have more of an impact with SCDA since there is no other way for information from the atmosphere to improve the ocean during model integration. In essence, the downstream observations improve the upstream domain.

These results contradict other findings (Lu et al., 2015a), but this could be explained by the fact that the data assimilation cycle period in the experiments presented here are different. Lu et al. used longer weekly or monthly timescales,

whereas 6 hourly scales are used here. With the longer data assimilation cycles the previous assumption of "the atmospheric observations assimilated into the ocean don't help much under SCDA in the extra-topics because the atmospheric observations are felt by the ocean anyway via the flux exchange during model integration" isn't necessarily true. If the coupled anomalies have a lag period of several days then an error in the atmosphere would have enough time to lead to an error in the ocean. This error could then be corrected by assimilating the atmospheric observations into the into the ocean with SCDA.

A number of problems were found with the SPEEDYNEMO-LETKF system that limits its use in studying short, 6-hour, strongly coupled data assimilation cycles as intended. These include:

- The atmospheric SPEEDY-LETKF can be unstable when imperfect models are used. In this case, the sea surface temperature boundary conditions provided by the ocean are inherently going to be imperfect. The LEKTF is then more likely to excite atmospheric waves at the lower levels that are then not properly damped by the model.

- The coupling frequency of SPEEDYNEMO is at the same timescale as the data assimilation length. It is possibly too infrequent for a proper cross domain covariance to form for use in SCDA.

- The low resolution of the ocean model causes very little ocean variability, especially in the extra-tropics, as shown in fig 2.13. Because of this lack of variability, it is difficult to maintain a proper ensemble spread even with large

values of relaxation. The small spread becomes even more problematic when atmospheric observations, which are more numerous than ocean observations, are assimilated into the ocean.

- Due to the low variability of the ocean, the errors in the ocean when assimilating any ocean observation is very low, meaning the atmospheric observations have very little ability to improve the ocean and in the case of experiments tried SCDA actually decreased the performance when all observations are used.

Due to the above mentioned difficulties with the SPEEDYNEMO-LETKF, its use has proven less than ideal for the 6 hour cycling experiments performed. Since the SPEEDYNEMO model itself was initially designed for climate length runs, the system may be better suited for studying strongly coupled data assimilation with other phenomena at longer timescales such as the Atlantic meridional overturning circulation (AMOC), Pacific decadal oscillation (PDO), and Atlantic multidecadal oscillation (AMO), assuming these phenomena are able to be represented by the model. Such an experiment was performed by Tardif et al. (2014), though with a more simplistic model. Choices that were made earlier in the development of the SPEEDYNEMO-LETKF make changing from a 6 hour cycle difficult, and is therefore easier to switch to a different model (CFSv2) for the remainder of the experiments.

Chapter 3: Full Model OSE: CFSv2-LETKF

The goal of this section is to develop a strongly coupled ocean atmosphere data assimilation system that is geared towards an operational quality model with a realistic 6 hour or daily data assimilation cycle, such as is used by many operational centers worldwide. The system that has been built for the simplified SPEEDYNEMO in chapter 2 will be modified to utilize a realistic model, namely NCEP's Climate Forecasting System v2 (CFSv2). This should additionally alleviate some of the problems that were occurring with SPEEDYNEMO, particularly the issues of the atmosphere blowing up when the ocean was assimilated into the atmosphere, and the low resolution and insufficient ensemble spread of the ocean. This will be tested, with real observations.

The LETKF used is a combination of the already existing GFS-LETKF (Lien et al., 2013) and MOM-LETKF (Penny et al., 2013). First, an initial ensemble will be spunup for several months. Then, WCDA will be used to further spinup the ensemble members. The cross-domain correlations found in the WCDA results will then be used to determine which observations should be used in the SCDA experiment. And finally, a SCDA cycle is run over the same period that was used for the WCDA run.

3.1 CFSv2 model

The second version of the Climate Forecast System (CFSv2) is a coupled atmosphere-ocean-land model that was made operational at NCEP in March of 2011 Saha et al. (2010, 2014). It has been used since by NCEP for seasonal forecasting, retrospective reanalysis, and reforecasts. The CFSv2 has shown significant skill improvements in seasonal forecasts for coupled ocean-atmosphere phenomena such ENSO and Madden-Julian Oscillation. It consists of three separate executables that are run simultaneously, the atmosphere-land model (GFS), the ocean model (MOM4), and the coupler. The CFSv2 is used as the forecast model for the Climate Forecast System Reanalysis (CFSR) that runs from 1979 to 2011 (Saha et al., 2010), and for the realtime Climate Data Assimilation System (CDAS) that is run to the present.

The atmospheric component consists of the Global Forecasting System (GFS) run at a reduced resolution than what the standalone atmosphere for weather forecasting is run. For the CFSR a resolution of T382 is used for the 9 hour forecast used as the background in the data assimilation cycle. The subsequent seasonal forecasts use a horizontal resolution of T126. There are 64 vertical levels, in a hybrid sigma coordinate system. The 4 level Noah land surface model (Ek, 2003) is run as part of the GFS component.

The ocean component consists of the GFDL Modular Ocean Model (MOM) version 4 run at a horizontal resolution of 1/2 degree with a latitudinal spacing of 1/4 degree near the equator to better capture the equatorial dynamics. The vertical

coordinates are 40 z* levels.

These two components are coupled at the ocean time step, every thirty minutes. The atmosphere accumulates net downward shortwave and longwave radiation, sensible and latent heat flux, wind stress, and precipitation. The ocean sends back the atmosphere SST and sea ice fraction. The land model is run as part of the atmospheric model, and so is not seen directly by the coupler. With coupling every 30 minutes, fluxes are exchanged 18 times in the 9 hour forecast generated for the background. This is compared with the 6 hour forecast for the SPEEDYNEMO that used a 6 hour period of the coupling. The strength of the resulting cross domain covariance should therefore be stronger and more useful for SCDA with the CFSv2.

3.2 Experiment Setup

The observation set to be used consists of the in-situ portion of the observations used operationally at NCEP for the CDAS. For the atmosphere this is a subset of the PREPBUFR data as described in table 3.1. Observations such as nexrad wind reports (VADWND) and wind profilers and acoustic sounders (PROFLR) are not utilized due to their high density and concentration over land. Ocean observations consist of the in-situ temperature and salinity profiles from the moored buoys (MRBs), expendable bathythermographs (XBTs), and Argo floats (ARGO). An example of the observation density over a given day is shown in fig 3.1.

The CFSv2 uses a relaxation to OISST in operations for a more accurate analysis. No relaxation is used in these experiments so as not to dampen the ensemble

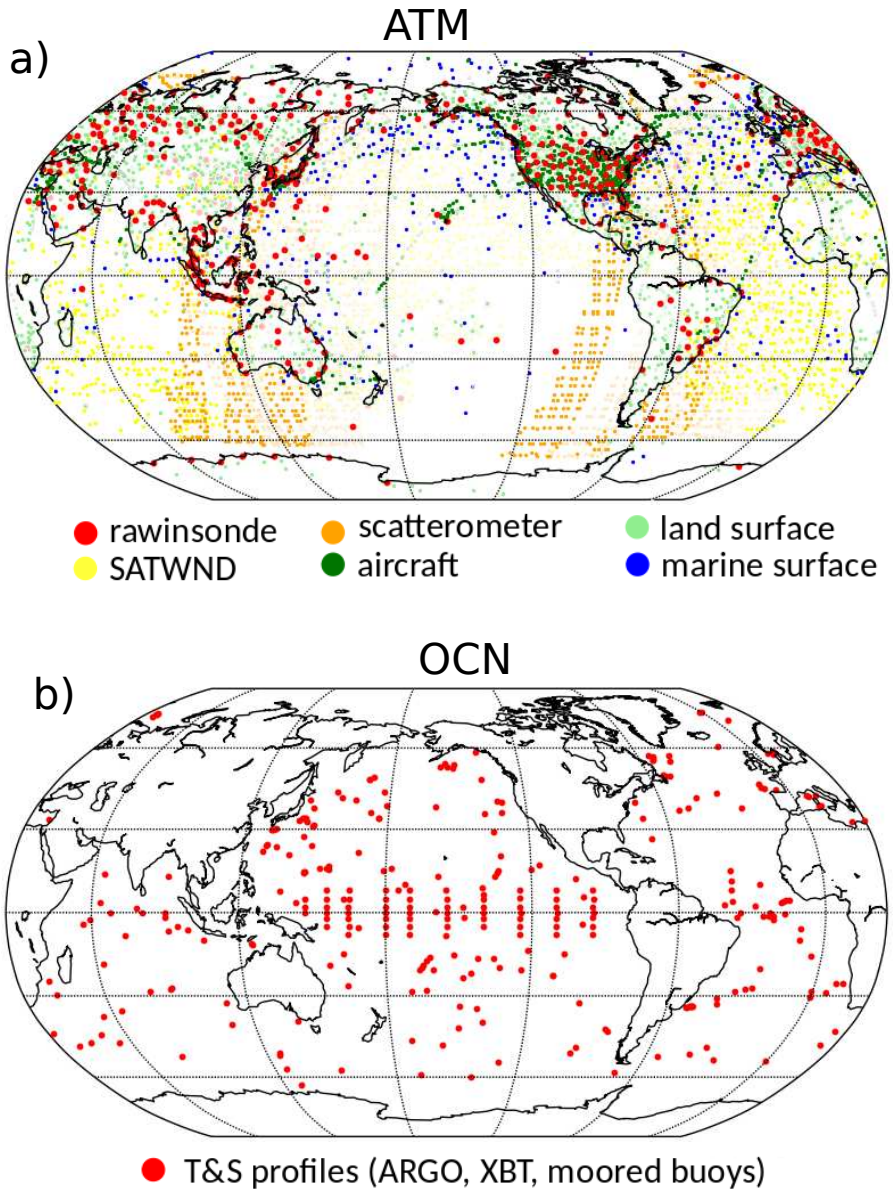

Figure 3.1: The subset of conventional observations used in the experiments. Shown are the atmospheric PREPBUFR obs over a 6 hour period (a) and ocean profiles over a 24 hour period (b).

ob type	description	count
	atmosphere observations	**288,938**
ADPUPA	upper-air reports, mainly rawinsondes	74,322
AIRCAR	MDCRS ACARS aircraft reports	53,953
AIRCFT	AIREP, PIREP, AMDAR, TAMDAR aircraft reports	49,076
SATWND	satellite derived motion vectors	20,496
ADPSFC	surface land reports (SYNOPTIC, METAR)	31,606
SFCSHP	marine surface reports (SHIP, BUOY)	26,569
QKSWND	scatterometer wind data from quickscat	11,906
SPSMIR	scatterometer wind data from SSM/I	21,010
	ocean observations	**14,949**
OCN T	ocean temperature	9,253
OCN S	ocean salinity	5,696

Table 3.1: Operational in-situ atmospheric observations used from PREPBUFR along with counts from a typical 24 hour period (Jun 1, 2005). SFCSHP are the only observations used by the ocean during strongly coupled DA. Each state variable and vertical level counts as a separate observations.

spread and weaken the cross domain correlations. Also, without SST relaxation, larger biases in the ocean are expected, which for the purposes of these experiments is good as it gives the strongly coupled DA another possible area to improve compared with WCDA.

To generate the initial ensemble the CFSR analysis from January 01, 2005 is used and run freely for each of the 50 ensemble members until sufficient spread in both the atmosphere and ocean develops. Tiny perturbations in the atmosphere quickly give rise to large spread after several weeks, and this difference in forcing causes spread in the subsurface ocean after several months (fig 3.2).

These members are run for a total of four months until May 01, 2005. Weakly coupled DA is then performed for one month until June 01, 2005. This is necessary so that the atmospheric ensemble is collapsed sufficiently so that large analysis increments from the atmospheric observations do not overly disturb the ocean during

the spin-up phase of the strongly coupled DA. The final output of the ensemble members at June 01, 2005 are used as the initial conditions for the subsequent experiments.

Other methods of generating an initial ensemble are possible, such as using a random sampling of CFSR analyses from the same month over the historical reanalysis period, which would help result in a larger and more realistic initial spread at depth. However, by initializing with random samples of CFSR analyses, it was found that the ocean had difficulty spinning up without generating large noise and ultimately diverging if the ensemble differed too greatly from the truth, and so a single CFSR analysis was used and run for sufficient time to generate ample spread and errors.

The following parameters are used by the LETKF for all experiments:

ensemble size	50
inflation	RTPS at 95%
ATM localization	1000km Hz, 0.4lnP vertical
OCN localization	720-200km Hz, 0.4lnP into ATM

Table 3.2: LETKF parameters used for CFSv2-LETKF experiments with real observations.

3.3 Control run OSE with WCDA

WCDA ensemble spread/rmsd/bias

The CFS-LETKF is run with weakly coupled DA for one month from May 1, 2005 to June 1, 2005. The ensemble spread of the ocean temperature and salinity

at various depths is shown in fig 3.2. Compared to the SPEEDYNEMO weakly coupled runs, the spread in the CFS is much larger and exhibits realistic patterns of increased spread along the equatorial thermocline and western boundary currents. This increased spread is likely due in large part to increase ocean model resolution from 2°to 1/2 °.

The WCDA configuration run is continued for several more months, and the resulting background minus observation (B-O) bias and RMSD averaged over the summer months are shown in fig 3.3. There are small upper ocean errors throughout most of the open ocean (less than 0.5°C), with larger errors (> 2°C) in the northern hemisphere mid latitudes and along the coastlines. These same areas of larger RMSD also have a large warm bias in the model, as seen from the SFCSHP and ocean B-O bias.

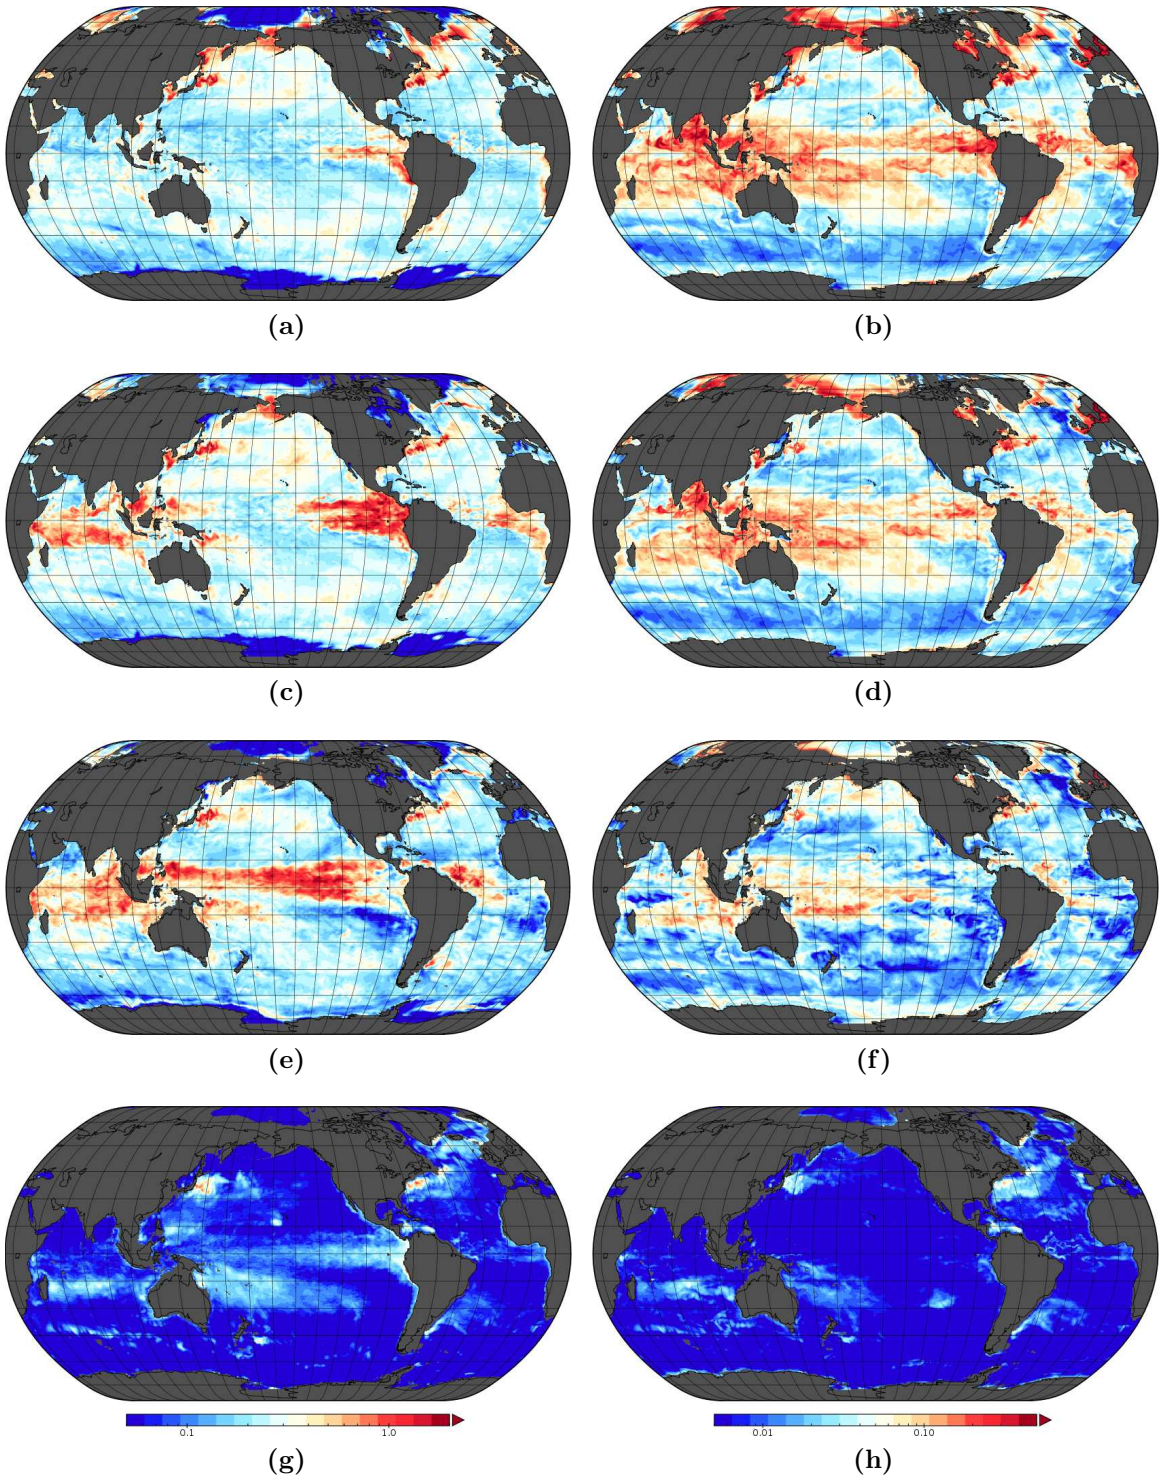

Figure 3.2: Ocean ensemble spread after the WCDA spinup for one month, ending Jun 1, 2005. Temperature in C (left) salinity in PSU (right) at 5m (a,b) 50m (c,d), 100m (e,f) and 500m (g,h) depths on a semi-logarithmic scale.

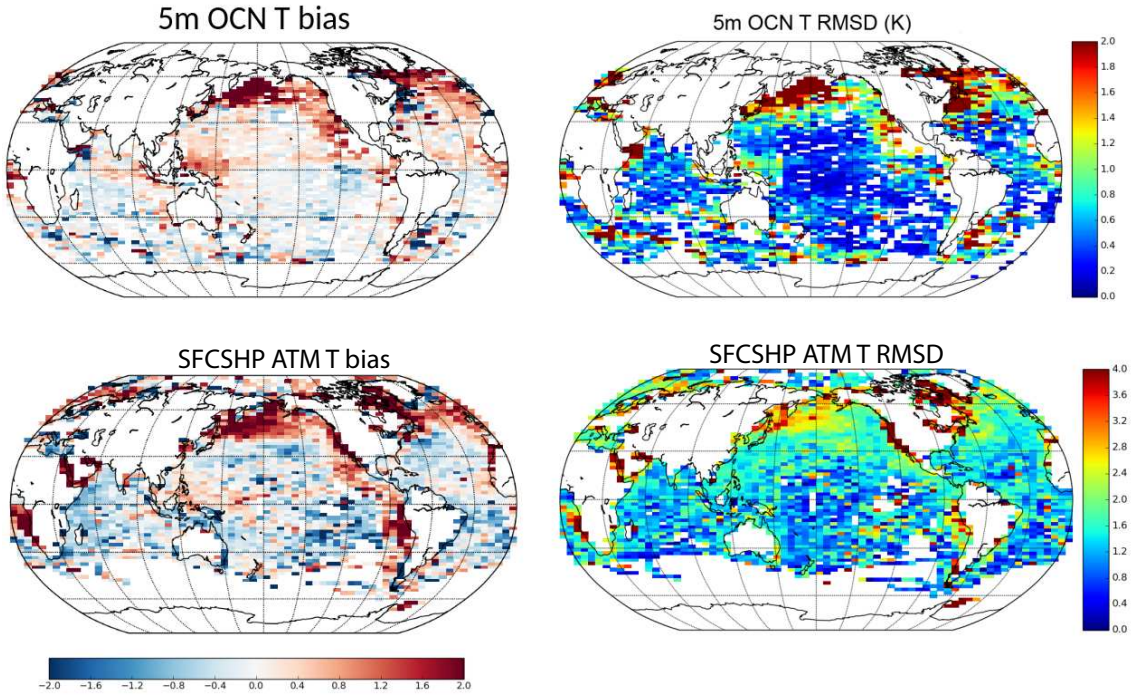

Figure 3.3: Background minus observation statistics of bias (left) and RMSD (right) averaged over JJA of 2005 with weakly coupled DA. Shown are the statistics for 5m ocean temperature observations (top) and SFCHP atmospheric temperature (bottom) in C.

WCDA ensemble cross-domain correlations

The background error cross-domain correlations from a single date in the NH summer (June 1, 2005) and a single date in the NH winter (December 1, 2005) are examined to determine how much information within the ensemble is available for use by SCDA. These data shown are from a single instantaneous time, and not a climatological average. If there is no correlation shown, SCDA would not be expected to provide any benefit, and may even harm the analysis if any correlations present are spurious. The strength of spurious correlations would be expected to change based on the ensemble size used.

The cross correlations in fig 3.4 and fig 3.5 are generated from the given variables by calculating the correlation between each atmospheric model level and the ocean surface level (top) and between each ocean level and the bottom atmosphere level (bottom) using the ensemble member perturbations. The results are then zonally averaged.

For the ocean and atmosphere temperature cross correlations, it can be seen in Fig 3.4 that temperature near the interface is highly correlated as would be expected. The highest correlation extends down into the ocean to the base of the mixed layer. The highest correlation extends up into the atmosphere to roughly the 0.9 sigma level. The correlations are stronger in the summer time (NH in fig 3.4a, SH in fig 3.4b). The plotted values have been normalized due to the differences in maximum correlation in June and December (0.64 vs 0.36). The June values are likely artificially large due to insufficient spinp-up time: the ensemble spread in the

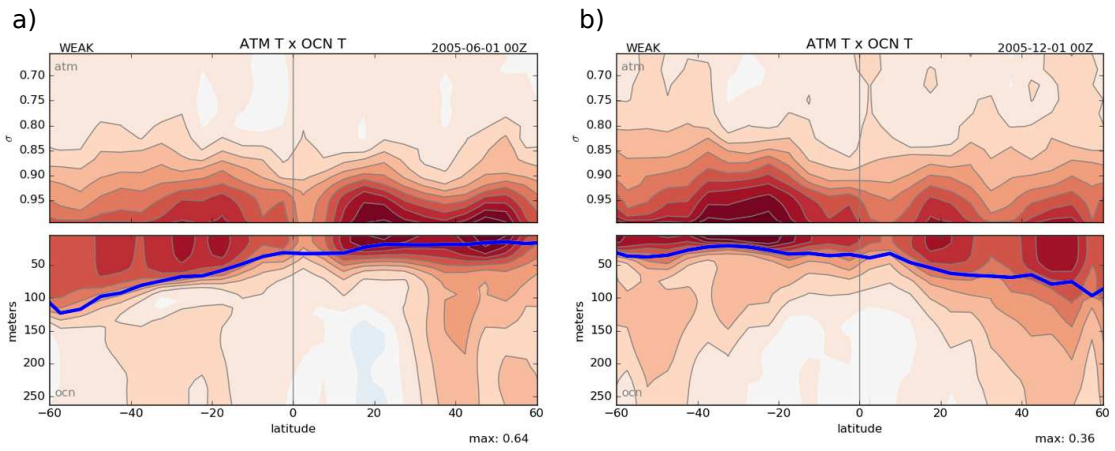

Figure 3.4: Zonally averaged instantaneous cross-domain background error correlations represented by the ensemble between ocean temperature at each level and surface atmospheric temperature (bottom) and atmospheric temperature at each level and surface ocean temperature (top). Dark blue line indicates location of the ocean mixed layer depth. NH summer (left) and NH winter (right) are shown. Plotted values have been normalized, maximum correlations are 0.64 for NH summer, and 0.36 for NH winter.

ocean is large in June, resulting in larger corresponding anomalies in the atmosphere. The 0.36 maximum correlation value in the winter is closer to the results found from Feng et al. (2018) with the ECMWF CERA for SST-T2m 3hr ensemble correlations.

Fig 3.5 shows the similar cross correlations against ocean temperature for atmospheric humidity (fig 3.5a) and wind speed (fig 3.5b), zonal wind speed and meridional ocean current(fig 3.5c), and meridional wind speed and zonal ocean current(fig 3.5d)in December. The overall pattern of OCN_T X ATM_Q is similar to that of OCN_T X ATM_T, although weaker. Atmospheric wind is shown to be negatively correlated with surface ocean temperature as a likely result of increase heat loss from the higher winds. The cross correlations with atmospheric winds are much weaker than those of temperature and humidity, and so are not likely to be represented will with a limited ensemble size. Winds are correlated with ocean cur-

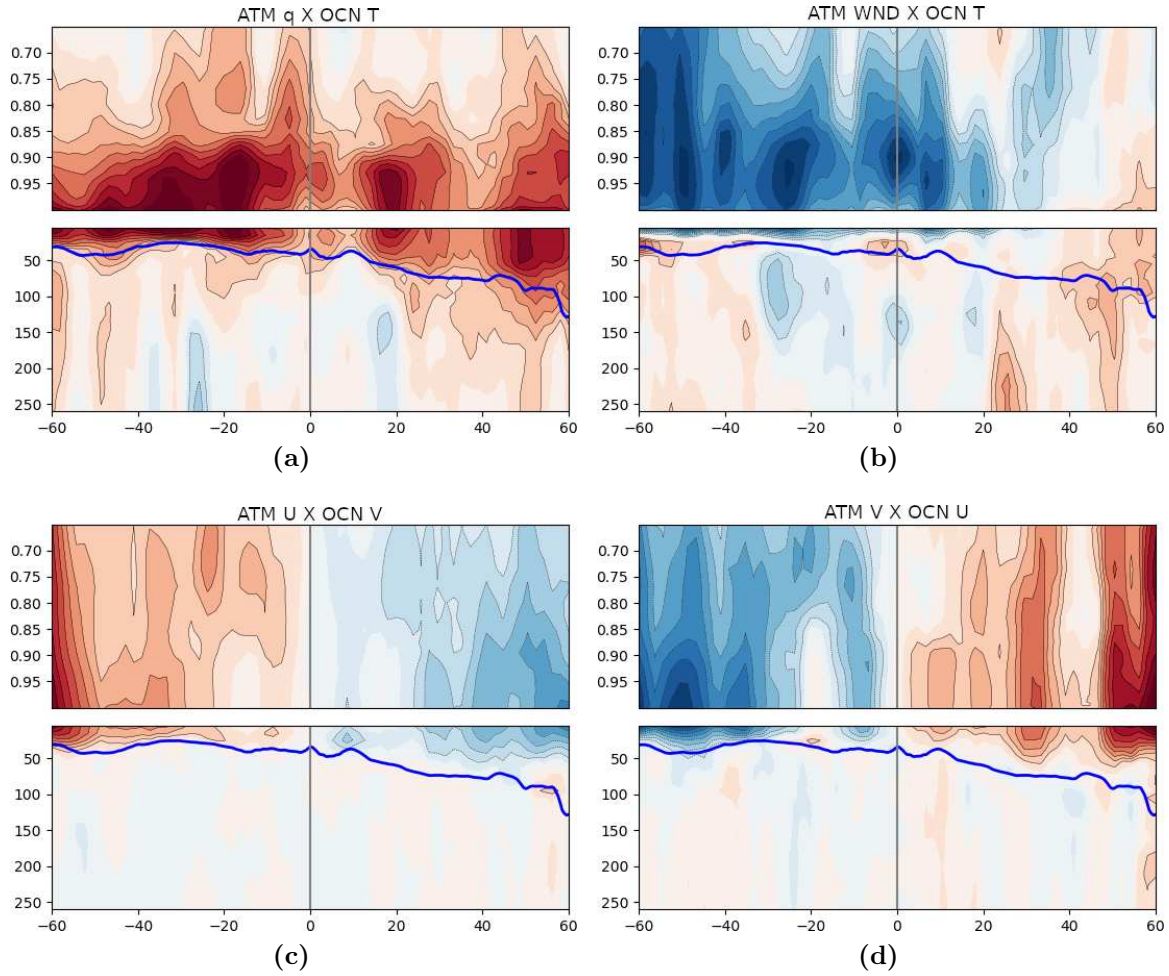

Figure 3.5: Zonally averaged instantaneous cross-domain background error correlations represented by the ensemble between the three-dimensional ocean value and the surface atmospheric value (bottom), and the three-dimensional atmospheric field and the surface ocean field (top). Variables shown are between ocean temperature and atmospheric humidity (a), ocean temperature and wind speed (b), ocean meridional currents and atmospheric zonal wind (c), and ocean zonal currents and atmospheric meridional wind (d). Dark blue line indicates location of the ocean mixed layer. Plotted values have been normalized, maximum correlations are given by table 3.3. The ensemble at a single snapshot at 2005-12-01 00Z is used.

		ATM					
		T	**q**	**u**	**v**	**wnd**	**vort**
	T	0.34	0.15	0.07	0.02	0.11	0.06
OCN	**S**	0.04	0.05	0.04	0.03	0.03	0.05
	u	0.08	0.06	0.17	0.14	0.11	0.05
	v	0.06	0.06	0.21	0.14	0.13	0.03

Table 3.3: The absolute values of the maximum zonally averaged cross correlation between the given atmospheric and oceanic variables between 60N and 60S on 2005-12-01 00Z.

rents that are at a 90° angle, due to the Ekman transport. Table 3.3 shows a more complete listing of the absolute value of the maximum cross-domain correlation for various variables, including derived atmospheric vorticity and divergence. Atmospheric temperature and humidity have the strongest correlations with the surface ocean.

Observation bias

When the SFCHP observation bias with respect to the model forecast is calculated over the summer months, and separated by observation hour (0, 6, 12, 18Z) it can be seen that the bias over large parts of the ocean are diurnally dependent (fig 3.6). The model is systemically colder than observations during the daytime, and warmer than observations at night. This type of bias is seen in the MERRA2 reanalysis also (James Carton, personal comm.), and is theorized to be a combination of model and observation biases. The SFCSHP temperature observations are obtained from ships and so are placed over a warm deck, and variables such as sensor position and deck color are not well accounted for. As a result the sen-

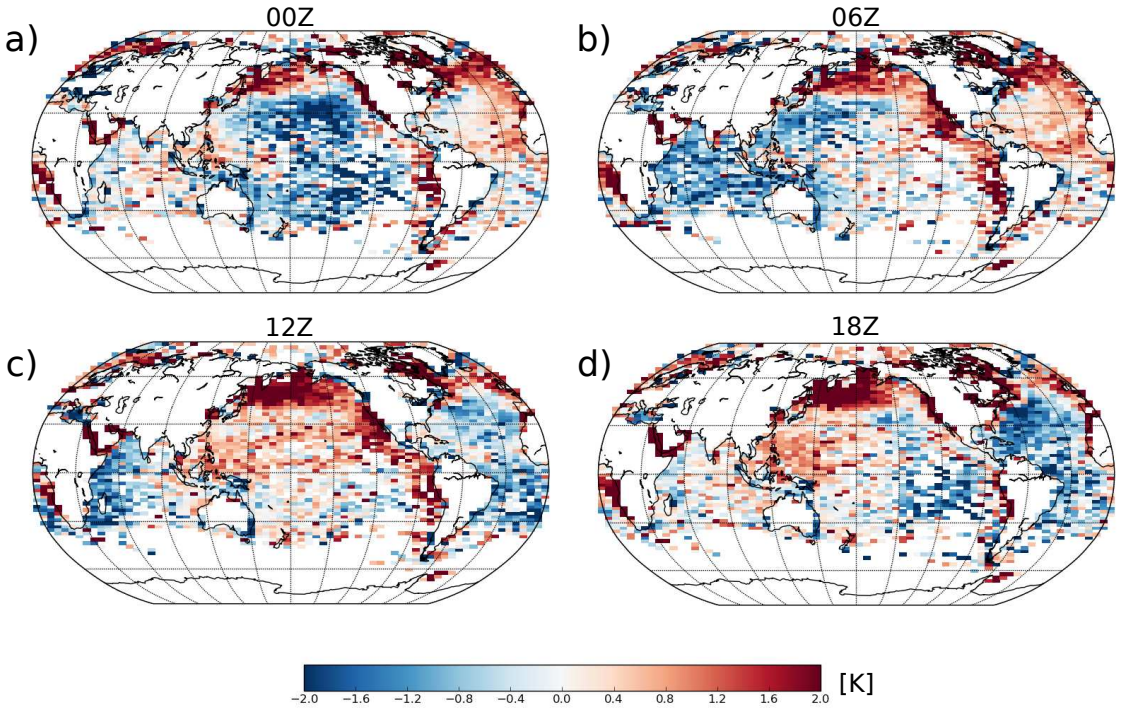

Figure 3.6: Background minus observation (B-O) diurnal bias shown by comparing CFS weak run with SFCSHP T observations averaged over JJA of 2005.

sors may read observations warmer than they would otherwise during the daytime. Also, the CFSv2 ocean model uses a top layer thickness of 10m, and so has a poor representation of the diurnal cycle.

3.4 SCDA OSE with SFCSHP obs

Single Observation Test

After 1 month of spinup with WCDA configuration, there are several areas, especially along the coastline, that exhibit large SST errors. The Yellow sea is chosen as a quick test of the SCDA configuration, the location of the observations and the background biases are shown in fig 3.7. By independently assimilating the

ocean observation or atmospheric observation with SCDA, a beneficial cross-domain analysis increment is generated (fig 3.8).

It is interesting to note that, at least in this one test case, the ocean observation was able to improve the atmosphere more than the atmosphere observation was able to. This can be attributed to the difference in the background and observation error ratio of the two domains. The ocean background spread (not shown) is larger than that of the atmosphere. The synthetic observations for both the atmosphere and ocean have a prescribed observation error of 1°C. Therefore, the ocean observation carries more information about the coupled system compared with the background than the atmospheric observation does. With SCDA, it is possible that there are inter-domain observations that are have more impact than intra-domain observations, depending on the background ensemble spread, cross domain correlations, and the observation error variance.

Experiment Setup

For the full SCDA experiment, only the SFCSHP temperature and humidity observations are assimilated into the ocean. These are the most abundant insitu atmospheric temperature observations over the ocean, and given that the ocean and atmosphere temperature cross correlations were the strongest in the WCDA run, other atmospheric observation types will not be used. By removing observations of model states shown to have small cross domain correlations, a form of variable localization (Kang et al., 2011) is essentially being used. Since the GODAS ocean

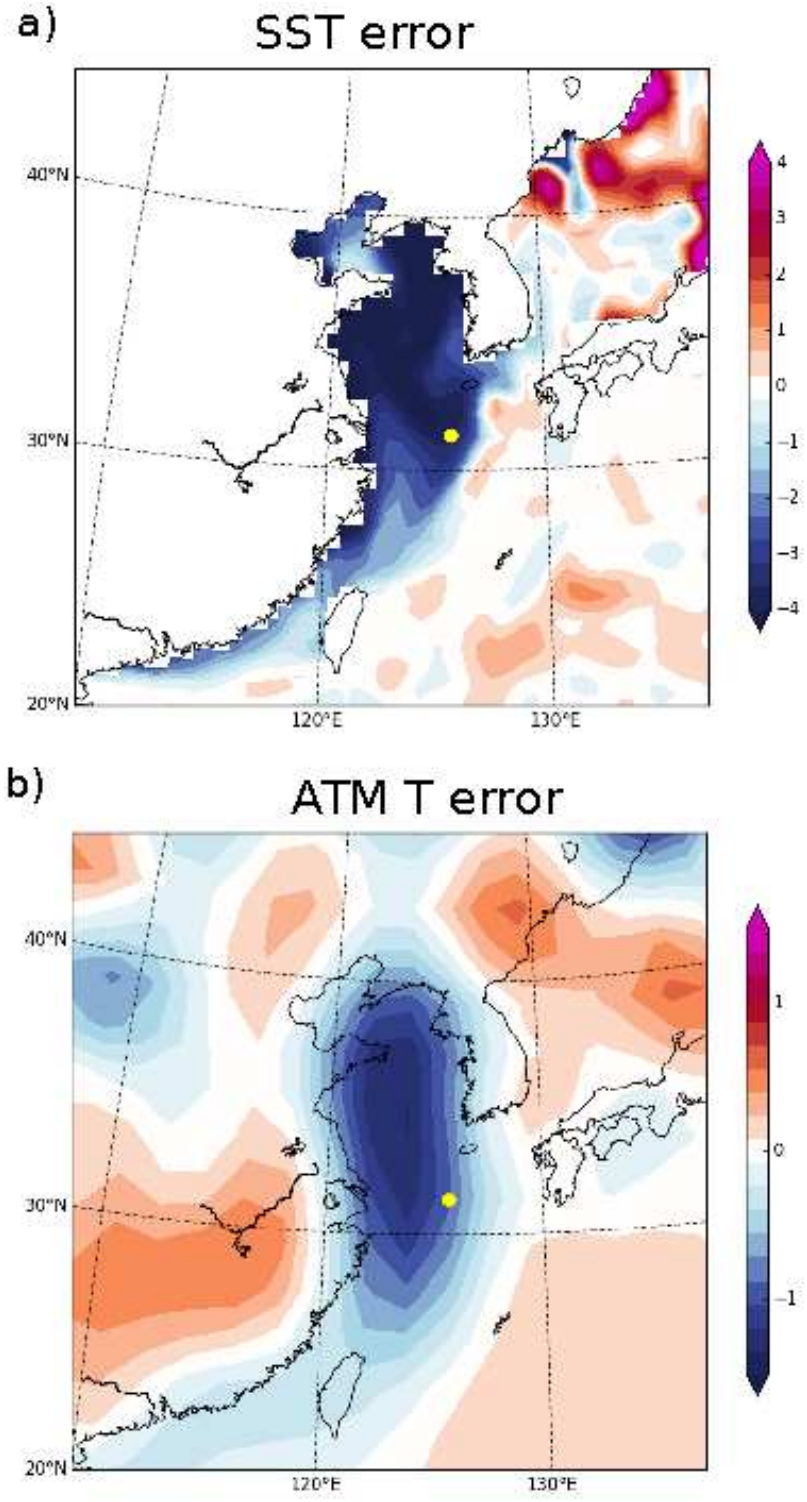

Figure 3.7: Location of the single test observations (yellow dots) and background bias for the ocean (top) and atmosphere (bottom).

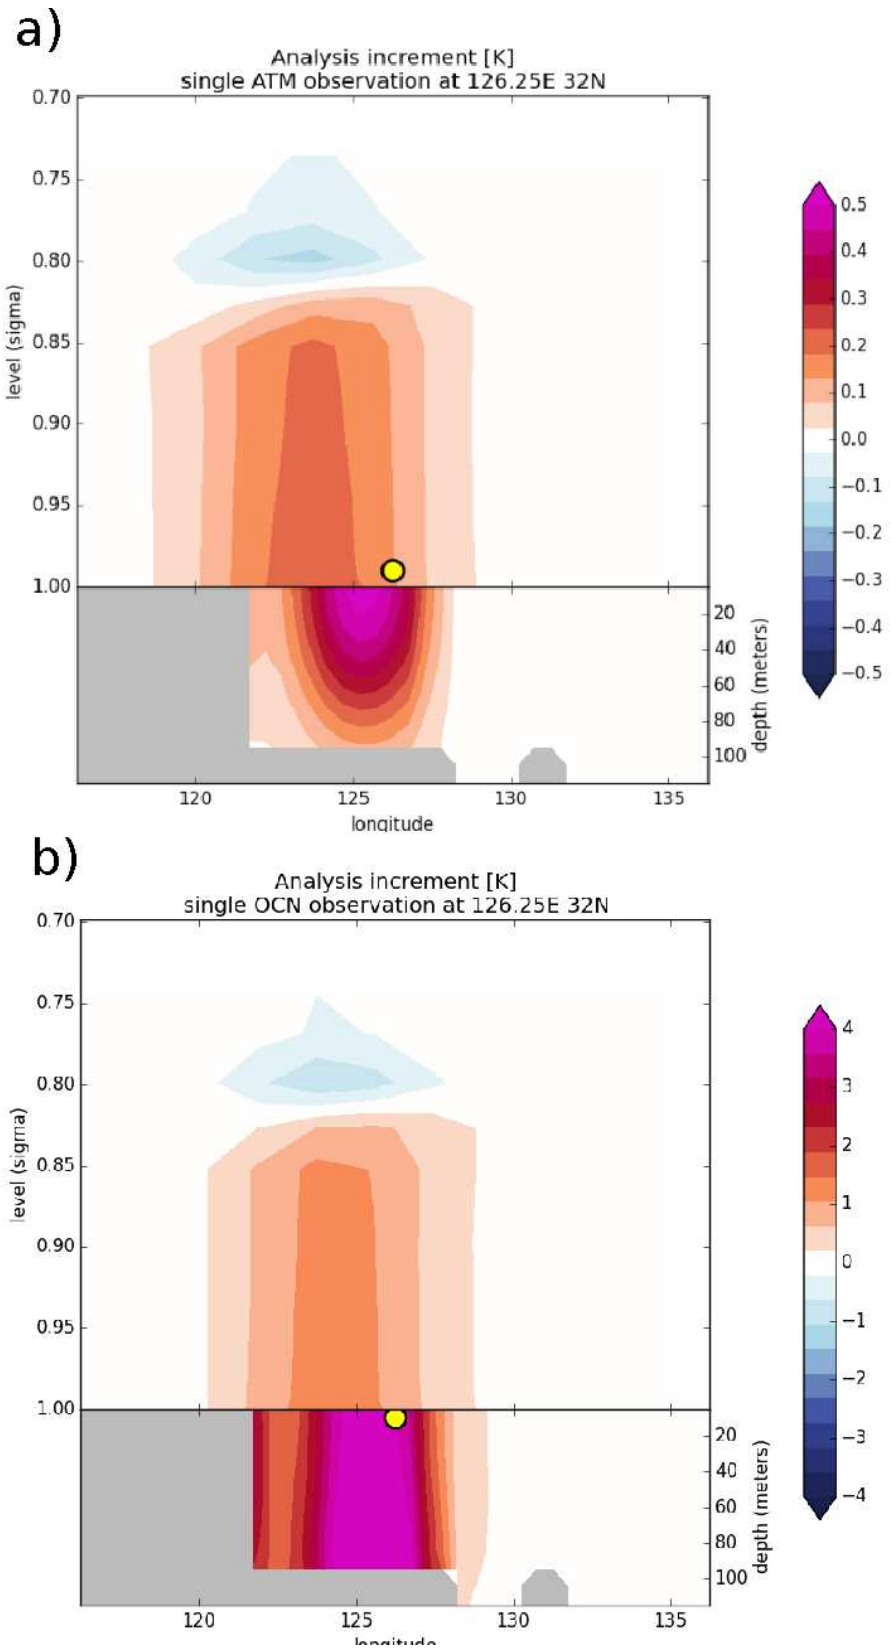

Figure 3.8: Vertical cross section of the SCDA analysis increment resulting from assimilating a single atmospheric temperature observation (top) and a single ocean temperature observation(bottom).

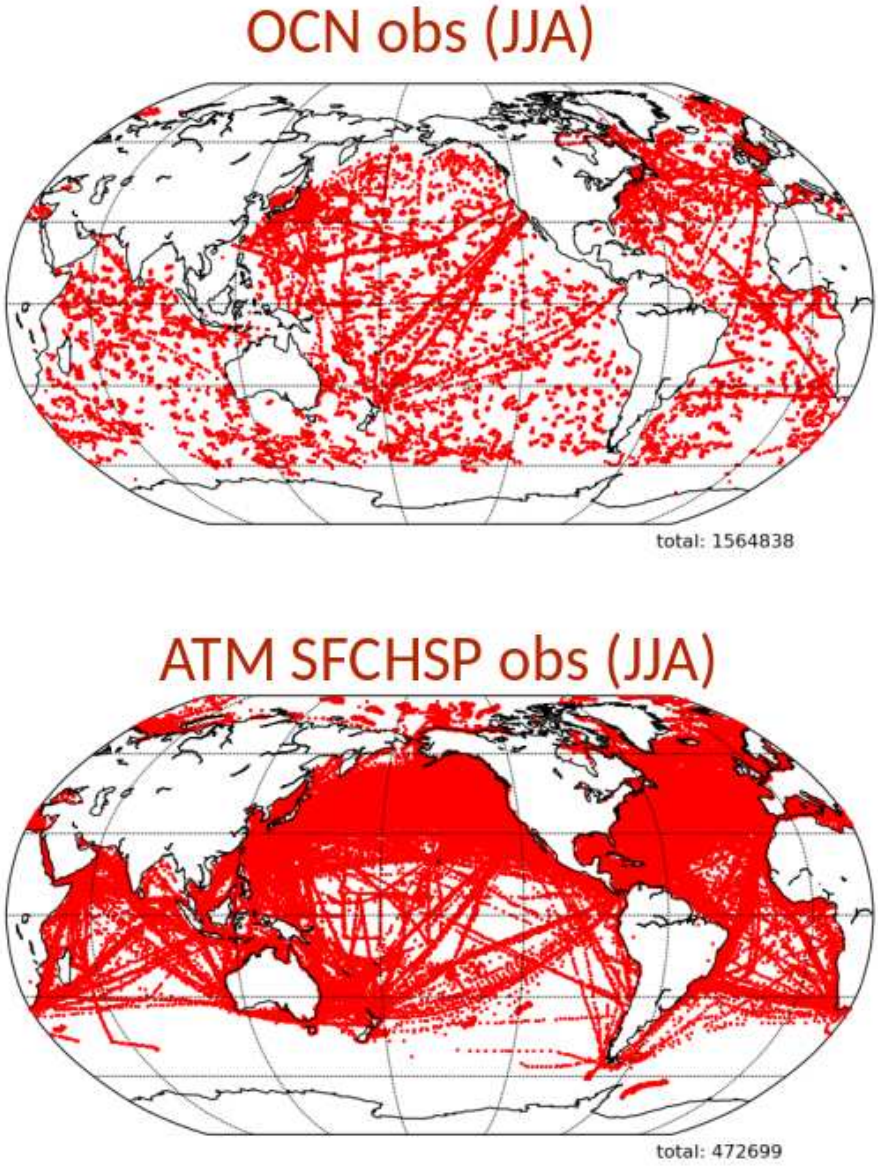

Figure 3.9: The subset of observations that are used across domains with SCDA. Shown are the top level ocean observation (top) and SFCSHP T and q (bottom) over the months of June, July, August in 2005.

observation dataset has already undergone daily averaging, the assimilation of the upper ocean observations into the atmosphere would be more difficult. So, as with the SPEEDYNEMO experiments, only a one-way atmosphere into the ocean SCDA experiment will be attempted with real observations.

Both the ocean and the atmosphere LETKF are running in 6 hour cycles. However, ocean observations are only available every 24 hours, whereas atmospheric observations are available for each 6 hour cycle. Given that the ocean DA is essentially only occurring every 24 hours, how often should atmospheric observations be assimilated into the ocean, 6 hours or 24 hours? Synchronous SCDA (fig 3.10a), is called so because the atmosphere observations are assimilated into the ocean at the same time as they are into the atmosphere. Initial experiments with the CFS failed with synchronous SCDA. This is likely due to the large diurnal bias discovered (fig 3.6) resulting in repeated shocks to the ocean by pulling it in opposite directions every 12 hours.

For SCDA to be successful, the model and observation biases could be addressed (which is important, but outside the scope of this work). Or, a longer window can be used for the ocean DA. Asynchronous SCDA (fig 3.10b), assimilated the atmospheric observations into the atmosphere every 6 hours, but those ensemble observation departures are saved up and then used in the ocean data assimilation every 24 hours, the same time that ocean observations are present. This is a similar concept to the lagged average coupled covariance of Lu et al. (2015b), except without any explicit averaging of the observations, leveraging the ability for the LETKF to operate as a 4D-EnKF.

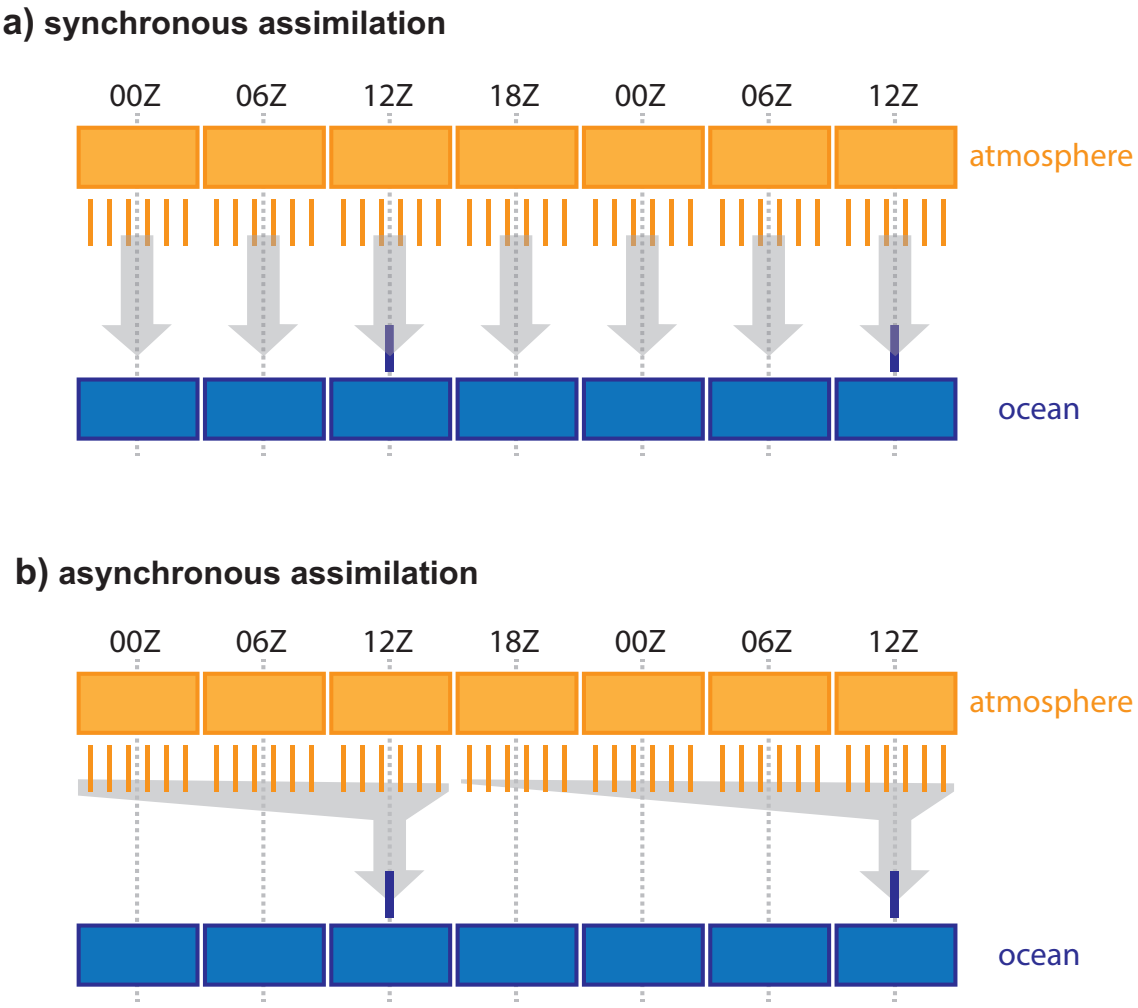

Figure 3.10: Synchronous SCDA (top) whereby atmospheric observations are assimilated into the ocean at the same time they are assimilated into the atmosphere (every 6 hours). **Asynchronous** SCDA (bottom) where the atmospheric observation are assimilated into the ocean at the normal ocean cycle time (every 24 hours).

It is not obvious that an asynchronous SCDA method would work, however it does if the background and analysis ensemble members stay "matched up" after the analysis step. The local ensemble Kalman filter, LEKF, (Ott et al., 2004), on which the LETKF is based, has the beneficial property that the distance between the background and analysis ensemble member is minimized (see Ott et al. (2004) Appendix A). This has been shown to be useful by Kretschmer et al. (2015), who has developed the climatologically augmented LETKF (CaLETKF). The CaLETKF splits the ensemble members into dynamic and static ensemble members, the first k_d members are always the dynamic members. Separating the ensemble members this way works because there is a natural correspondence between the perturbation direction of a given ensemble member's background and analysis with the LETKF.

Results

Comparing the resulting SCDA run to WCDA over the summer months in fig 3.11, the errors in the 6 hour background for atmospheric temperature are greatly reduced in the northern hemisphere where there is the greatest density of observations. SFCSHP temperature RMSD reduction is 13% by the end of August in the northern hemisphere, with small improvements of 3.8% and 2% seen in the southern hemisphere and the tropics, respectively.

Shown spatially (fig 3.12), both the ocean and atmosphere are shown to have improvements in the northern hemisphere in the same regions shown to have the large biases. There is some degradation in the ocean near the coasts. For better

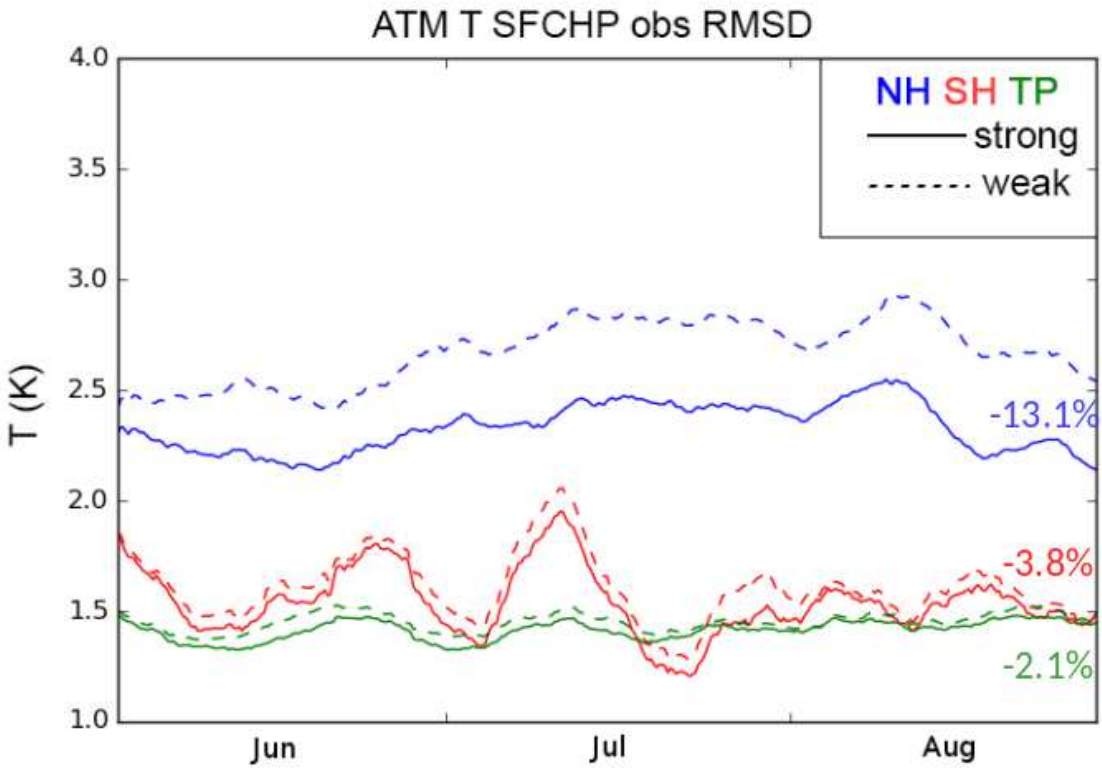

Figure 3.11: The RMSD for atmospheric SFCSHP surface ship temperature observations comparing the 6 hour forecast to the observed values. Shown are results with WCDA (dashed) and SCDA (solid) averaged over the northern hemisphere (NH), tropics (TP), and southern hemisphere (SH)

performance, the SFCSHP observations close to land should be excluded.

The SCDA-WCDA O-F RMSD at various depths is shown in fig 3.13. It can be seen that the RMSD of the tropical ocean is reduced between the surface down to 100m. In the northern hemisphere however, only the upper ocean is improved, and is in fact degraded below 15m. This is likely due to spurious correlations between the surface and the deeper ocean due to insufficient ensemble size.

The ocean in the extra-tropics during the summer time have a very shallow mixed layer (fig 3.14). As was shown with the WCDA correlations, the strongest correlations between the ocean and the atmosphere exists within the mixed layer.

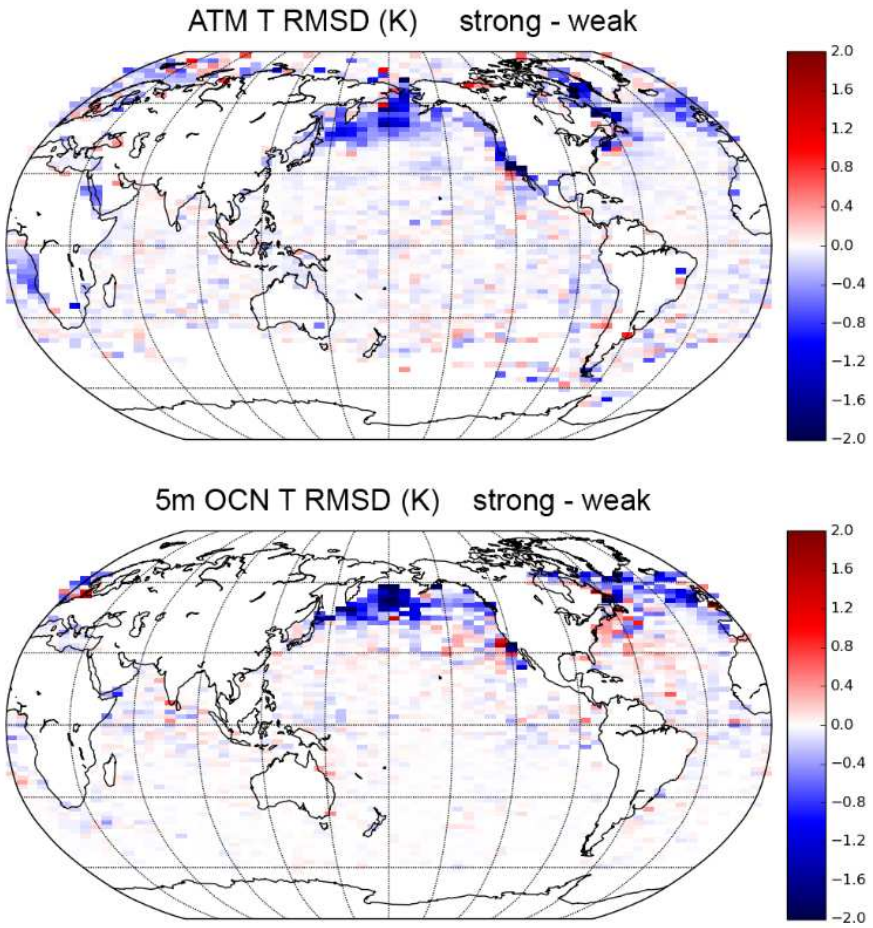

Figure 3.12: STRONG - WEAK change in observation minus forecast (O-F) RMSD for atmospheric temperature at the lowest model level (top) and upper ocean temperature (bottom).

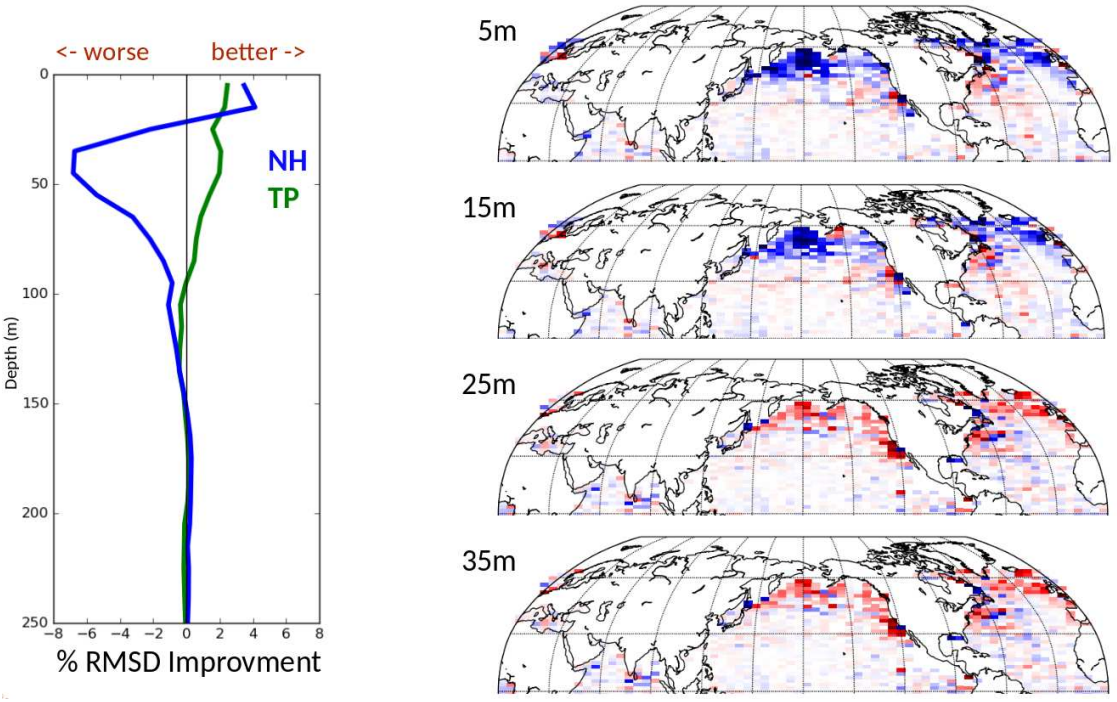

Figure 3.13: STRONG - WEAK change in observation minus forecast (O-F) RMSD for ocean temperature. Averaged over the tropics (TP) and Northern Hemisphere (NH) at various depths (left) and shown spatially (right). For the spatial plot blue is an RMSD improvement, red is a degradation.

Therefore, to further improve performance, the vertical localization needs to be applied to limit the impact of the atmospheric observations to the mixed layer. The tropics exhibit a slightly deeper mixed layer, which explains why the tropical ocean has a consistent reduction of RMSD to a deeper depth.

3.5 Summary

The strongly coupled data assimilation experiments performed with the CFSv2 and real observations has demonstrated that SCDA in an operational context might be possible if done carefully, however model and observational biases likely need

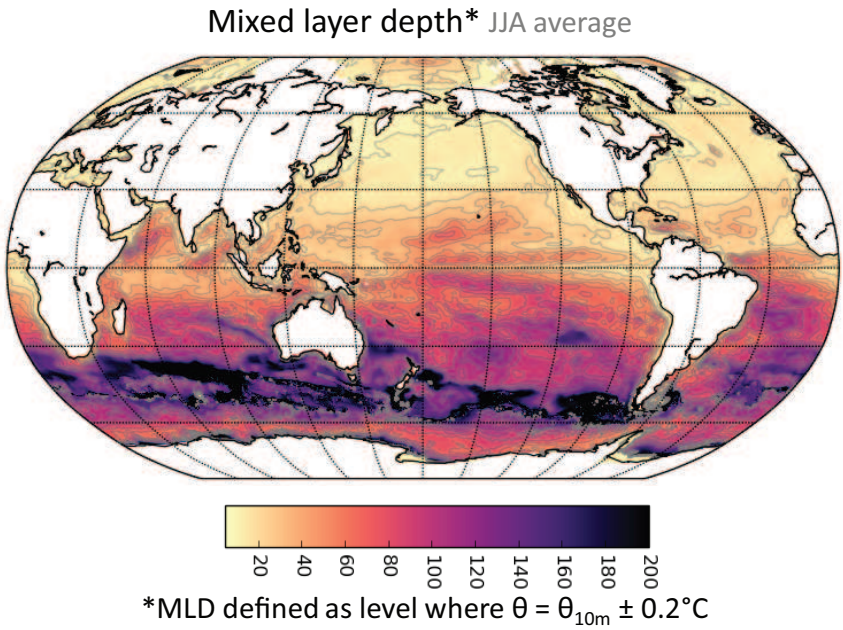

Figure 3.14: The depth of the mixed layer averaged over the JJA months for 2005 in the WCDA experiment as calculated by a change in surface temperature of 0.2C.

to be addressed first before any substantial improvements can be had with real observations.

A coupled run with weakly coupled DA is beneficial for identifying the state variables and observation types that have a chance to provide beneficial impacts to the SCDA analysis. Using vertical and variable localization (Kang et al., 2011) as a form of correlation cutoff method (Yoshida and Kalnay, 2018) is vital when using a limited ensemble size. These cross-domain state variables with small correlations need to be removed in the data assimilation step in order to avoid the detrimental impact of spurious correlations in the ensemble.

The weakly coupled CFSv2-LETKF runs exhibited a strong bias in parts of the northern hemisphere, when examining the observation minus forecast statistics.

These areas were the areas most easily helped by using SCDA to assimilate the atmospheric surface temperature and humidity into the ocean. Areas without a large ocean bias did not improve as substantially.

The CFSv2-LETKF was difficult to get working well in a SCDA setting. This was due to the very strong diurnal signal in the atmospheric observation minus forecast biases. For a well tuned ocean-atmosphere SCDA to work correctly, the vertical resolution of the ocean model needs to be increased to allow for better representation of the surface diurnal cycle, and bias correction of the atmospheric observations needs to be performed. As a work around, asynchronous SCDA was utilized. With asynchronous DA, the 6-hourly atmospheric observation ensemble departures were collected for several cycles and then assimilated into the ocean at the 24-hour interval the ocean observations were present. This method smoothed out the diurnal bias of the atmospheric observations allowing assimilation into the ocean to work. This has a similar effect as the lagged average coupled covariance (LACC) method of Lu et al. (2015b) increasing the strength of atmosphere-ocean correlations. If diurnal bias issues are resolved, synchronous SCDA may be beneficial.

Chapter 4: Towards Operational SCDA: Hybrid-GODAS

A new ocean data assimilation system is currently being developed at NCEP called the Hybrid-GODAS. This system will serve as a replacement for the global ocean data assimilation system (GODAS) used at NCEP for realtime ocean monitoring (Behringer and Xue, 2004; Behringer, 2007). GODAS is run in two configurations, one is as part of the weakly coupled data assimilation used for the CFSv2, the other configuration is a stand-alone ocean monitoring system driven by offline atmospheric fluxes. Hybrid-GODAS will serve as an upgrade to the stand-alone GODAS (described in this chapter), but a similar upgrade will be made to the coupled system in the near future.

While the work presented does not directly involve strongly coupled ocean-atmosphere data assimilation, the upgrade of GODAS presented will outline the ways in which considerations are made for future use as a strongly coupled system. The future coupled model should alleviate some of the problems noted with the CFSv2 in chapter 3 that are preventing the system from easily being used with SCDA. Notably, the ocean vertical resolution will be much higher at the surface, allowing for a better diurnal cycle representation in the ocean. Also, Hybrid-GODAS is based on the LETKF, meaning the SCDA methods presented in chapters 2 and

3 can easily be implemented in the future with it.

The current GODAS is based on the original 3DVAR algorithm for the ocean developed by Derber and Rosati (1989) using a state-space variational solver. The background error covariance model is simple compared with more modern ocean DA systems used operationally. All background error covariances are univariate: temperature and salinity analysis are essentially performed independently. This can be compared with other systems, such as NEMOVAR, which have balance operators to allow for a single temperature increment to update all other state variables (salinity, zonal, and meridional current).

GODAS has served well over the years, and is a popularly downloaded dataset. However, the last major update was in 2003, and the system is beginning to show its age. The system often does not perform as well as other centers (Xue et al., 2017). Several design deficiencies compared with other operational centers are noted: GODAS only assimilates insitu temperature, and does not take advantage of a wide range of other observation platforms (e.g. salinity, satellite altimetry, satellite SST, drifter positions, and ocean currents from ADCP). The 3DVAR algorithm that is implemented is computationally inefficient for the ocean, due to its state space formulation and the use of diffusion operators for the horizontal background error correlation model.

The GODAS system is being replaced by the Hybrid-GODAS, described in this chapter, with a target operational deployment at NCEP by the end of 2019. The new system is a complete upgrade (no code spared!) and aims to improve all aspects of the system: observation platforms and their quality control, ocean

and ice model, atmospheric surface forcing, and data assimilation method (fig 4.1). Hybrid-GODAS will become the next real-time ocean monitoring system used by the CPC. It is also expected to form the basis for future work at NCEP with coupled ocean/ice/wave data assimilation, and parts of the advancements made will be available for the coupled data assimilation system using MOM6/FV3 under the Joint Effort for Data Assimilation Integration (JEDI). JEDI is expected to be the replacement data assimilation system used at NCEP, and is being developed by the Joint Center for Satellite Data Assimilation (JCSDA)

All code for Hybrid-GODAS is already publicly available online for use by the research community, even though active system development is still ongoing, (Sluka, 2018a).

4.1 Hybrid-GODAS

A summary of the major changes between GODAS and hybrid-GODAS are given in Table 4.1 and are expanded upon in the text following.

4.1.1 Model

The Modular Ocean Model 6 (MOM6), is the latest generation of ocean model produced by the NOAA Geophysical Fluid Dynamics Laboratory (GFDL), and is a substantial change from the previous MOM models. MOM6 uses a new algorithm, the arbitrary Lagrangian-Eulerian (ALE) algorithm, to allow for any type of vertical grid without having to worry about violating CFL conditions. MOM6 is coupled

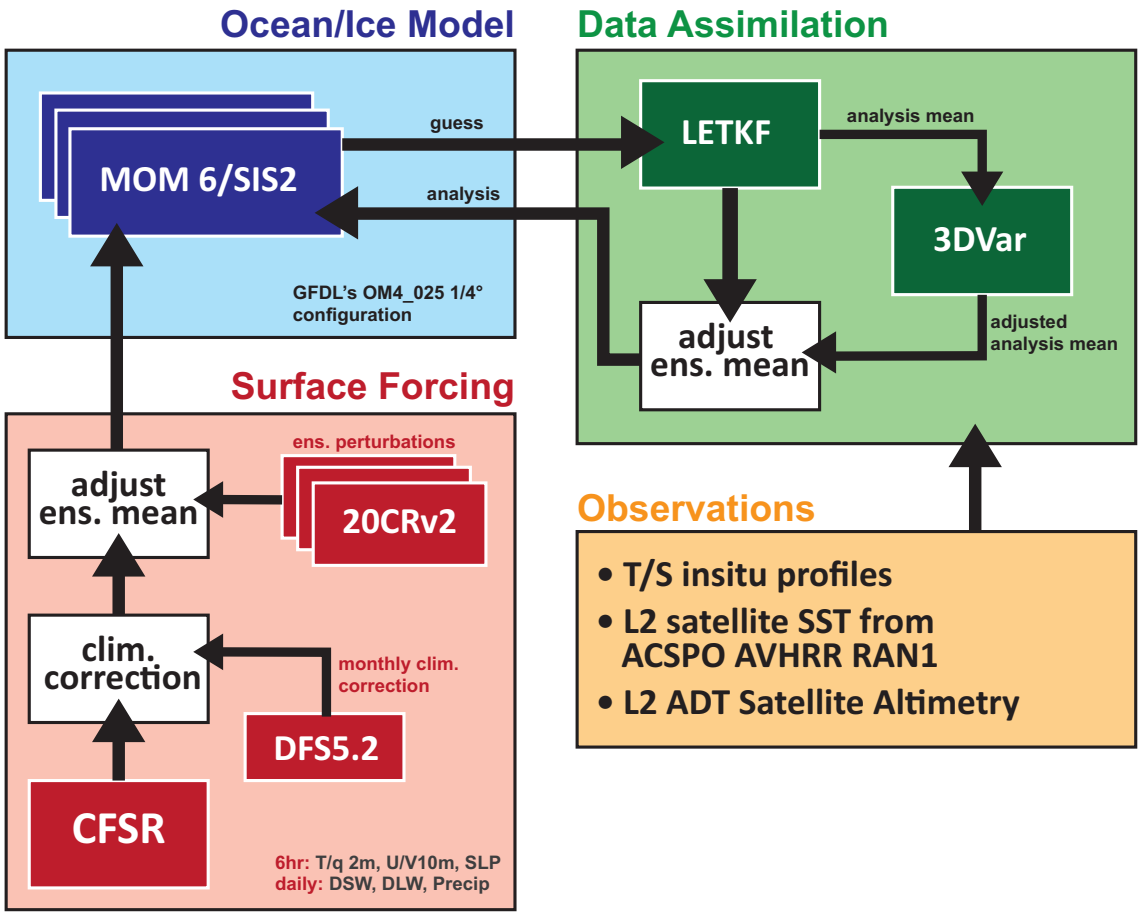

Figure 4.1: Overview schematic of the Hybrid-GODAS ocean data assimilation system, with updates to all components.

	GODAS	Hybrid-GODAS
model		
model	MOM4p1	MOM6
ice model	-none-	SIS2
hz resolution	1 deg	1/4 deg
vt resolution	40 levels, top 10m thick	75 levels, top 2m thick
forcings		
atm. forcing	R2 daily fluxes (wind stress, heat & freshwater flux)	CFSR: 6hr U/V/T/q/MSLP daily DSW, DLW, Precip.
clim. correct.	-none-	DFS52: U/V/DSW/DLW/Precip
ens. purturb.	-none-	20 members from 20CRv2
SST relaxation	OISST, 5 day scale	-none-
SSS relaxation	WOA, 10 day scale	Monthly WOA clim. 166 mm/day
river runoff	Dai and Trenberth annual clim.	Dai and Trenberth monthly clim. + stochastic perturb.
observations		
T profile	NCEP BUFR tank	WOD & NCEP BUFR tank for operations
S profile	synthetic salinity from clim.	observed S profiles
SST	relaxation to Reynolds SST	Night time ACSPO L2 SST
Altimetry	-none-	L2 ADT altimetry (in progress)
data assim.		
method	univariate 3DVar	Hybrid EnKF/3DVar
ensemble size	1	20
bkg. err. var.	vertical gradient of bkg	ensemble spread, vertical gradient of bkg, horizontal T O-F at surface
vertical scales	model level thickness	mixed layer depth at surface transitioning to model level thickness

Table 4.1: Major differences between GODAS and the new Hybrid-GODAS

with the sea ice simulator 2 (SIS2) ice model.

The configuration used here is the OM4 1/4 degree configuration from GFDL, which is the configuration to be used in their next coupled model system. The horizontal grid is 1/4 degree, and the vertical grid consists of 75 levels, with the top levels at 2m thickness down to 10m. The default vertical coordinate system used in OM4 is the hybrid z-isopycnal coordinate. However, due to issues with spurious spread that have been found (Steve Penny, personal comm.) current work with MOM6 here is being done with z* coordinates. Hybrid vertical coordinates are shown to have better model performance in the deep ocean and in areas of high vertical stratification, and so we do wish to switch from Z* to the hybrid- vertical coordinates as soon as the spurious spread issue is resolved.

The increased horizontal resolution allows the model to be eddy permitting in the mid latitudes, and therefore should represent western boundary currents better than the ocean model in the current CFSv2. Also, the increase in vertical resolution at the surface should allow for a better diurnal cycle representation.

4.1.2 Surface forcing

The stand-alone GODAS is forced with a set of atmospheric fluxes from the NCEP-DOE AMIP-II Reanalysis (R2) (Kanamitsu et al., 2002). R2 provides the momentum (zonal and meridional wind stress), heat (sensible, latent, radiative), and freshwater (precipitation minus evaporation) fluxes from 1979 to the present on a T62 resolution grid. Since the ocean model is driven by a set of fluxes calculated by

an offline atmospheric reanalysis, a strong relaxation to a sea surface temperature product must be used in order to provide the negative feedback needed to keep the ocean temperature from drifting. The weekly Reynolds (Reynolds et al., 2002) SST product is used with a relaxation timescale of 5 days, and the World Ocean Atlas climatological salinity (Conkright and Coauthors, 1999) is used with a relaxation timescale of 10 days. Without relaxation of the SSS and SST to these products, the ocean surface state would begin to drift away quite rapidly.

The surface forcing for Hybrid-GODAS uses a bulk formulation based on Large and Yeager (2004). With this formulation the fluxes are calculated from the SST of the ocean model and the surface fields of the offline atmosphere. Since the model's SST is considered in these calculations, there exists a negative feedback the prevents the model SST from drifting too far away from nature. For this reason relaxation to an SST product is no longer as important and can be removed entirely if the model biases are small enough. Since there is no similar negative feedback for ocean salinity, an SSS restoration term is still required.

The atmospheric forcing for Hybrid-GODAS uses a combination of information from three different sources. These components are a 1) relatively high resolution mean forcing from the CFSR, 2) a climatological correction to the CFSR from DFS5.2, and 3) a low resolution set of ensemble perturbations from the 20th Century Reanalysis.

4.1.2.1 Mean Forcing from CFSR

The CFSR provides the mean surface forcing for the ocean from 1979 to 2011 at a resolution of T382. The CFSv2 operational analysis provides the mean forcing from 2011 to the present, however, due to it's higher resolution at T574 the surface fields are remapped to the lower resolution of T382 for consistency. For the Hybrid-GODAS, daily averaged fields of downward shortwave (DSW), downward longwave (DLW), and precipitation rate (rain + snow) are calculated. The ocean model calculates an artificial diurnal cycle using the daily averaged radiation fields when the model is run. Also, 6 hourly instantaneous fields for 2m temperature and humidity, mean sea-level pressure, and 10m zonal and meridional winds are used. Daily averaged winds had been used initially in the testing of Hybrid-GODAS, but this produced wind stresses that were too weak in the extra tropics. Hourly fields were tested as well, but the difference with 6 hourly fields was determined to be minimal for the purposes of this system.

Since precipitation rate in CFSR is provided as the total sum of liquid and frozen precipitation, the default MOM6 configuration is changed to partition frozen and liquid precipitation from the total based on the 2m atmospheric temperature.

4.1.2.2 Climatological Correction

There exist very large known biases in the CFSR fluxes. The CFS, as well as many other coupled climate models, fails to produce correct marine cloud patterns in several key regions. The eastern ocean basins are home to persistent marine

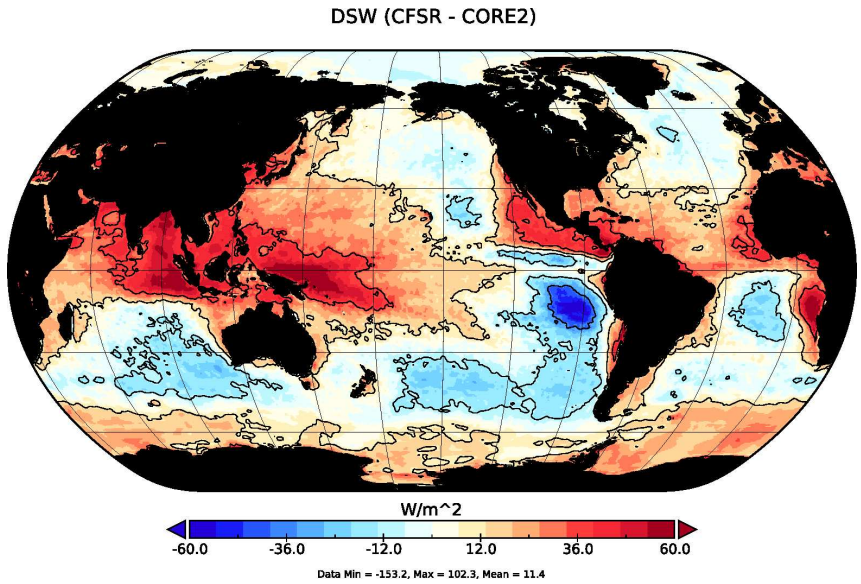

Figure 4.2: The CFSR exhibits a very large bias in the shortwave radiation of over 60 W/m^2. Shown is CFSR - CORE2 for the annually averaged downward shortwave in 2005.

stratocumulus clouds which the model does not produce well. As a result, shortwave radiation is too high, and longwave radiation is too low in these regions. In the western tropical pacific and Indian ocean, the downward shortwave radiation is too high by as much as 60 W/m^2. (Imagine an extra incandescent light bulb sitting over the ocean every meter!). This can be seen compared to other reanalysis products that have been calibrated to fit observations, fig 4.2.

The original GODAS had very strong relaxation to an SST product, which likely masked the effect of atmospheric forcing biases. Since the Hybrid-GODAS does not use any SST relaxation, these forcing biases must be handled as best as possible. The climatology of the DRAKKAR forcing set (DFS52) (Dussin et al., 2016) is used to correct the climatology of the CFSR. DFS52 is a product that is based on the ERA-interim reanalysis (Dee et al., 2011) and uses various observa-

tional datasets to correct the radiation, precipitation rate, and surface winds.

Monthly climatologies are calculated for the 1980 to 2015 period for both the CFSR and the DFS5.2. A multiplicative correction factor (Eq 4.1) is calculated for each month in the period for precipitation rate, downward shortwave, and downward longwave, and 2m humidity. An additive correction factor (Eq 4.2) is calculated for the winds and temperature.

$$corr_{mul} = (CFSR_{clim} - DFS52_{clim})/CFSR_{clim} \tag{4.1}$$

$$corr_{add} = CFSR_{clim} - DFS52_{clim} \tag{4.2}$$

The choice of an additive or multiplicative correction factor both result in identical climatologies, but the actual daily fields will be different. The multiplicative factor is used for fields that should not become negative (radiation, precipitation, humidity) when applying the correction. The winds receive an additive correction so that high wind events in the storm tracks don't receive an overly large increase in intensity as they would with a multiplicative correction factor.

The CFSR, as well as many other reanalysis products, exhibits serious shifts in its climatology due to abrupt changes in the observation platforms being assimilated. The largest known shift occurs between 1998 and 1999 due to the assimilation of the Advanced TIROS Operational Sounder (ATOVS). This jump resulted in a marked increase in global precipitation rate (Zhang et al., 2012). Many other atmospheric surface fields are affected in the tropics, including temperature, humidity,

field	description	freq.	clim. corr.
DSWRF	downward shortwave	daily	multiplicative
DLWRF	downward longwave	daily	multiplicative
PRATE	precipitation rate	daily	multiplicative
PRES	sea level pressure	6 hr	NONE
TMP	2m temperature	6 hr	NONE
SPFH	2m humidity	6 hr	NONE
UGRD	10m zonal wind	6 hr	additive
VGRD	10m meridional wind	6 hr	additive

Table 4.2: Overview of the atmospheric forcing fields supplied to the Hybrid-GODAS. Shown for each field are the short names used within the code, description, whether daily averaged or 6 hourly instantaneous frequency, and the type of climatological correction applied.

and winds. As a result, two climatology periods 1980-1998 and 1999-2015 are used when calculating the correction factors. These are shown in Fig 4.3 and Fig 4.4. It should be noted though that monthly correction fields are calculated and used by the model, though only an annual average is shown in these figures for simplicity.

Using the corrected fields with a free running model ocean (as described later), the model SST exhibits a cold bias especially in the tropics. As has also been found in applications of DFS5.2 forced SODA (James Carton, personal comm.) it appears that there is a bias in the 2m temperature and humidity in the DFS5.2 that results in too much cooling of the ocean through latent and sensible heat fluxes (fig 4.5). For this reason, the temperature and humidity corrections that have been calculated for CFSR from the DFS52 are not applied. By only applying a correction to the radiative, precipitation, and wind fluxes, an SST in the tropics with less bias is produced. A summary of the final configuration for the atmospheric forcing files is given by Table 4.2.

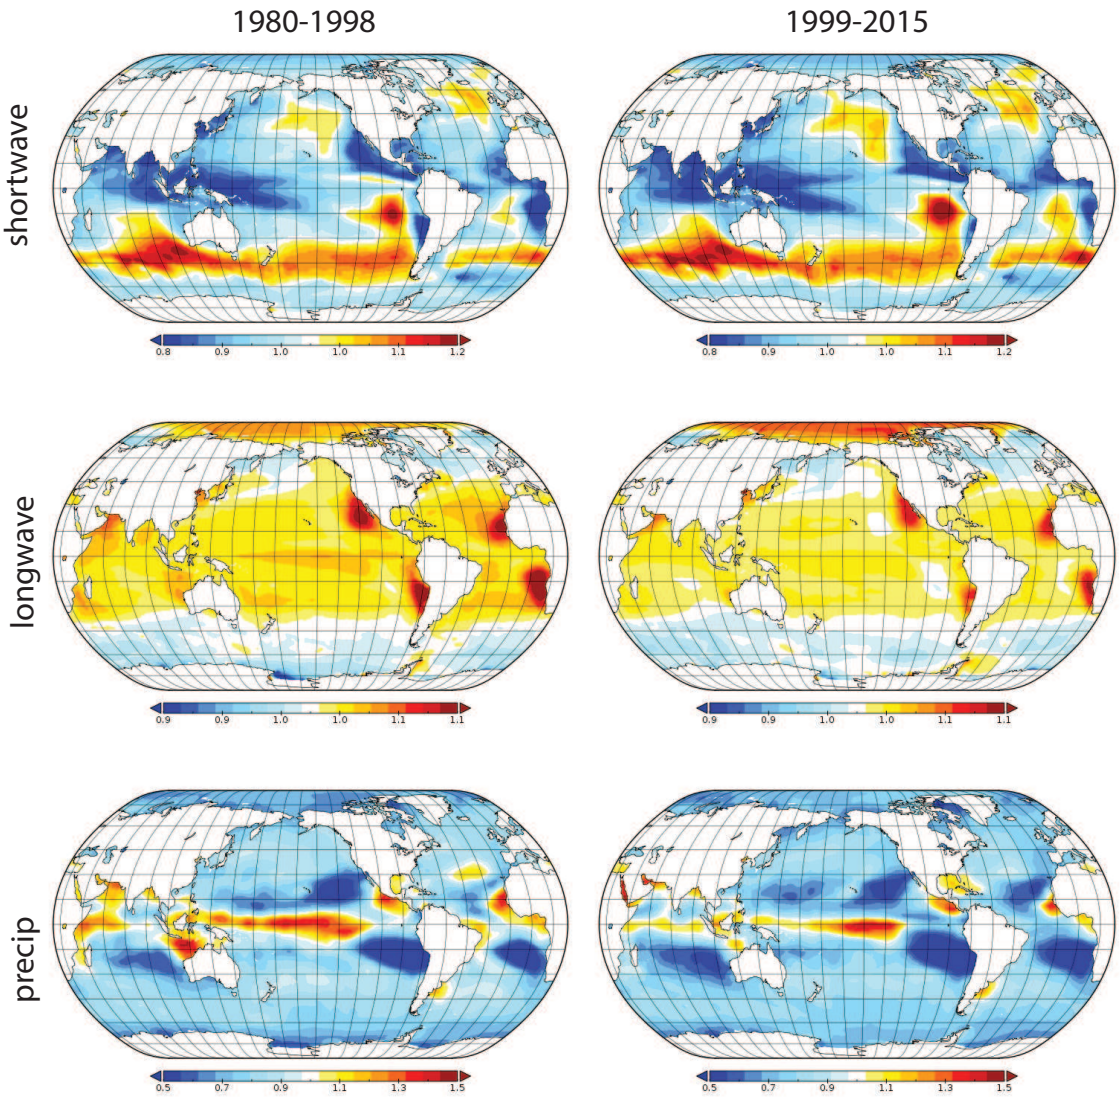

Figure 4.3: The multiplicative bias corrections that are applied to the downward shortwave, downward longwave, and precipitation rate, as calculated from the climatology difference between CFSR and DFS52. Monthly bias corrections are generated, but only the annual average is shown here.

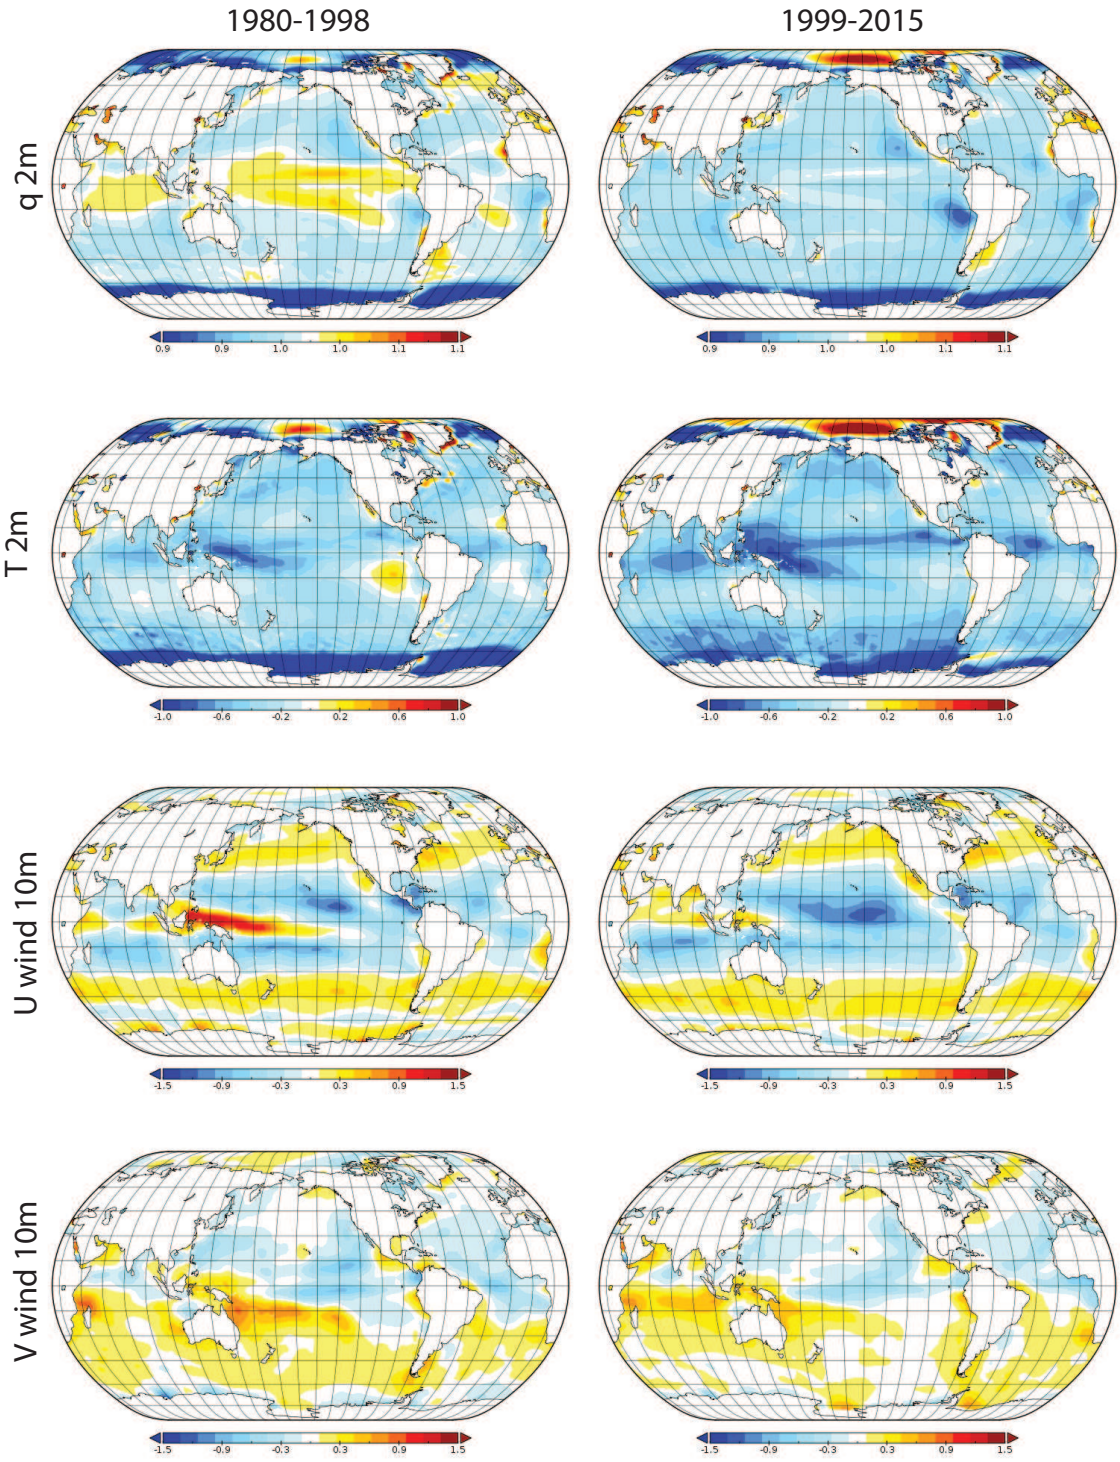

Figure 4.4: The multiplicative bias corrections for 2m humidity, and the additive bias corrections for 2m temperature and 10 meter winds, as calculated from the climatology difference between CFSR and DFS52. Monthly bias corrections are generated, but only the annual average is shown here.

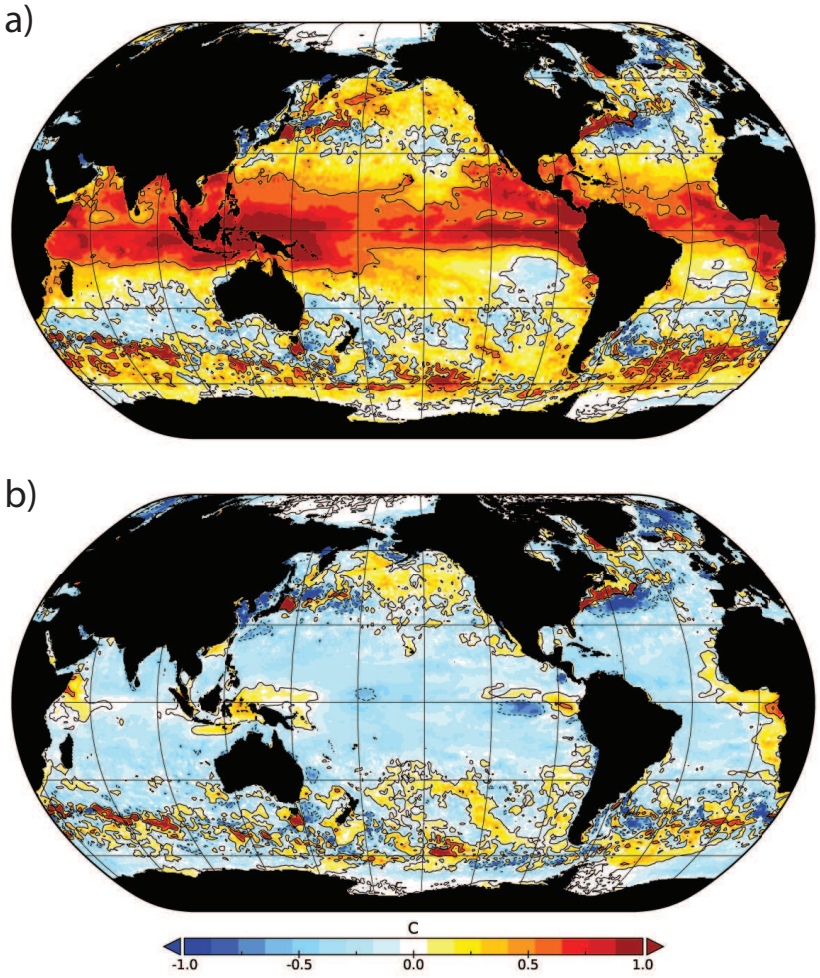

Figure 4.5: Difference in the ocean SST for forced run without data assimilation compared with DA analysis, averaged over 1 year. Shown is a forced run using uncorrected CFSR fluxes (a) and CFSR fluxes with climatology corrected by DFS5.2 (b).

4.1.2.3 Ensemble Perturbations

After the climatological correction is applied to the CFSR fields, ensemble perturbations for each ocean ensemble member are applied. These perturbations are derived from the 6 hour forecast fields of the 20th Century Reanalysis v2, 20CRv2 (Compo et al., 2006). The 20CRv2 fields are not used directly because of the low resolution of the reanalysis (T62). The 20CRv2 is a 56 ensemble member reanalysis that assimilates only surface pressure. It is chosen for its long timeline available (1851-2014).

As can be seen in Fig 4.6, most of the spread in the 20CRv2 is located in the tropics. The extra-tropics are suspected of having insufficient spread in the T2m, Q2m, and wind fields for our purposes. This is expected given the coarse resolution of the 20CRv2, and the fact that the ocean SST driving the 20CRV2 ensemble members are very similar. Future improvements to Hybrid-GODAS will use a higher resolution reanalysis (20CRv3, which is in the works) or the ensemble reanalysis from a coupled model.

In addition to the atmospheric surface forcings, the ocean also requires river and land water runoff fields. The monthly Dai-Trenberth (Dai, 2016) climatology is used for this purpose. In initial tests Hybrid-GODAS had shown difficulty in maintaining the salinity spread, and seemed especially weak near the coastlines. Stochastic perturbations to the runoff fields are therefore used to help add salinity spread to the ocean.

The monthly climatological variability of the Dai-Trenberth dataset is first

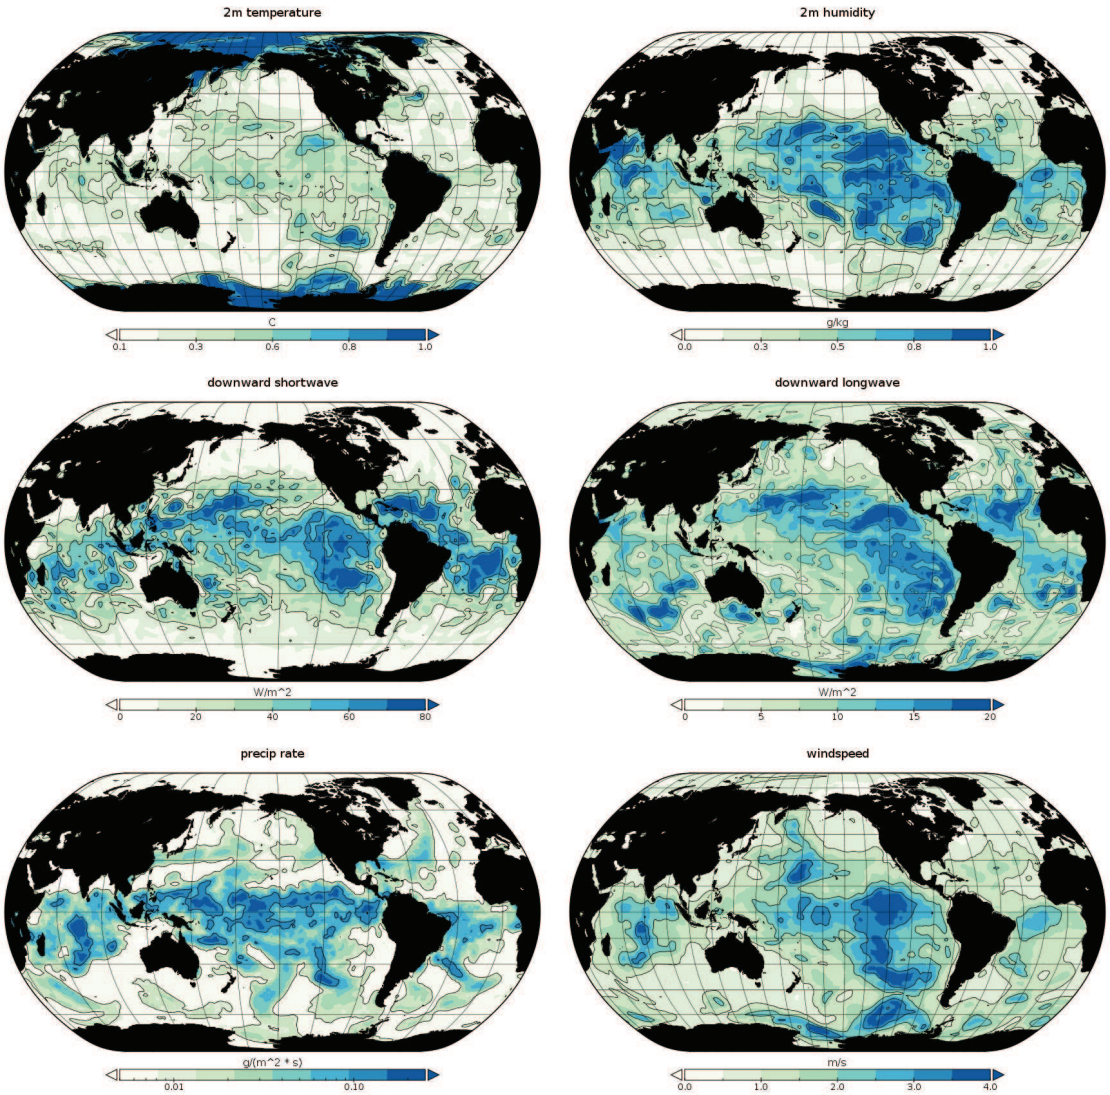

Figure 4.6: Ocean surface forcing spread from the 20CRv2 for a single date, 2003-03-15, for 2m temperature, humidity, downward shortwave, downward longwave, precipitation rate, and windspeed.

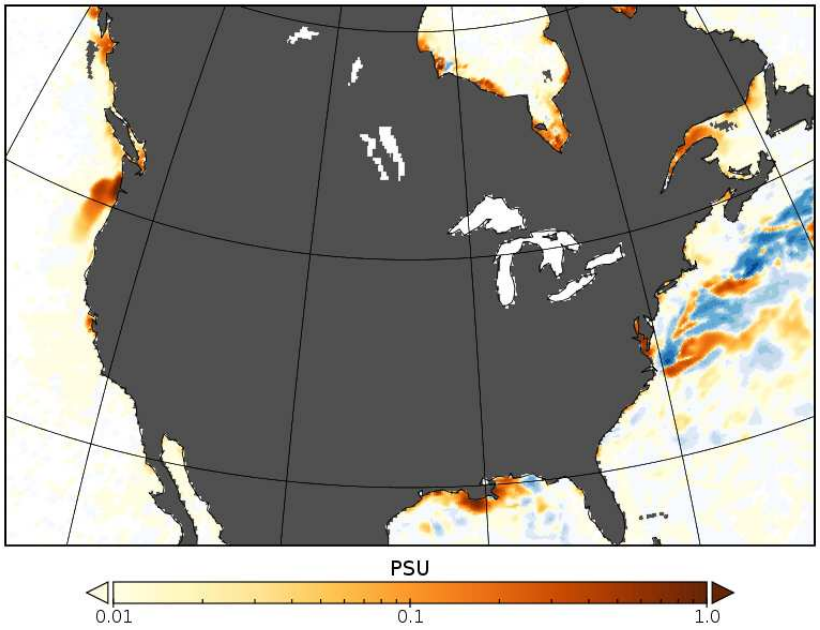

Figure 4.7: The increase in surface salinity spread after 1 month data assimilation cycle when turning on stochastic land/river runoff perturbations.

calculated using the available interannual forcing files from 1948 - 2007. Then a Perlin noise is generated for each ensemble member that slowly varies from month to month. This noise generates a field varying from -1.0 to 1.0 with per-ensemble member horizontal and temporal correlations. Each ensemble member's stochastic noise is then multiplied by the climatological variability, and then added to the climatology. In this way, short of obtaining a better estimate of the actual monthly river and continental runoff in real-time, this method adds a slightly more realistic spread to the runoff climatology for the ocean ensemble.

There is a resulting increase in the salinity spread along the coastlines, and especially near rivers, Fig 4.7. It is not clear how much of an impact this actually has in the ensemble ocean DA system, as the spread dissipates rapidly as the water

mass leaves the coastline, especially when the nearly-global coverage of satellite SST observations are used. However, this method is likely to be of more importance when ocean data assimilation is applied on regional scales.

4.1.3 Observations

The inclusion of additional observations is another area of significant progress with the Hybrid-GODAS. The previous GODAS only assimilated observed temperature profiles. Salinity was constrained by assimilating a synthetic salinity that was calculated from an observed climatological temperature / salinity relationship. While this served well for purposes of analyzing temperature, the salinity fields were always very close to climatology and exhibited very little inter-annual variability. Hybrid-GODAS has been upgraded to use insitu temperature, insitu salinity, and along track satellite SST retrievals. Future plans also include the use of satellite altimetry.

Insitu temperature and salinity profiles are obtained from the world ocean database (WOD) (Boyer et al., 2013). Only the profiles with the highest quality control flags are used. There are multiple platform types in WOD, but only a subset are used here: expendable bathythermographs (XBT), moored buoys (MRB), and profiling floats (PFL). The profiles are temporally and spatially averaged so that multiple profiles from a single platform, in a single grid point, in a single day, are averaged together. These observations are suitable for testing and reanalysis, but observation quality control procedures are being developed so that ocean profiles

from the NCEP BUFR data tanks can be used going forward.

The previous GODAS used satellite SST indirectly by relaxing the top layer of the ocean model toward the weekly Reynolds SST product (Reynolds et al., 2007). This method resulted in an adequately accurate SST from GODAS, but was not taking full advantage of all the information SST observations could offer. Hybrid-GODAS directly assimilates the SST by using the along track retrievals from NOAA's Advanced Clear Sky Processor for Ocean, ACSPO, (Ignatov et al., 2016). When assimilated, the 3DVar and LETKF use the satellite SST to impact the entire mixed layer. This results in an instantaneous correction to the mixed layer with Hybrid-GODAS, whereas GODAS would take much longer to impact the mixed layer since only the top layer is being relaxed toward observed SST. Additionally, the LETKF produces a multivariate update, and so SST observations can impact the salinity and ocean currents, which will be shown to be important in maintaining more accurate western boundary currents in the model background during the DA cycle.

The definition of "SST" is an ambiguous term. For our purposes, the SST is represented by the top level of the model, at 1m. However, due to stratification in the upper centimeters and even millimeters of the ocean, satellites observe a different SST depending on the time of day, underlying stratification of the ocean, and the amount of mixing from the wind. Satellites actually observe a skin-SST, which for infrared satellites is at a depth on the order of a centimeter. During the day, this skin-SST will be significantly warmer than the temperature just below this at 1m. During the night time, this warm skin disappears, and the skin-SST is more similar

to the bulk SST, although a cool skin layer is still present due to evaporative cooling. For this reason, only the night time tracks of the SST retrievals are assimilated.

Initially, the Hybrid-GODAS reanalysis was planned to use the AVHRR Pathfinder dataset, version 5.3 (Casey et al., 2010). Pathfinder presents an SST retrieval spanning from the beginning of the satellite era (1979) to near present (2012). The best quality AVHRR satellite at any given time is used, and efforts are made to bias correct among the satellites for a continuous reliable record. Due to the extensive record presented by this dataset, it is widely used.

However, large biases were discovered in Pathfinder while developing the Hybrid-GODAS, which were negatively impacting the resulting analysis. A negative bias exists in the tropics, and a positive bias in the extra-tropics, which can be seen by comparing the insitu and satellite SST observation minus forecast (O-F) statistics. As can be seen in Fig 4.8a, the insitu observations were constantly trying to cool the model due to excess warming from the SST observations. Additionally, when comparing Pathfinder to other SST retrievals such as Reynolds, Fig 4.8b, these biases are more apparent. Pathfinder is known to have shifts in the bias when the predominant satellite is changed, and has poor performance when dealing with aerosols and cloud contamination. Most importantly, Pathfinder is not maintained in real-time.

Hybrid-GODAS was switched to the Advanced Clear Sky Processor for Ocean (ACSPO) produced by NOAA/NESDIS (Ignatov et al., 2016). This SST retrieval is superior to Pathfinder in that the satellite SST minus insitu temperature differences are far smaller, as confirmed by NESDIS with their quality control routines, and by examining O-F statistics in the Hybrid-GODAS. The best two satellites at any

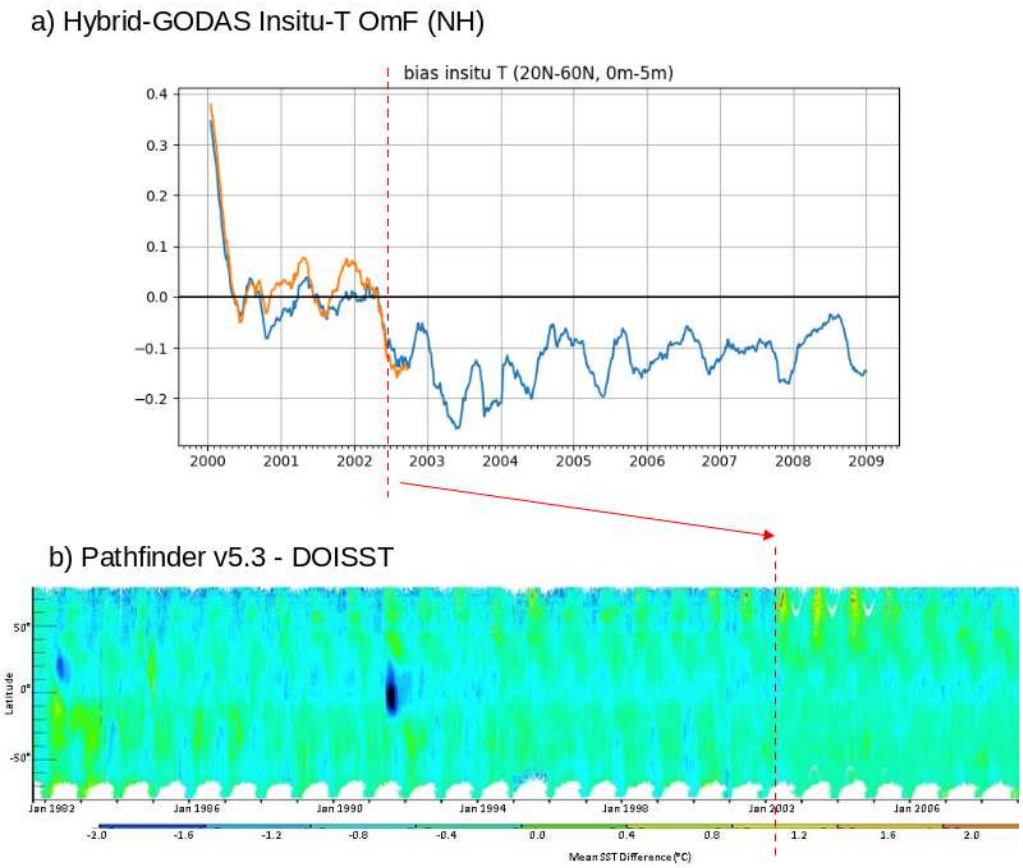

Figure 4.8: Insitu O-F bias in an early Hybrid-GODAS run (a) and the difference between Pathfinder and daily OISST (from https://www.nodc.noaa.gov/SatelliteData/pathfinder4km53/). Pathfinder exhibits large biases that change when satellites are changed (e.g. vertical dashed lines on July 2002).

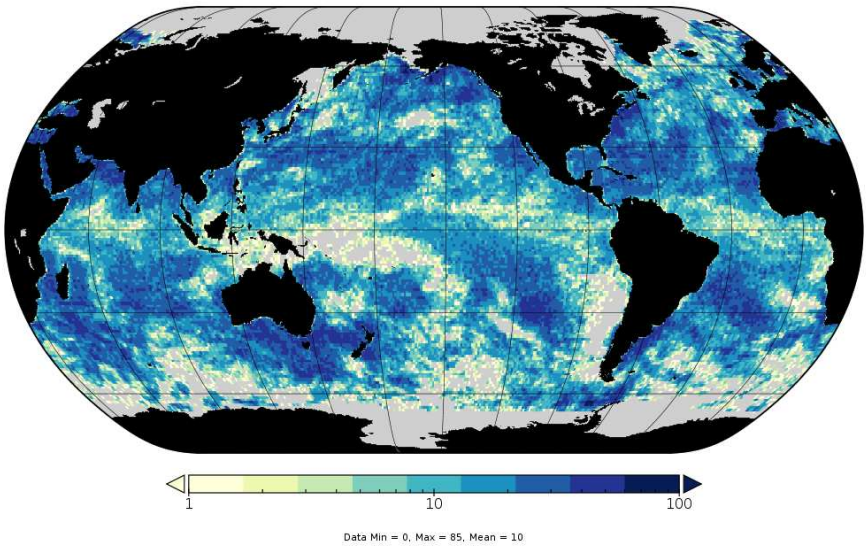

Figure 4.9: Satellite SST observation density for a single 5 day cycle (January 1 to January 5 of 2004). Shown is the number of satellite SST observations, after superobbing, per 1°.

given time are used with ACSPO, and a wider scan angle of the AVHRR satellites is used. This increases the number of available observations in the dataset compared with Pathfinder, and end up providing nearly global coverage of every model grid box in any given 5 day data assimilation cycle (Fig 4.9), except in areas of persistent cloud cover.

As with the Pathfinder dataset, only the observations at nighttime and with the highest quality control flags are used. Another benefit of the ACSPO dataset is that an estimate of the skin-SST to 1m-SST bias correction is given, as well as an observation error estimate. The observation error estimate primarily takes into consideration the errors from cloud contamination and the reduced quality of the SST observations near the edge of the satellites swath. The provided bias correction

term was found to be adequate to remove biases between the satellite SST and insitu top level temperature observations within Hybrid-GODAS.

Before being assimilated, the satellite SST observations undergo superobbing so that there is at most 1 observation per grid-box per day. Otherwise, the over abundance would lead to difficulty in maintaining the ocean ensemble spread. To help account for errors of representativeness, the variance of the observations going into each grid-point before the superobbing is calculated. The final observation error variance then equal to the variance of the superobbed observations plus the estimated error variance from the ACSPO dataset.

4.1.4 Data assimilation

The upgrade to the data assimilation system, from a simple 3DVar to the hybrid gain EnKF/Var (Penny, 2014), is the central motivation for the Hybrid-GODAS project. In the atmosphere, operational centers have moved to more advanced hybrid EnKF/Var systems (Kleist and Ide, 2015; Bonavita et al., 2012), relying on ensemble perturbations from an EnKF system to provide the dynamic part background error covariance used in the variational solver. In such systems, the variational solver is the true work-horse of the system. With the hybrid-gain, the opposite is true in that the EnKF solver is the system's work-horse. Described in more depth in Penny (2014), the essence of the hybrid gain solver (green box in fig 4.1) is that the EnKF creates an analysis, which then has its mean partially corrected by a variational solver that is run using the analysis created from the EnKF as it's background.

By using the 3DVAR after the EnKF, the 3DVar allows the system to explore parts of the model state space that are not represented by the ensemble. This is essential when using a small ensemble size, as we are doing with Hybrid-GODAS's limited ensemble of 20 members. Model biases and highly non-linear regions of the ocean that are not captured well by the ensemble such at near the Gulf Stream could ultimately cause the EnKF to undergo filter divergence and begin to drift from the observations, but the 3DVar helps mitigate this and allows the data assimilation to remain stable with a much smaller ensemble size than would normally be possible with just an EnKF implemented.

A significant benefit of switching from 3DVAR to hybrid-gain is that the data assimilation system will now be able to create multivariate analysis increments. GODAS is only able to update temperature and salinity independently. Hybrid-GODAS will update all state variables (temperature, salinity, currents), likely resulting in a better analysis and a better balance in the analysis, as shown in Penny et al. (2015). As shown in the NEMOVAR ocean data assimilation system (Mogensen et al., 2012) maintaining a proper temperature/salinity balance near the thermocline is important for a good quality analysis. GODAS was not able to do this, but the EnKF portion of Hybrid-GODAS now should be able to.

Hybrid-GODAS makes use of the LETKF and a new observation space 3DVar. The LETKF used is the UMD-LETKF implementation (described in chapter 5) (Sluka, 2018b), that aims to become the standard LETKF implementation used at NCEP. The UMD-LETKF is a generic model-independent solver, and seeks to make future SCDA easier by ensuring the various domains of a coupled model are using an

identical LETKF code base. The current LETKF configuration for Hybrid-GODAS uses relaxation to prior spread of 60% for covariance inflation, and a latitudinally dependent localization that varies from 600km at the equator to 100km in the high latitudes for insitu observations, and a smaller 200km EQ to 50km high latitudes for satellite SST observations. The smaller localization radius for satellite SST observations is important, otherwise the large number of SST observations leads to an overly small ensemble spread.

The observation space formulation solver for the 3DVar is modeled after the Navy's NCODA (Daley and Barker, 2001; Cummings and Smedstad, 2013), and is described in more detail in Appendix A. The observation space formulation allows the solver to run much faster than the previous state space GODAS, given that ocean observations are very sparse compared with the number of 3D grid points. One current shortcoming though is the lack of multi-variate balance operators, so for the time being this new 3DVAR is essentially performing separate temperature and salinity analyses such as GODAS was doing. This will be remedied in the future by migrating toward the JEDI ocean 3DVAR system, which is still under development, but will be a more sophisticated multivariate 3DVar with background error covariance model similar to NEMOVAR (Mogensen et al., 2012). Within the hybrid-gain framework, the 3DVAR uses the analysis mean from the LETKF analysis, and applies a percentage of the analysis increment (currently configured as 50%) to adjust the analysis ensemble mean.

Hybrid-GODAS has been built so that switching between data assimilation methods (freerun, 3DVar, EnKF, hybrid EnKF/Var) is done simply through the

configuration scripts.

4.2 Initial Results

Iterative tests are being conducted with a 20 ensemble member Hybrid-GODAS to tune the system before implementation into operations at NCEP. To setup the test experiments, the ocean model ensemble members are initialized with identical climatological temperature and salinity from the World Ocean Atlas 2013 (Locarnini et al., 2013; Zweng et al., 2013). The hybrid data assimilation is then run from Jan 1, 2003 to Jan 1, 2004, and the ensemble members at the end date are saved as the initial conditions for subsequent experiments. All ensemble members start with initial conditions, however due to the atmospheric forcing perturbations, the spread in the ocean ensemble quickly grows after several months. At the beginning of the run, the 3DVar will be doing all of the work, since the initial members are all nearly identical and therefore LETKF will have no impact on the analysis. As the data assimilation cycle progresses, the spread increases and the LETKF begins to have more of an impact. This gradual increase in LETKF impact prevents large initial shocks to the system as was experienced with the CFS-LETKF experiments. This initial spinup method is used for simplicity for these experiments, though other methods of initial ensemble generation could be used.

From these Jan 1, 2004 initial conditions, experiments with several configurations are conducted. Hybrid-GODAS is run with 1) no data assimilation, just prescribed atmospheric forcing 2) 3DVar data assimilation mode, and 3) hybrid-gain

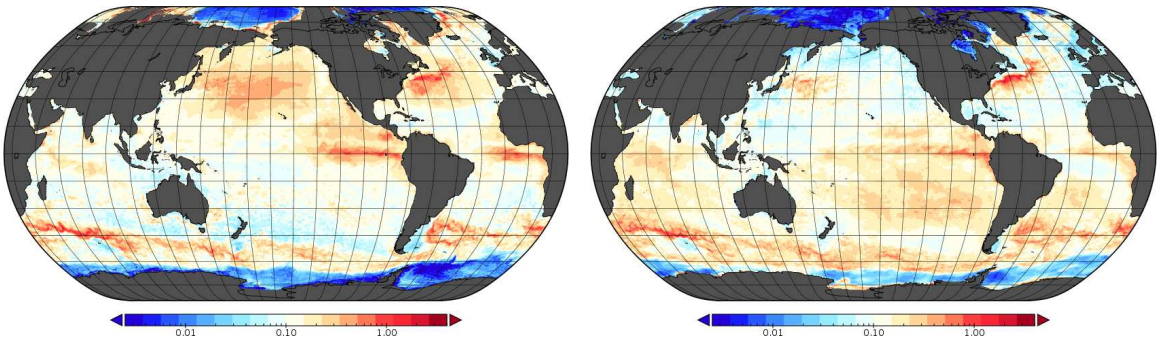

(a) (b)

Figure 4.10: Hybrid-GODAS ensemble spread at the surface averaged over June (left) and December (right) of 2004. Surface spread is higher in the summer months.

EnKF/3DVar data assimilation. The results are then compared against several data sets, including the original GODAS, the UK MetOffice EN4, OISST, and OSCAR surface current estimates.

The resulting ensemble spread for the hybrid-gain run is shown in Figure 4.10 and Figure 4.11. After tuning the LETKF localization and satellite observation thinning parameters, the resulting spread looks reasonable. There is increased spread, and corresponding uncertainty in the analysis, in the mid-latitude ocean surface during the summer months, along the western boundary currents, and along the thermocline in the tropics.

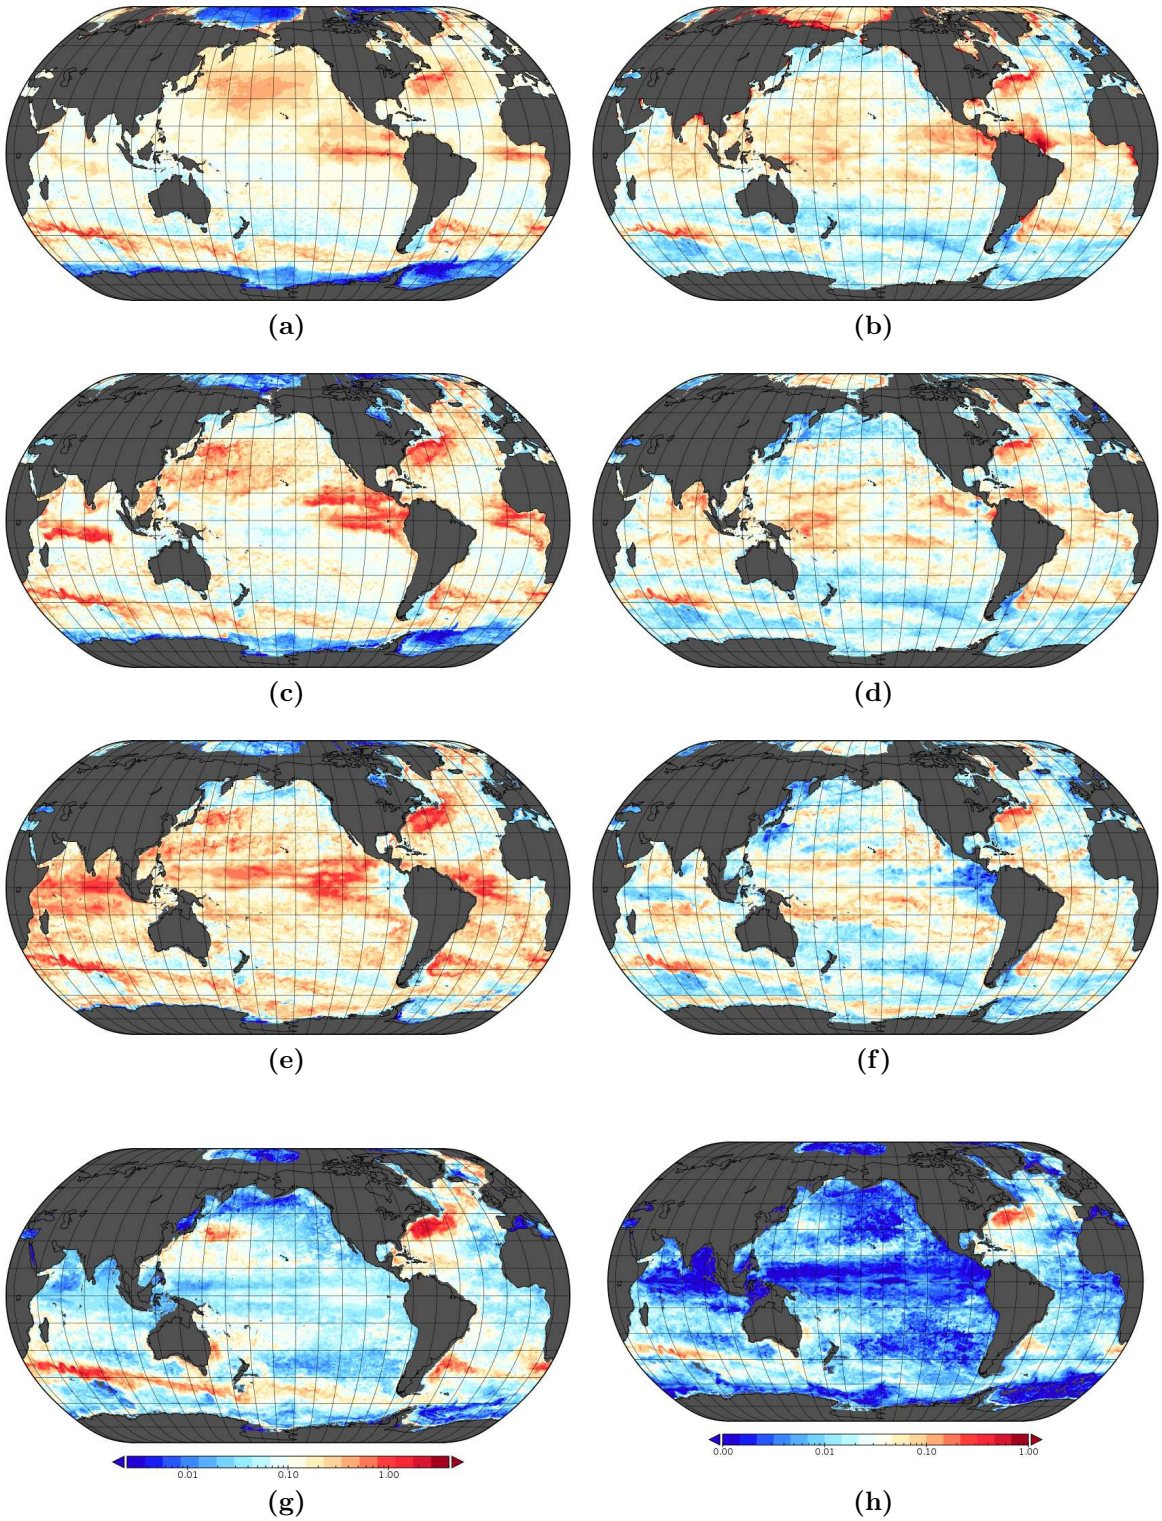

Figure 4.11: Hybrid-GODAS ensemble spread averaged over the June 2004 for temperature in C (left) salinity in PSU (right) at depths of 1m (a,b), 50m (c,d), 100m (e,f), and 500m (g,h).

4.2.1 Comparison to GODAS/EN4/CMC

To evaluate the performance of Hybrid-GODAS, it is first compared against the pentad analysis from GODAS and either the UK MetOffice EN4 objective analysis (Good et al., 2013) for temperature and salinity at depth, or the SST product from the Canada Meteorological Center (Brasnett, 2008) for sea surface temperature. The following issues should be noted while examining the results:

- For technical reasons in the file processing, the 5 day **background** average (not the analysis), was used for the Hybrid-GODAS. This is due to the fact that IAU has not yet been implemented for Hybrid-GODAS.

- The Hybrid-GODAS freerun (marked fr.02 in some of the figures), was initialized on 2004-01-01 from a spinup with hybrid DA, and so a subsequent longer freerun is likely to show a bigger difference between the freerun and data assimilation run.

- The EN4 and CMC datasets are not truth, they have their own errors, and so other methods of verification (e.g. O-F statistics) will be discussed later

It will be seen though that despite these considerations, Hybrid-GODAS largely outperforms GODAS in several key areas.

To evaluate the performance of sea surface temperature for Hybrid-GODAS, the RMSD compared with the daily SST product from the Canadian Meteorological Center (CMC) are calculated over the year 2004 and shown in fig 4.18. Unsurprisingly, Hybrid-GODAS has better SST due to the fact that it directly assimilates

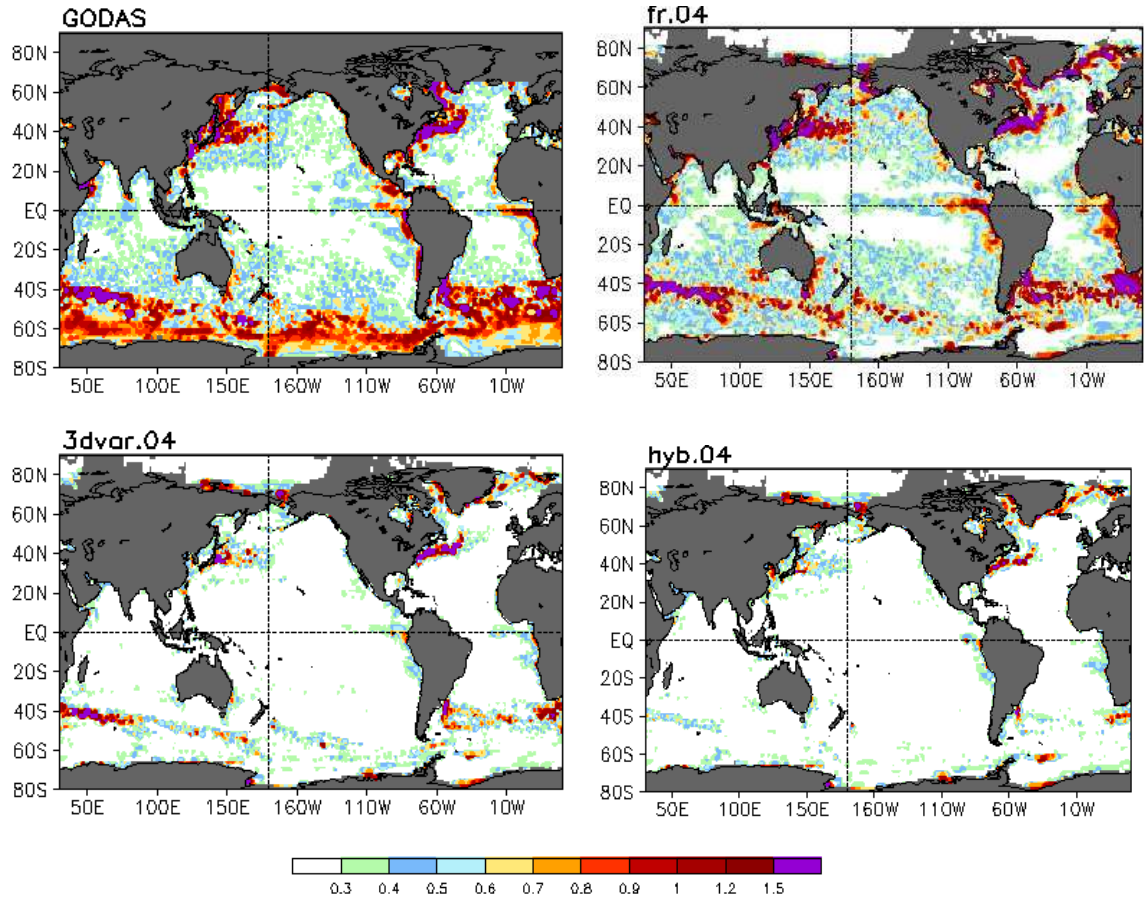

Figure 4.12: SST RMSD (C) as compared with CMC SST averaged over 2004 for GODAS (top left), Hybrid-GODAS forced run with no DA (top right), 3DVAR component only of Hybrid-GODAS (bottom left), and the full 3DVAR/LETKF Hybrid-GODAS (bottom right).

the along-track SST retrievals from ACSPO, whereas GODAS relaxes to OISST. Hybrid-GODAS shows marked improvement globally, especially in the Southern Oceans, the eastern coastline of the Pacific and Atlantic basins, as well as along the Western boundary currents. The Hybrid-GODAS is also an eddy-permitting model, at 1/4 $^\circ$, and so better resolves the position of the Western boundary currents, especially once altimetry data is assimilated.

The timeseries of model minus OISST RMSD averaged over the globe (fig 4.13)

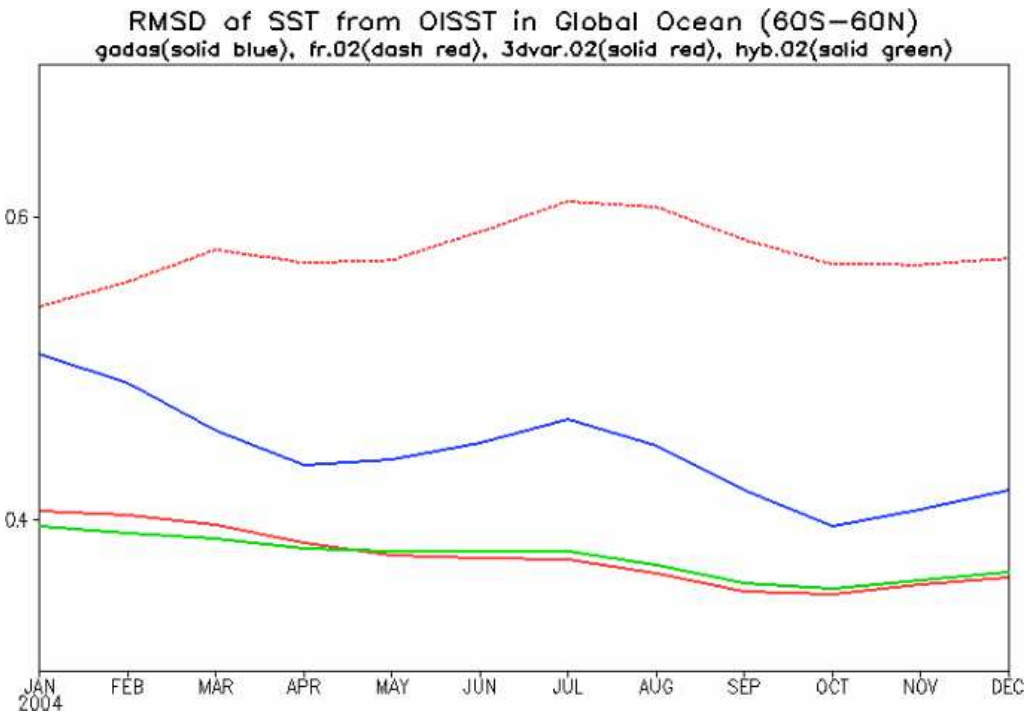

Figure 4.13: SST RMSD compared with OISST for 2004 averaged globally over 60S to 60N. Shown are GODAS (blue) Hybrid-GODAS (red), Hybrid-GODAS with only 3DVar on (green) and Hybrid-GODAS with no DA (dashed red).

shows again that Hybrid-GODAS has less error over time, with an average SST RMSD of 0.4C. Even without any data assimilation, the SST RMSD only grows as high as 0.6C, showing that it is well constrained by the atmospheric surface forcings.

The heat content in the upper 300 meters (fig 4.14), unfortunately, does not show marked improvement yet. There is some improvement with Hybrid-GODAS in the Southern Ocean, however, there are currently known issues being addressed that are degrading performance in the western boundaries and open ocean. These issues, mainly the lack of vertical localization with the LETKF for satellite SST observation, and quality control of profile observations, are being addressed and should hopefully be fixed in the next iteration of Hybrid-GODAS being tested before implementation

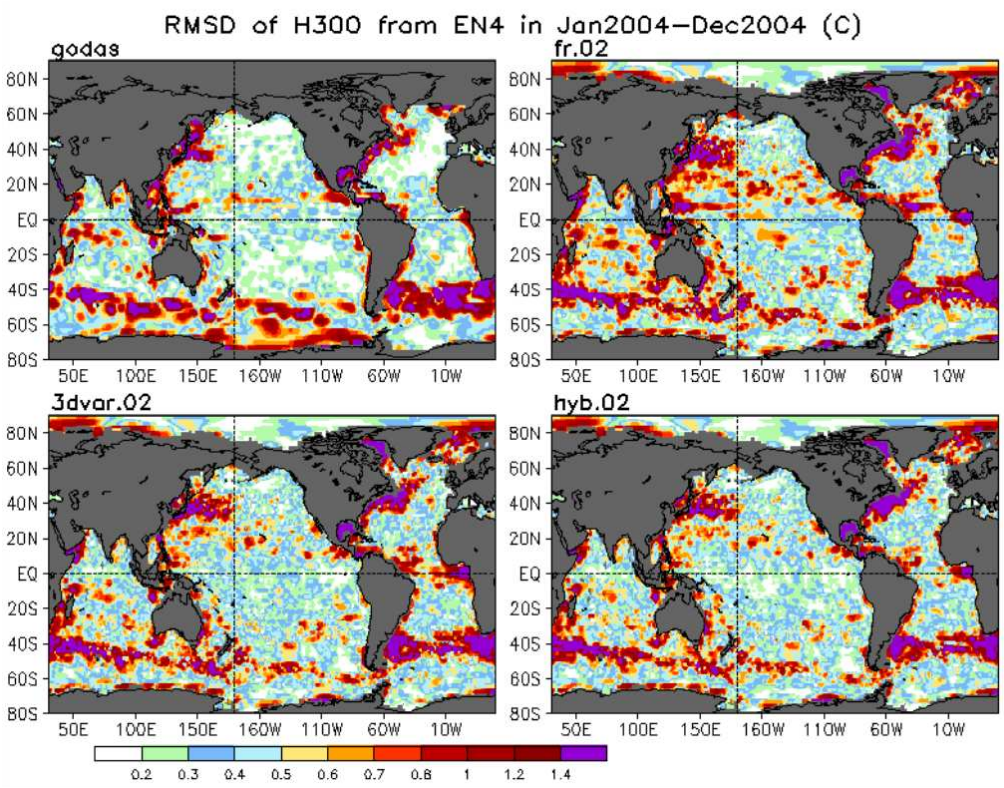

Figure 4.14: Heat content in upper 300m RMSD as compared with EN4 averaged over 2004 for GODAS (top left) Hybrid-GODAS with no DA (top right), Hybrid-GODAS with only 3DVar (bottom left), and Hybrid-GODAS with EnKF/3DVar (bottom right).

into operations.

The sea surface salinity RMSD compared against EN4 (fig 4.15) shows that Hybrid-GODAS is clearly improved. This is not surprising given that we now assimilate actual salinity observations, whereas GODAS only used synthetic salinity. Noticeable decreases in RMSD are visible in the tropics, however, occasional large salinity errors in Hybrid-GODAS along the west coast of equatorial Africa are apparent. It is suspected that the cause of this is the same cause of the H300 underperformance (lack of LETKF vertical localization, and observation quality control), and will therefore hopefully be improved in the next Hybrid-GODAS test iteration.

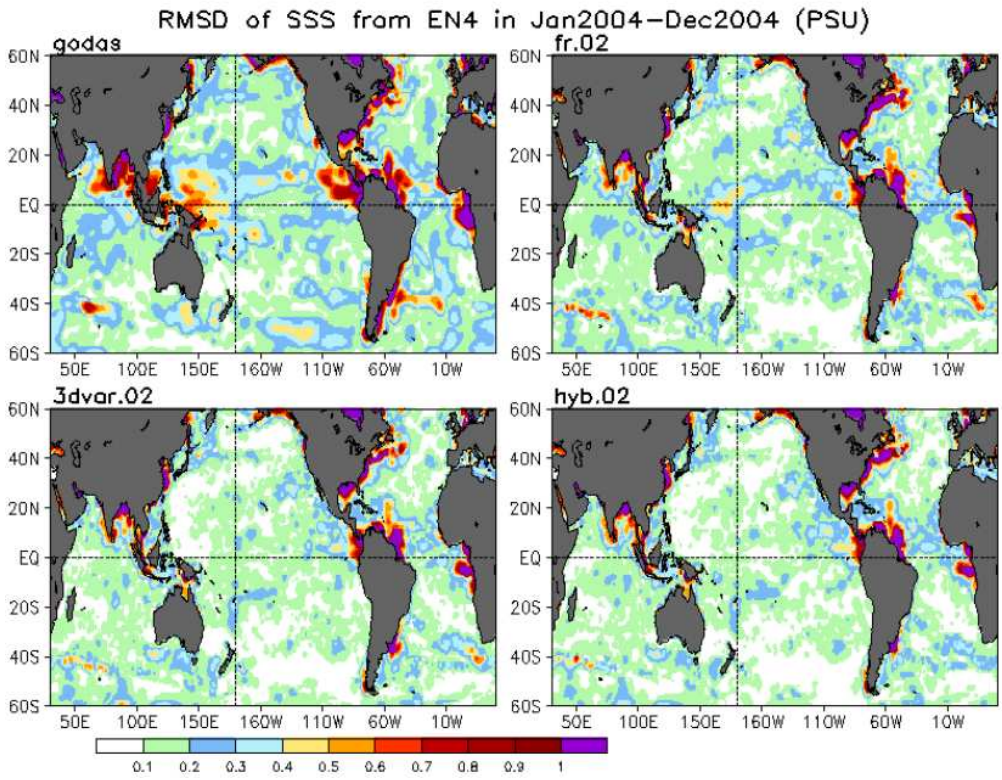

Figure 4.15: Sea surface salinity RMSD as compared with EN4 averaged over 2004 for GODAS (top left) Hybrid-GODAS with no DA (top right), Hybrid-GODAS with only 3DVar (bottom left), and Hybrid-GODAS with EnKF/3DVar (bottom right).

Salinity RMSD compared against EN4 in the deeper ocean levels of 300m to 750m (fig 4.16) again shows the improvement in Hybrid-GODAS especially in the extra-tropics. There is a decrease in performance near the Gulf Stream and equatorial Atlantic, likely due to the previously mentioned outstanding DA issues. It should also be noted that the difference between the hybrid-DA, 3DVar, and freerun versions of Hybrid-GODAS are small. This similarity in results is due to the fact that all three of those experiments were initialized with the same initial conditions on Jan 1, 2004, from an ensemble run made with hybrid-DA over 2003.

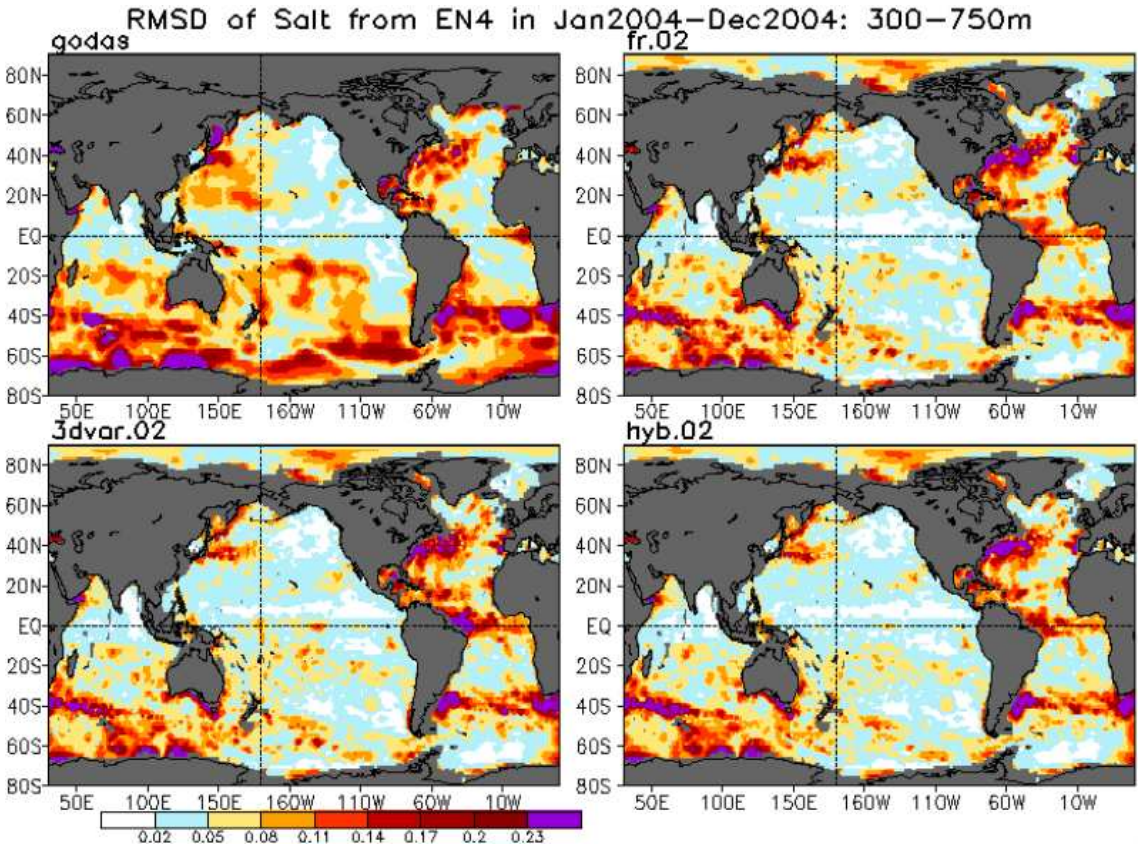

Figure 4.16: Ocean salinity RMSD as compared with EN4 averaged over 2004 between 300m and 700m for GODAS (top left) Hybrid-GODAS with no DA (top right), Hybrid-GODAS with only 3DVar (bottom left), and Hybrid-GODAS with EnKF/3DVar (bottom right).

The results would expect to diverge more as the experiments are run into later years.

4.2.2 Comparison to OSCAR

An important test of the data assimilation system is to compare against observations that are not assimilated into the system. Here we look at the surface currents from OSCAR (Ocean Surface Current Analysis Real-time) (Bonjean et al., 2002). OSCAR estimates surface ocean currents from indirect observations such as

sea surface height, SST, and surface winds, and then using dynamical balances such as geostrophic balance to calculate surface currents. A 1/3 degree grid is produced every 5 days.

The western boundary currents can be difficult for an ocean data assimilation system to properly capture. The Kuroshio current, as observed by OSCAR, is shown in fig 4.17a for a single day in Oct 2004. The GODAS analysis (fig 4.17b) is not capable of reproducing the Kuroshio because of the coarse horizontal resolution of the model, 1 degree. There is a resemblance of the currents in the GODAS analysis, however it is too weak and diffuse, with a maximum current speed less than 0.5 m/s, compared with the observed $>$ 1.5 m/s. The Hybrid-GODAS is eddy-permitting due to its 1/4 degree horizontal grid. As a result, a forced run will spontaneously produce a Kuroshio current (fig 4.17c) even though it is not in the correct place. Using Hybrid-GODAS with 3DVar-only enabled (fig 4.17d) is still not able to put the current in the correct location. The Hybrid-GODAS with full hybrid DA, however, is able to place the Kuroshio current (fig 4.17e) with reasonably accurate location and speed of the meanders. There are even some eddies present in OSCAR that the hybrid DA is trying to place in the model. This is a good result, especially considering that altimetry is not being assimilated yet in Hybrid-GODAS. Once altimetry is being assimilated, it is expected that the placement of the western boundary currents and some eddies could be even better.

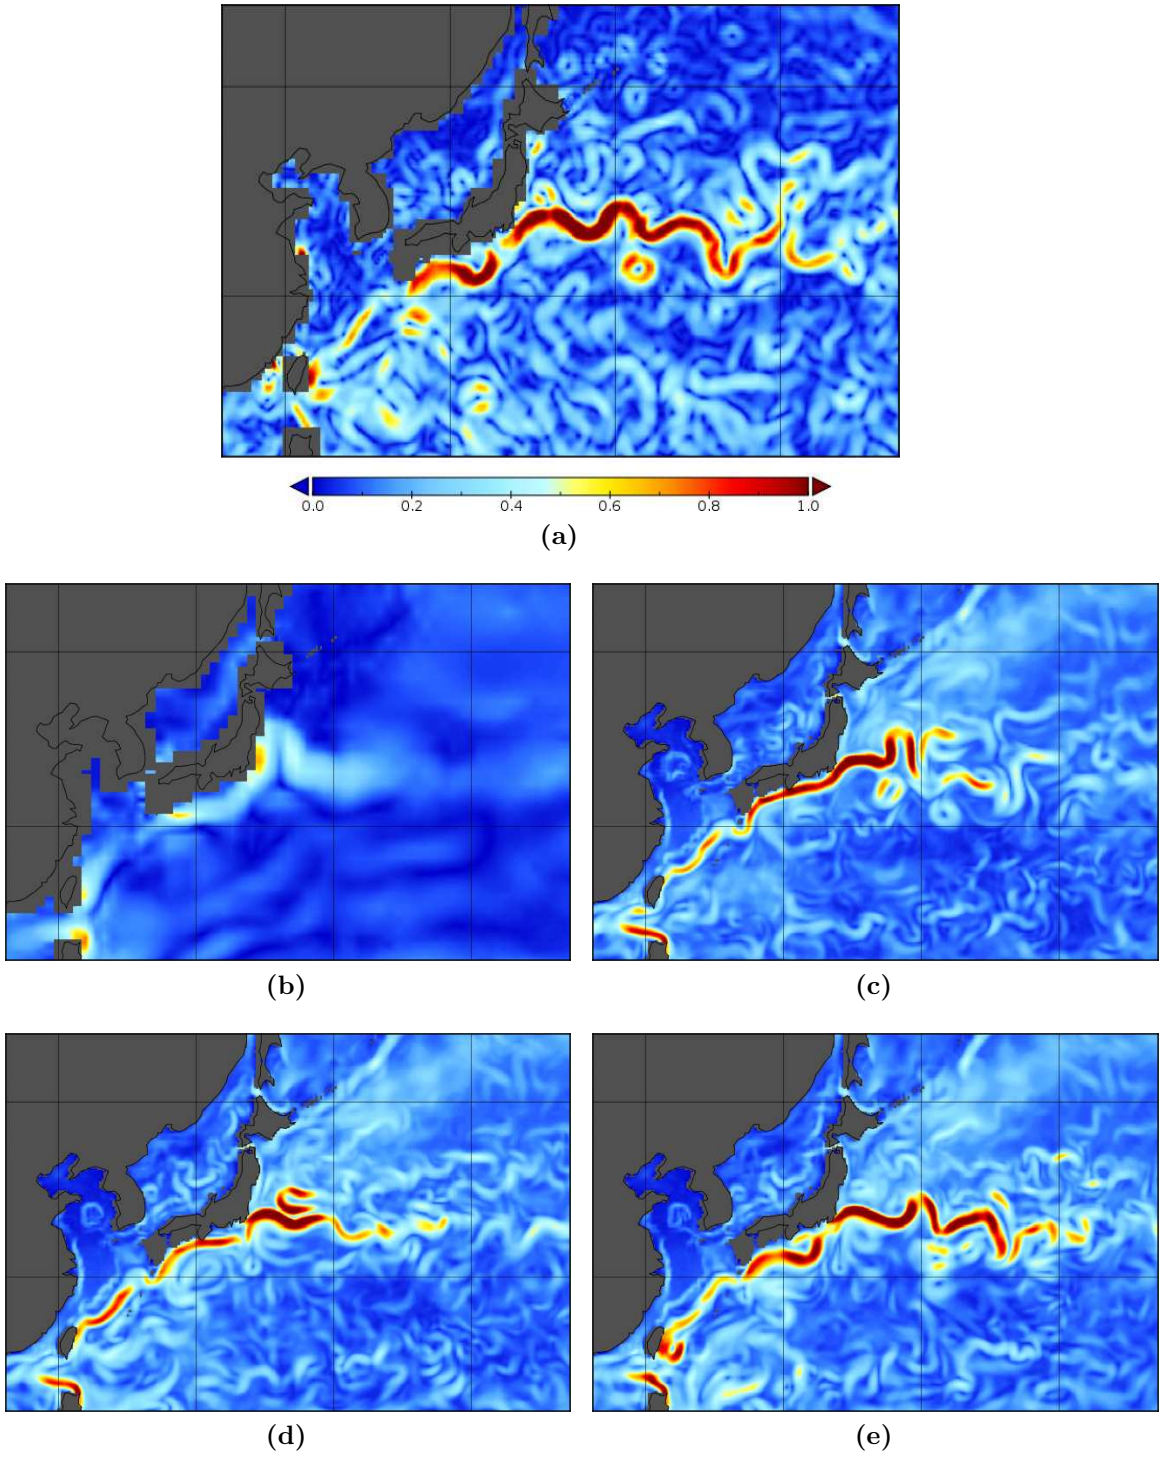

Figure 4.17: Top level ocean current speed on Oct 6, 2004, for OSCAR observations (a) GODAS (b) Hybrid-GODAS with no DA (c) Hybrid-GODAS with only 3DVAR (d) and Hybrid-GODAS with EnKF/3DVar (e).

4.2.3 Comparison of O-F RMSD/bias

The remaining system evaluations are performed using the observation minus forecast (O-F) statistics and compare only the various configurations of the Hybrid-GODAS runs. The RMSD for SST observations is shown in fig 4.18 over the 1 year data assimilation experiment. The forced run (exp_none.02) begins to diverge from the initial conditions, but stays relatively constrained by the atmospheric surface forcing, reaching a peak error of slightly over 1° C in the northern hemisphere summer. The summer time is typically when the highest surface RMSD occurs. The 3DVAR (exp_var.02) and hybrid DA (exp_hyb.02) are substantially better than the freerun, although there is not much difference between 3DVAR and hybrid RMSD in the northern hemisphere and tropics, although hybrid performs slightly better. The hybrid DA does perform quite a bit better than 3DVAR in the southern hemisphere, reducing the SST RMSD from just under 0.6C to 0.4C.

The global SST O-F bias is shown in fig 4.19. An earlier configuration of the Hybrid-GODAS using DFS5.2 climatology corrections on all atmospheric variables (denoted exp.none.01) shows that such a configuration resulted in overly cold model SST that the observations were constantly trying to warm. By removing the climatology correction on 2m temperature and humidity (denoted exp.non.02) the global SST bias still exhibits small seasonal variations, but is now centered around 0°C. The hybrid data assimilation results in SST bias very close to 0°C.

Vertical profiles averaged over the 1 year experiment are shown for temperature (fig 4.20) and salinity (fig 4.21), showing that the hybrid DA has a marked

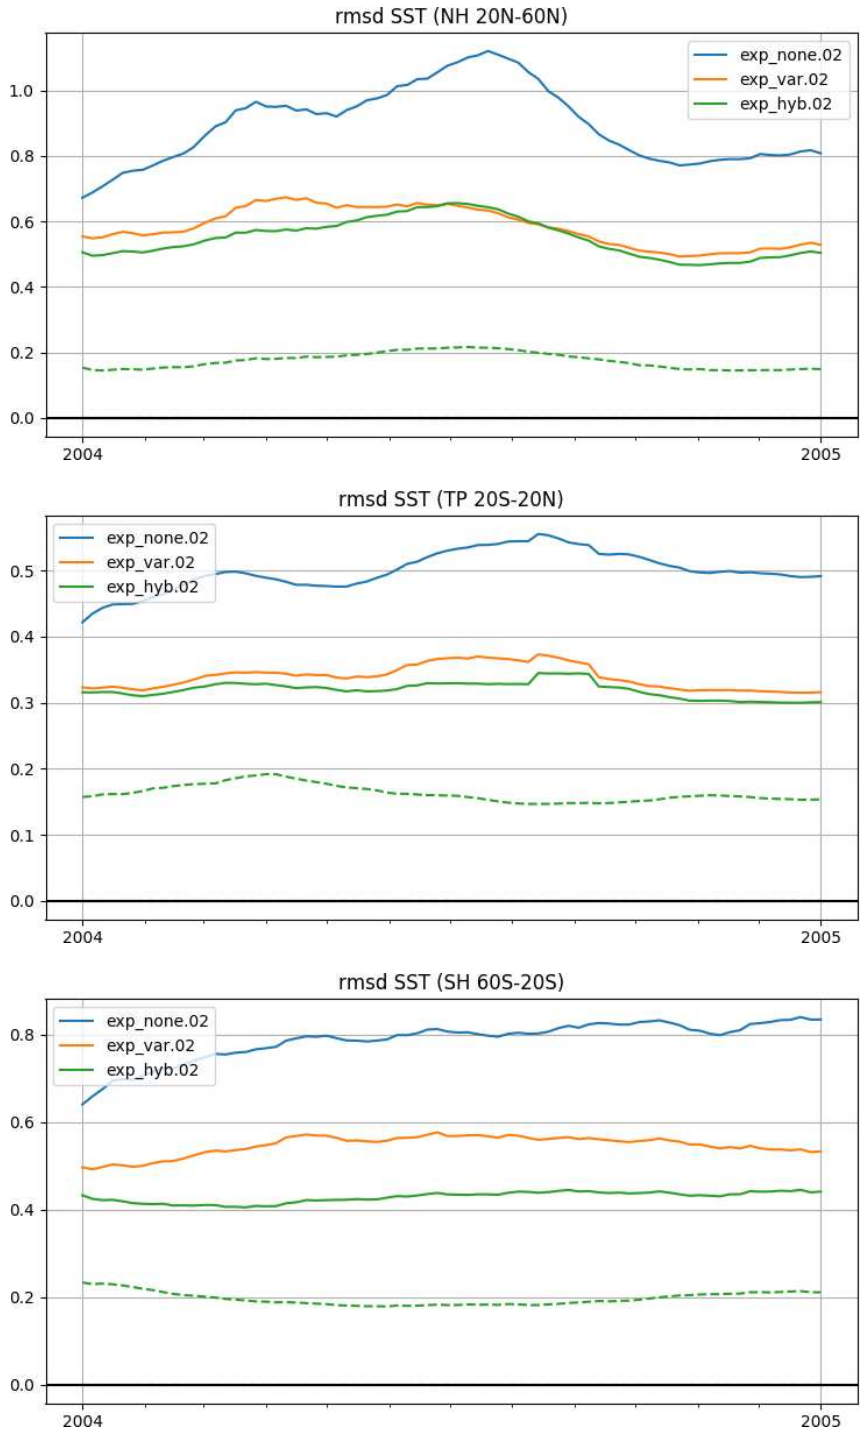

Figure 4.18: Hybrid-GODAS SST O-F RMSD averaged over northern hemisphere (NH), tropics (TP), and southern hemisphere (SH). Shown are forced run with no data assimilation (blue), 3DVar only (orange) and hybrid EnKF/3DVar (green). Ensemble spread is shown in dashed line. A 1 month moving average applied to smooth data.

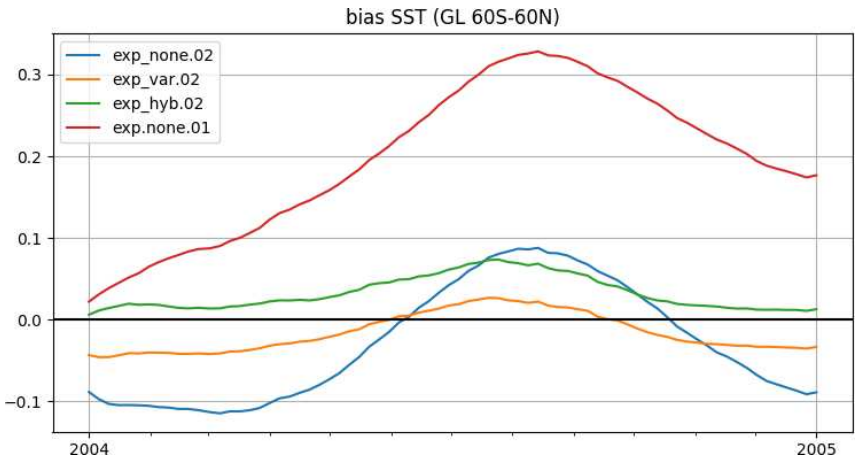

Figure 4.19: Hybrid-GODAS SST O-F bias averaged over the globe between 60S and 60N. Shown are forced run with no data assimilation (blue), 3DVar only (orange) and hybrid EnKF/3DVar (green). Original forced run with full DFS5.2 corrections (red) also shown. A 1 month moving average applied to smooth data.

benefit over 3DVar, except for salinity in the northern hemisphere, which is likely due to the previously mentioned known issues with the system. It is again worth mentioning that all three experiments were initialized from the same Jan 1, 2004 initial conditions, which would likely mean that the 3DVar and no data assimilation runs would appear worse compared with the hybrid DA once the experiments are run for longer.

4.2.4 Deficiencies

The results shown are only the second iteration of hybrid DA comparison experiments integrated over a full year, and so the system is still being tuned for optimal performance as issues are found. Several areas of the system under performing have already been mentioned, but the most striking deficiency is currently with the salinity at depths in the tropical Atlantic as shown in fig 4.22. Salinity

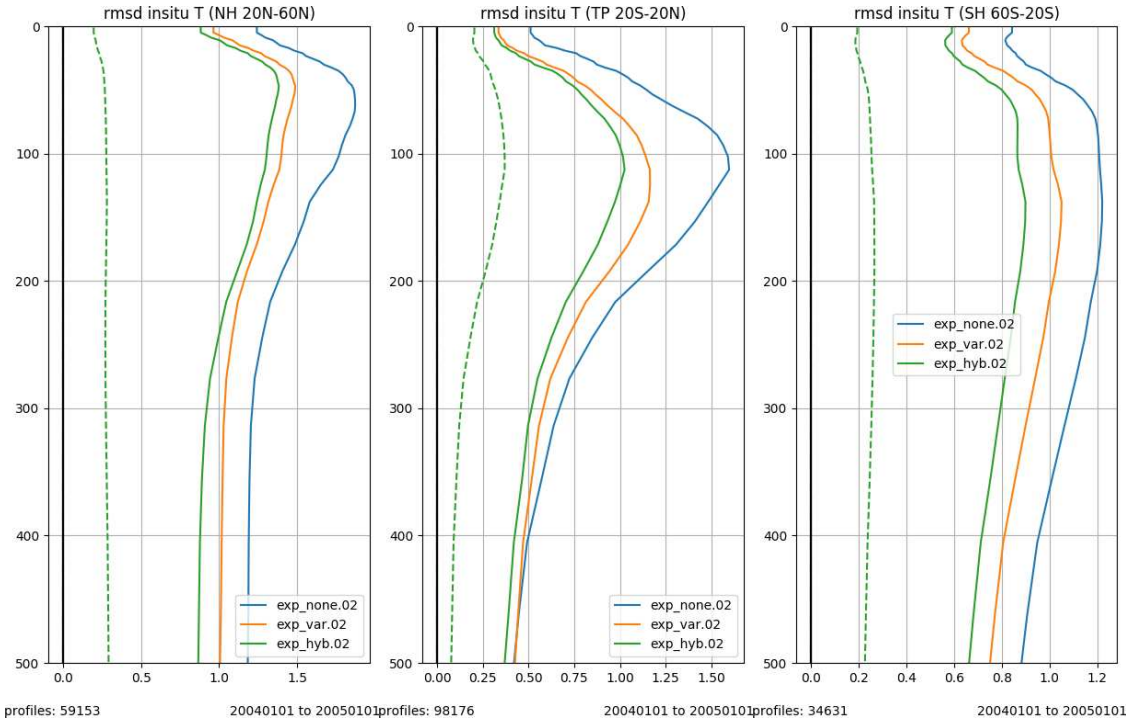

Figure 4.20: Hybrid-GODAS temperature profile O-F RMSD for the northern hemisphere (left) tropics (center) and southern hemisphere (right). Shown are forced run with no data assimilation (blue), 3DVar only (Orange) and hybrid EnKF/3DVar (green).

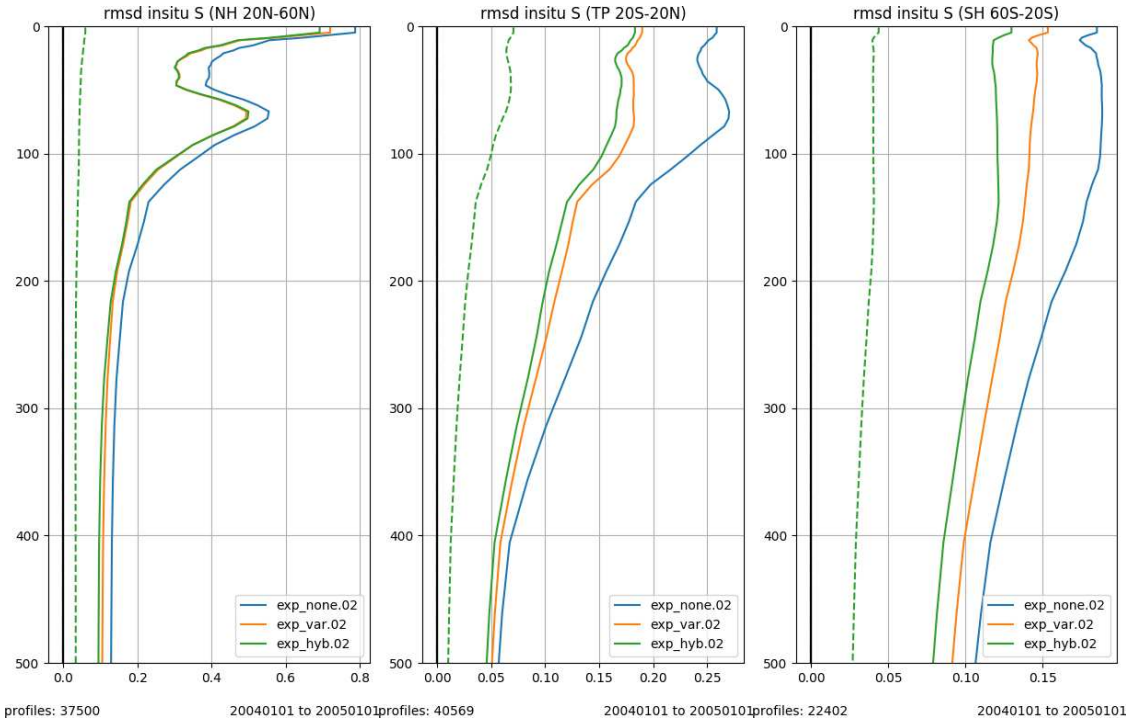

Figure 4.21: Hybrid-GODAS salinity profile O-F RMSD for the northern hemisphere (left) tropics (center) and southern hemisphere (right). Shown are forced run with no data assimilation (blue), 3DVar only (Orange) and hybrid EnKF/3DVar (green).

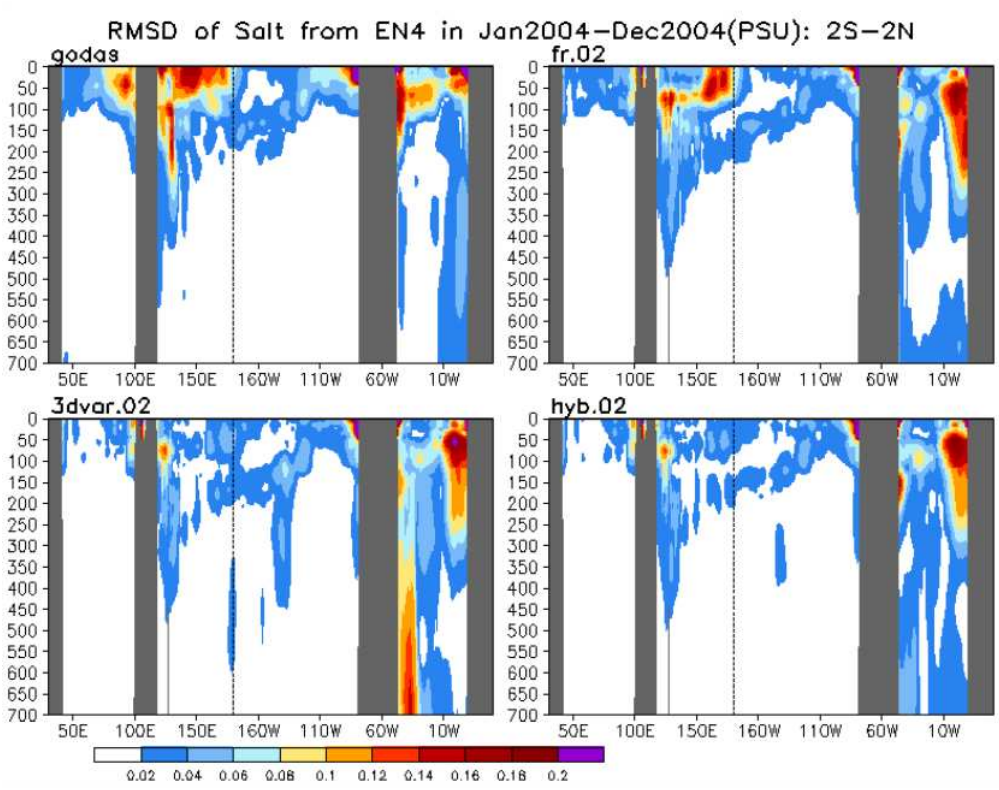

Figure 4.22: Ocean salinity RMSD as compared with EN4 averaged over 2004 along the equator between 0m and 700m depth for GODAS (top left) Hybrid-GODAS with no DA (top right), Hybrid-GODAS with only 3DVar (bottom left), and Hybrid-GODAS with EnKF/3DVar (bottom right). Large errors are introduced in the tropical Atlantic with Hybrid-GODAS

is improved most everywhere with Hybrid-GODAS, except for large errors that are introduced by the 3DVar in the western tropical Atlantic below 750m. Several bad salinity observations have been found that have slipped past the WOD quality control procedures, which is likely the cause of this problem. Custom quality control code that is currently being developed will be used before the next test experiments are started. Results are also expected to be better, especially near the boundary currents, once altimetry observations are being assimilated.

4.3 Operational plans

As has been shown here, the Hybrid-GODAS is already superior to the current operational GODAS in several major areas (salinity, SST, surface currents). There are however some areas (upper 300m heat content in the Atlantic) that are underperforming with respect to the operational GODAS. These deficiencies are expected to be addressed as the system continues to be tuned, and known issues (such as lack of LETKF vertical localization for satellite observations, and observation quality control) are fixed.

Immediate plans for Hybrid-GODAS are to prepare it for implementation into operations with a target date of late 2019. Several changes will be made to accommodate this, including replacing the SIS2 ice model with CICE5 to maintain an identical model configuration to that which will be used for the upcoming NCEP coupled model. Profile observation sources will be switched from the World Ocean Database to NCEP's in house data tanks. If time permits, the observation operators used by Hybrid-GODAS will be replaced by the JEDI unified forward operators (UFO). Several benefits will be had by switching to UFO, including having an online observation operator, meaning all observation operators will be performed at their observed time and at the closest model time step. Currently daily averages are used as input to the observation operator, having an at time observation operator from UFO should provide improvements near the surface where the diurnal cycle is important, and could allow for the assimilation of not only night time SST data, but also daytime SST. Last, the source of atmospheric perturbations will be switched to

a different source that is run in realtime at NCEP, as the 20CRv2 is not available in real-time.

But how does this all fit into the goal of strongly coupled data assimilation? Hybrid-GODAS is expected to form the basis of work at NCEP for marine data assimilation. This includes not only the ocean, but wave and ice DA. All three of these domains are expected to use the UMD-LETKF, which has been designed with strongly coupled DA in mind. In this way, after independent development on the data assimilation for these components has been completed. Strongly coupled configurations can be 'switched on' in the UMD-LETKF for experimentation with SCDA.

Chapter 5: Conclusion

In summary, strongly coupled data assimilation with the LETKF has been shown to be easy to implement from a software engineering standpoint, and may soon be practical in operational settings for limited sets of observations used across domains. SCDA was able to have a small, but beneficial impact on the CFSv2 using real observations. Currently, an effective implementation of SCDA is made difficult by biases in the observations and the models. However, near-term upgrades to the operational coupled models may alleviate this. The observations that can be used in SCDA depend heavily on the timescale of the desired DA cycle, with the use of strong vertical and variable localization required for short timescales

As a result of the work presented here, several software packages have been developed and are publicly available on GitHub.

- SPEEDYNEMO-LETKF
- CFSv2-LETKF (https://github.com/UMD-AOSC/CFSv2-LETKF)
- UMD-LETKF (https://github.com/travissluka/UMD-LETKF)
- Hybrid-GODAS (https://github.com/UMD-AOSC/hybrid-godas)

5.1 Unified Multi-Domain LETKF (UMD-LETKF)

While developing the strongly coupled ocean-atmosphere data assimilation systems for both the SPEEDY-NEMO, and the CFSv2, it became apparent that the current structure of the local ensemble transform Kalman filter code was problematic. Typically, when someone was creating an LETKF for a new model, they would have to use an existing LETKF for a similar model and then go through and replace sections of hard-coded logic. This is somewhat tedious and prone to error. If an improvement is then made to the LETKF by someone else (for example updating to include a different inflation scheme), the user would then have to manually find a way to include this improvement into their own code.

The Ocean-LETKF by Steve Penny mitigates this to some degree by pulling out code that is common to multiple ocean models. However, for a strongly coupled DA system, it would still be beneficial to have a base LETKF that can be used by any domain (ice, ocean, land, atmosphere, etc.) and any specific changes required by the domain's model are kept completely separately.

Most of the code for the LETKF is not in the core algorithm itself (which only take up at most 100 lines of code). Instead, most of the code is in the support for the core LETKF algorithms (localization, reading/writing observations and the model state, distributing the state and observations in an intelligent and scalable manner).

As a result of the CFSv2-LETKF, SPEEDYNEMO-LETKF, and MOM6-LETKF development, I have developed a completely independent Universal Multi-Domain

LETKF library (UMD-LETKF) that should be able to be used as-is for nearly any geophysical model (Sluka, 2018b). The only assumption made by the code is that the domain of interest is represented by a latitude/longitude grid, with one or more state variables of arbitrary vertical structure. The "geophysical grid" assumption would unfortunately make the library as it currently stands more difficult to use with models such as the Lorenz95. For most cases, all configuration can be done through configuration files, and the provided LETKF driver can be used, resulting in the user not needing to touch any code for the LETKF to port to a new model. If the user desires changes to the code, they simply need to use the LETKF library and provide function callbacks for the places where they wish to modify the behavior of the code.

A summary of the design choices made when creating the UMD-LETKF are as follows:

- **model independent library** - provides a single LETKF library that is compiled once and can be used for all systems. Eliminates redundancies in code. Most specialization for a given domain is done though configuration file, a generic driver is provided for simple use cases, and if the user is willing to constraints such as having all I/O be in NetCDF format. A custom driver can easily be built to interface with the library with minimal code required.

- **object oriented design** - classes for observation I/O, state I/O, and localization behavior have a default implementation that are capable of providing for most use cases, but can be overridden with user specified code if required.

- **multi-model strong coupling** - The code should allow for easy transition from weakly coupled to strongly coupled DA with no changes required by the code, everything is done within the configuration files.

Special attention has been paid to improving memory and computational efficiency in the MPI code, as well as a complete generalization of model state definition and state / observation I/O.

The code, and further documentation of it, is publicly available on github at https://github.com/travissluka/UMD-LETKF.

5.2 Possible future work

SPEEDYNEMO was originally designed for climate length runs and has proved less than ideal for the 6 hour cycling runs I performed. However, given its low computational cost, it may be very useful in examining the impacts of strongly coupled DA on climate timescales. For example, Tardif et al. (2014) found with a simple model that the Atlantic meridional overturning circulation cant be initialized correctly using only ocean observations. Similar results may be found with this and other longer timescale phenomena such as the Pacific decadal oscillation and Atlantic multidecadal oscillation. SPEEDYNEMO has been used to examine the Atlantic forcing of Pacific decadal variability (Kucharski et al., 2015). It is not known if such decadal variability results spontaneously from the model, such as the case with ENSO where flux corrections had to be performed in order to get an ENSO signal to appear. It would be a good project to see if SPEEDYNEMO produces semi-realistic decadal variability, and if so use the SPEEDYNEMO-LETKF to examine the role of strongly vs weakly coupled DA on the model's performance.

The CFSv2-LETKF experiments with real observations only used conventional observations due to the added difficulty of using radiances. The Community Radiative Transfer Model (CRTM) is being integrated with the GFS-LETKF by others and should be available for use shortly. This will allow an investigation of the effects of strongly coupled DA on the assimilation of radiances. Since the ocean sea surface temperature is required by the CRTM when computing radiances from atmospheric temperatures, and conversely, corrections due to atmospheric conditions need to be

Figure 5.1: Coupled Earth system data assimilation

considered when using infrared brightness temperatures to examine SST, it is an inherently coupled observations, and so strongly coupled DA could provide an extra benefit in this case.

The successor to the CFSv2 being developed by NCEP will be using an EnKF of some flavor for all components of the data assimilation system. More specifically, the land, ice, wave, and ocean components will be using the LETKF. This provides an ample opportunity to test the concepts of strongly coupled DA on this future

model version.

Appendix A: Observation-space 3DVar for Hybrid-GODAS

The development of the data assimilation system for Hybrid-GODAS depends on the combination of two systems, the LETKF and an ocean 3DVAR. Originally, this task was to be accomplished using the existing GODAS 3DVAR. However, the existing system was quite slow and did not scale well due to the use of a diffusion operator to model the horizontal correlations and from the algorithm's state-space formulation. It was expected that the speed of the 3DVAR would become a bottleneck at the planned 1/4 degree model resolution. Therefore, a new 3DVAR was created by combining concepts from the operational ocean data assimilation systems of the Navy (Cummings and Smedstad, 2013), UK MetOffice (Waters et al., 2015), and ECMWF (Mogensen et al., 2012).

Hybrid-GODAS is designed expecting that the LETKF performs most of the heavy lifting. The accompanying 3DVAR is there mainly to correct any biases that the LETKF cannot handle on its own. A lightweight, fast, 3DVAR is therefore the target, and while on its own its performance might not match that of the other operational centers, it should perform well when combined with the LETKF. This new ocean 3DVAR is seen as temporary, ultimately it will be replaced with the marine data assimilation being developed under the JEDI for NCEP.

State-Space Formulation

The cost function which is to be minimized is given by

$$J = \frac{1}{2}\left(\mathbf{y} - H\mathbf{x}\right)^T \mathbf{R}^{-1}\left(\mathbf{y} - H\mathbf{x}\right) + \left(\mathbf{x} - \mathbf{x}_b\right)^T \mathbf{B}^{-1}\left(\mathbf{x} - \mathbf{x}_b\right) \tag{A.1}$$

where $\mathbf{x}$ is the resulting analysis, $\mathbf{y}$ is the observations, H the observation operator to take state-space to observation-space, and $\mathbf{B}$ and $\mathbf{R}$ are the background and observation error covariances, respectively. Typically, this is solved in an incremental 3DVAR with a state-space formulation

$$\left[\mathbf{B}^{-1} + \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H}\right]\delta\mathbf{x} = \mathbf{H}^T\mathbf{R}^{-1}\delta\mathbf{y} \tag{A.2}$$

where $\delta\mathbf{x}$ is the analysis increment that is iteratively solved for through some type of minimization algorithm, and $\delta\mathbf{y} = \mathbf{y}^o - H\left(\mathbf{x}^b\right)$ are the observation increments.

When solved using a preconditioned conjugate gradient decent method, the $\mathbf{B}$ matrix is used as the preconditioner, eliminating the need to explicitly calculate $\mathbf{B}^{-1}$. The matrix to effectively be inverted by the iterative solver is $\mathbf{A} = \mathbf{H}^T\mathbf{R}^{-1}\mathbf{H} + \mathbf{B}^{-1}$. This is the method that was used in the existing GODAS 3DVAR.

A.1 Observation-Space Formulation

Using the Sherman-Morrison-Woodbury formula, Equation A.2 can be rearranged to obtain the following observation-space formulation

$$\delta\mathbf{x} = \mathbf{B}\mathbf{H}^T \left[\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R}\right]^{-1} \delta\mathbf{y} \tag{A.3}$$

This formulation is known otherwise as PSAS (Physical Space Assimilation System) at NASA/Goddard (Cohn et al., 1998). The matrix to effectively be inverted by the iterative solver is $\mathbf{A} = \mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R}$. If the number of observations is less than the number of grid points, which is easily the case with the ocean, the $\mathbf{A}$ matrix is smaller with this formulation, and therefore computationally faster. Satellite based observations can be thinned or superobbed in order to maintain the assumption that there are fewer observations than model grid points.

The observation-space formulation can be divided into two steps, the calculation of an intermediate vector $\mathbf{z}$

$$\left[\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R}\right] \mathbf{z} = \delta\mathbf{y} \tag{A.4}$$

followed by the projection of $\mathbf{z}$ vector back into state-space

$$\delta\mathbf{x} = \mathbf{B}\mathbf{H}^T\mathbf{z} \tag{A.5}$$

the computational expense of calculating $\mathbf{z}$ by an iterative algorithm depends on the

number of observations. So, for the ocean, the minimization here should be faster than with state-space. The subsequent matrix multiplication required project $\mathbf{z}$ onto the model-space is more computationally expensive, but is only required once, not iteratively.

Preconditioned Conjugate Gradient Solver

An outline of the preconditioned conjugate gradient algorithm used to solve the observation-space formulation follows, with $\mathbf{A}^{*^{-1}}$ defined as an approximate inverse to $\mathbf{A}$, and $\mathbf{A} = \left[\mathbf{HBH}^T + \mathbf{R}\right]$. The following variables are first initialized:

$$\begin{aligned} \mathbf{z}_0 &= 0 \\ \mathbf{r}_0 &= \delta\mathbf{y} \\ \mathbf{s}_0 &= \mathbf{A}^{*^{-1}}\mathbf{r}_0 \\ \mathbf{p}_1 &= \mathbf{s}_0 \end{aligned} \tag{A.6}$$

The iterative solver is run until the solution converges, usually when the residual decreases by two orders of magnitude. At each step:

if $k > 1$:

$$\begin{aligned} \beta_k &= \frac{\mathbf{r}_{k-1}^T \mathbf{s}_{k-1}}{\mathbf{r}_{k-2}^T \mathbf{s}_{k-2}} \\ \mathbf{p}_k &= \mathbf{s}_{k-1} + \beta_k \mathbf{p}_{k-1} \end{aligned} \tag{A.7}$$

then for each iteration with $k > 0$:

$$\begin{aligned}
\mathbf{q}_k &= \mathbf{A}\mathbf{p}_k \\
\alpha_k &= \frac{\mathbf{s}_{k-1}^T \mathbf{r}_{k-1}}{\mathbf{p}_k^T \mathbf{q}_k} \\
\mathbf{z}_k &= \mathbf{z}_{k-1} + \alpha_k \mathbf{p}_k \\
\mathbf{r}_k &= \mathbf{r}_{k-1} + \alpha_k \mathbf{q}_k \\
\mathbf{s}_k &= \mathbf{A}^{*^{-1}} \mathbf{r}_k
\end{aligned} \tag{A.8}$$

it should be noted that in the actual code, only the most recent vectors for $\mathbf{r}_k$ and $\mathbf{s}_k$ are stored, and so only the final dot products of $\mathbf{r}^T\mathbf{s}$ are kept from two previous steps for calculating β_k.

The preconditioning matrix $\mathbf{A}^{*^{-1}}$ is calculated by a block diagonal approximation to the full $\mathbf{A}^{-1}$ matrix. This is done by dividing observations into blocks so that each block contains a reasonably small number of observations (on the order of 1000). Then a Cholesky decomposition is performed on each block in parallel to directly invert the matrix. The use of this preconditioning step speeds up convergence of the solver substantially and is relatively inexpensive to perform as $\mathbf{A}^{*^{-1}}$ is calculated once at the beginning of the solver and stored for subsequent applications.

A.2 Background Error Covariance Model

For the observation-space formulation of the 3DVAR, the background error covariance between two observation locations (**HBH**), and the background error covariance between an observation location and a model grid location (**BH**) are

required. These are nearly identical functions, however, there are small differences in the code for performance reasons, which will not be discussed here.

The covariance is decomposed into a background error variance and correlation. The correlations are further separated into the horizontal, vertical, and variable correlations, as well as an SSH gradient tensor and coast distance tensor to modulate the horizontal correlation. Since the background error covariance model used here is rather simple, there is no balance operator between temperature and salinity, and so the covariance between observations and model variables of different types is zero.

All distance based correlations are calculated using a compact spline (Gaspari and Cohn, 1999) given the distance between two points, d, and the desired length scale, L, (equivalent to 1 standard deviation of a Gaussian). Using $r = d/(L\sqrt{10/3})$, a distance based correlation value is given by:

$$\begin{aligned} gc(r) &= -r^5/4 + r^4/2 + 5r^3/8 - 5r^2/3 + 1, && 0 \leq r \leq 1, \\ &= r^5/12 - r^4/2 + 5r^3/8 + 5r^2/3 - 5r + 4 - 2/3r, && 1 \leq r \leq 2, \\ &= 0, && r > 2 \end{aligned} \tag{A.9}$$

Horizontal Correlation

The correlation lengths in the horizontal are calculated as an anisotropic 2D Gaussian approximated by the Gaspari-Cohn function. These lengths vary with latitude as a function of the Rossby radius. The correlation length in the meridional

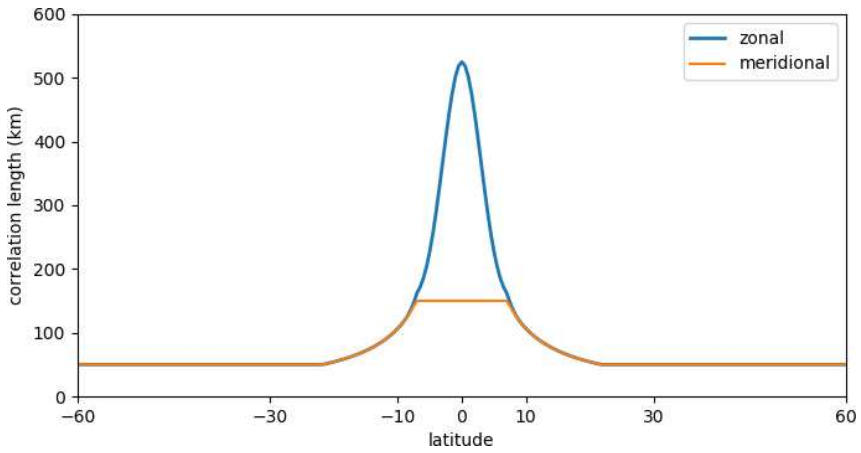

Figure A.1: The zonal and meridional horizontal correlation length scales used in the 3DVAR.

direction is given by

$$L_{\phi}(\phi) = \max\left(\min\left(\frac{c}{2\Omega\left|\sin(\phi)\right|}, L_{\phi_{max}}\right), L_{\phi_{min}}\right) \tag{A.10}$$

where $c = 2.7$ m/s is the assumed gravity wave speed. The value is clamped to a minimum of 50 km in the high latitudes, and a maximum of 150km along the equator. The correlation length in the zonal direction is equal to L_{ϕ} in the extra-tropics, but is stretched longer near the equator within 10°to a maximum of 525km as given by

$$L_{\lambda}(\phi) = L_{\phi}(\phi)\left(1 + gc\left(2\phi/\phi_{L}\right)\left(\delta - 1.0\right)\right) \tag{A.11}$$

The value of the parameters used are given in Table A.1 and the resulting horizontal correlation length scales are shown in Figure A.1.

$L_{\phi_{min}}$	50 km
$L_{\phi_{max}}$	150 km
ϕ_L	10°
δ	3.5
c	2.7 m/s
Ω	7.29×10^{-5}

Table A.1: horizontal correlation length scale parameters

Coast Distance and SSH Gradient Tensor

In order to introduce an element of crude flow dependence into the background error covariance model, the horizontal correlations are modulated by two gradient tensors based on 1) the distance from the coast and 2) the SSH field. This method is borrowed from Cummings and Smedstad (2013).

The horizontal correlation is modified by multiplying by the following calculation

$$C_{coast} = \max\left(C_{coast_{min}}, 1 - \frac{|r_1 - r_2|}{r_c}\right) \tag{A.12}$$

where r_c is the distance from the coast at which this effect begins, r_1 and r_2 are the distances to the coast for the two grid points of interest (clamped to a maximum value of r_c). $C_{coast_{min}}$ is the strength of the effect, 1.0 being off and 0.0 forcing grid points on the coast to be completely uncorrelated with those points r_c away from the coast. Hybrid-GODAS uses $r_c = 75km$, $C_{coast_{min}} = 0.3$. For an observation that is next to the coast, this has the effect of spreading the analysis increment along the coastline, and not out into the open ocean.

An additional similar modulation is performed to the horizontal correlations

based on the difference in sea surface height of the two points. In this way, analysis increments are spread along fronts instead of across them, a feature that is particularly important along the western boundary currents, as can be shown in Figure A.2. Any surface field can be used, SSH, SST, etc, but the benefit of this surface gradient tensor relies on the surface fields being represented relatively realistically. For Hybrid-GODAS, SSH is used, with a strength of 10cm. As can be seen in Figure A.2, increments are stretched out along the Kuroshio Current, and not across it. This method is also useful for preventing analysis increments across bodies of water separated by land, such as between the Pacific and Atlantic oceans near Panama, since the SSH is different enough between the two locations to isolate analysis increments.

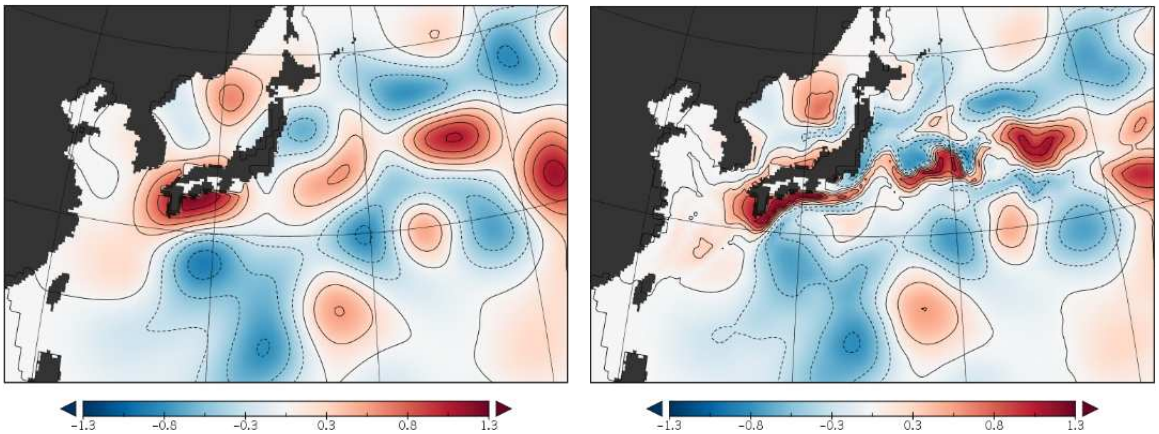

Figure A.2: Surface temperature analysis increment for a single cycle with normal horizontal correlations (a) and horizontal correlations modulated by an SSH gradient tensor (b)

Vertical Correlation

Three dimensional vertical correlation lengths were originally calculated by scaling the vertical density gradient, as done in Cummings and Smedstad (2013). However, if not done very carefully and with ample smoothing of the resulting field, this would lead to instances where $\mathbf{B}$ was not positive definite, and the Cholesky decomposition for the preconditioner would fail.

Instead, the vertical correlation length is first set equal to the model level thickness. The top ocean model level has its vertical correlation length set equal to the mixed layer depth (MLD) defined as the depth where a reduction in density of 0.125 kg/m^3 occurs. If this depth is deeper than 250m it is clamped to 250m. Between the top model level and the base of the mixed layer, the vertical correlation length is linearly interpolated between the values at the two locations. In this way, satellite SST observations are able to impact the entire mixed layer, but below the mixed layer observation profiles impact primarily only the model level in which they occur.

The 3D vertical correlation length field is then smoothed using a recursive filter using the horizontal and vertical correlation lengths that were found. This helps ensure stability of the solver, as $\mathbf{B}$ can fail to be positive definite if the correlation lengths vary spatially at a rate faster than their own correlation length. A cross section along the equation for an example date is shown in Figure A.3.

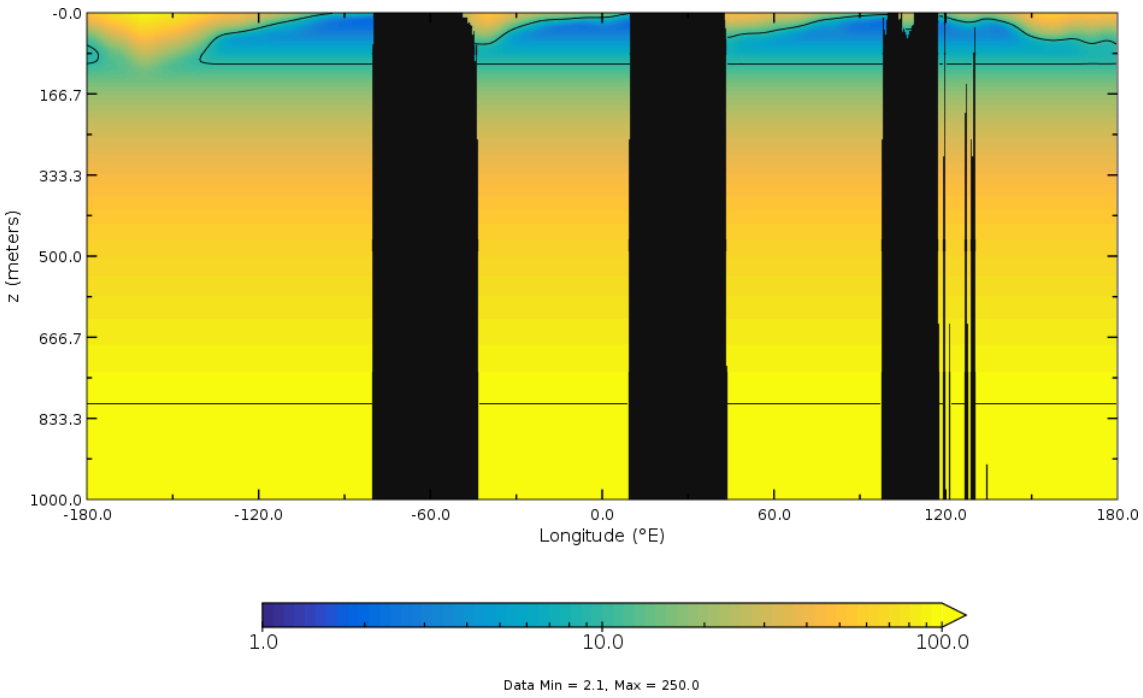

Figure A.3: Cross section of vertical correlation lengths along equator a single example date. Vertical correlation length is equal to model level thickness, except when within the mixed layer.

Background Error Variance

The method of calculating the 3D background error variance for the temperature and salinity was adopted from the 1/4 degree NEMOVAR data assimilation system used at the UK MetOffice (Waters et al., 2015), and is similar in concept to the method used for the original GODAS. First, the vertical temperature gradient is calculated and multiplied by a scaling factor δz

$$\sigma_{vg} = \delta z \left| \frac{dT}{dz} \right| \tag{A.13}$$

The resulting 3D field is smoothed by the previously calculated horizontal and vertical correlation lengths. The final standard deviations of the background error variance are given by

$$\sigma = \max\left(\sigma_{min}, \min\left(\sigma_{vg}, \sigma_{max}\right)\right) \tag{A.14}$$

where σ_{max} is a defined constant, and σ_{min} is calculated from

$$\sigma_{min} = \sigma_{do} + \left(\sigma_{surf} - \sigma_{do}\right) \exp\left[\frac{d\left(1\right) - d\left(k\right)}{L}\right] \tag{A.15}$$

The minimum background error varies vertically with a maximum at the surface of σ_{surf} decreasing exponentially with a length scale of L to a minimum of σ_{do}. The value for σ_{surf} for salinity is a fixed constant, however for temperature a 2D field is generated based on a scaled and clamped climatological average of the O-F RMSD from the satellite SST observations of a previous 3DVar run (Figure A.4).

The other constant parameters used in the calculation of the background error variance for temperature and salinity are given by Table A.2.

	Temp	Salt
L	500m	250m
σ_{max}	1.8	0.25
σ_{do}	0.1	0.02
σ_{surf}	2D field	0.3
δz	20	2.5

Table A.2: Parameters used by 3DVar background variance calculation

The 2D σ_{surf} field for temperature is needed in order to increase the background error variance where the vertical temperature gradient does not capture all

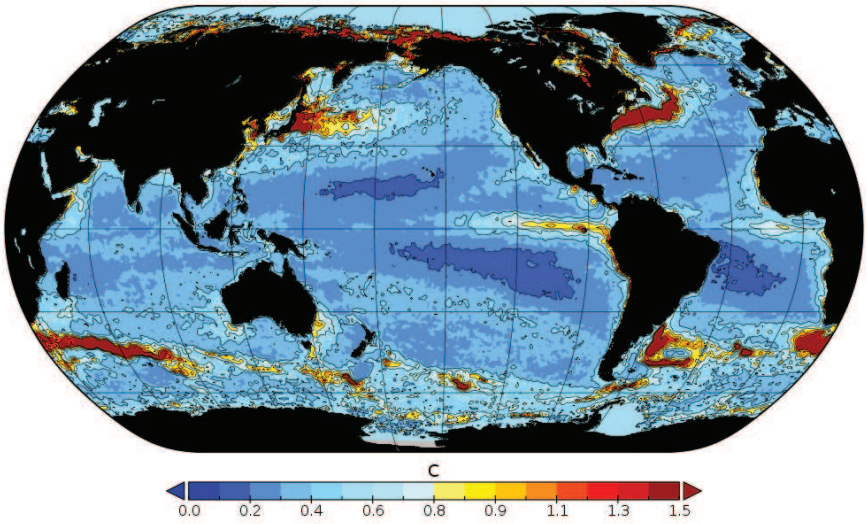

Figure A.4: The minimum surface temperature background error covariance used in the Hybrid-GODAS 3DVar. Calculated from annual average of satellite SST O-F RMSD.

the error, such as near the western boundary currents. This effect can be seen in Figure A.5.

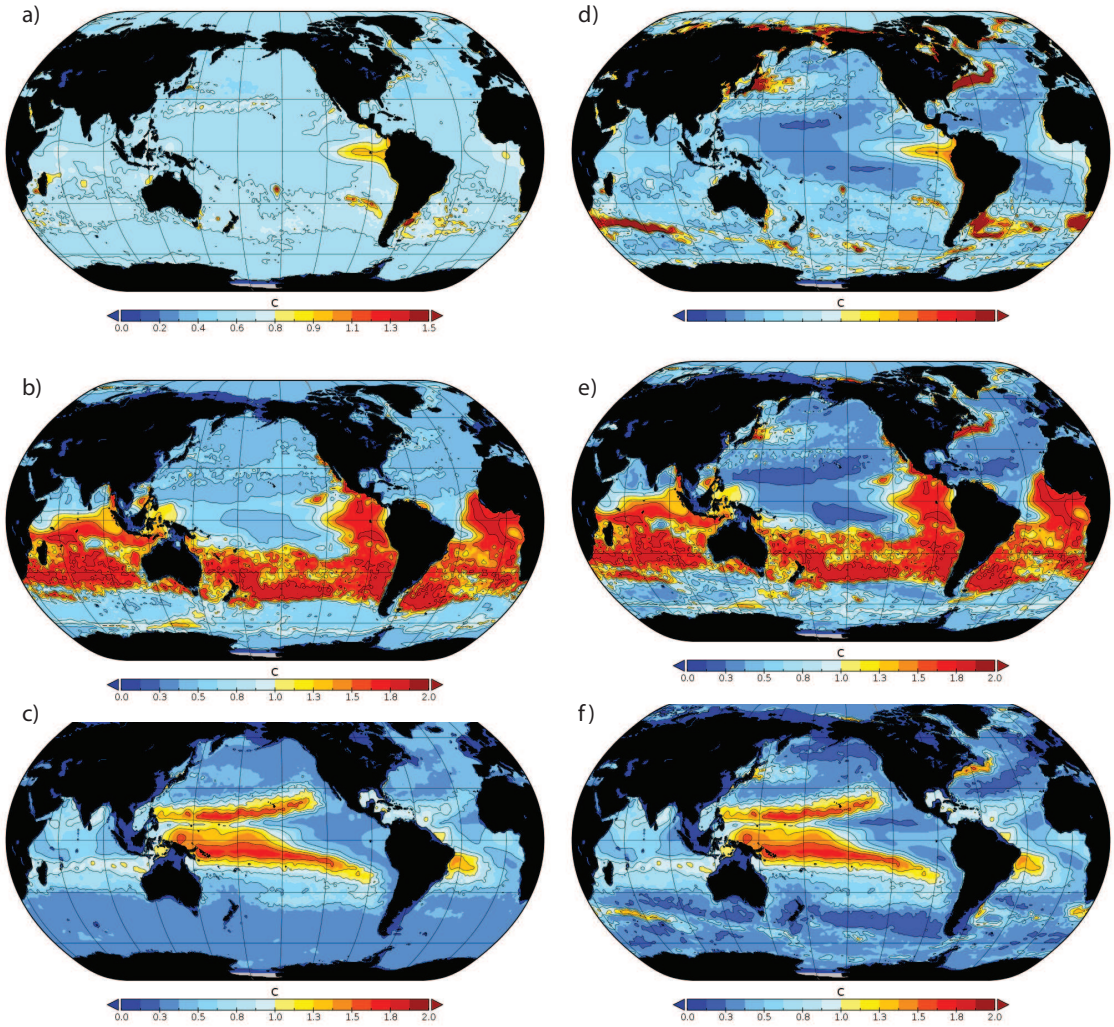

Figure A.5: The 3DVar background error variance for 2004-04-01 at depths of 1 meter (top), 50 meters (middle), and 200 meters (bottom) using minimum surface background error variance of 0.5C as done by GODAS (left) and a 2D varying modified background error variance as in Hybrid-GODAS (right).

Bibliography

Anderson, J. L. (2001). An Ensemble Adjustment Kalman Filter for Data Assimilation. *Monthly Weather Review*, 129(12):2884–2903.

Balmaseda, M. and Anderson, D. (2009). Impact of initialization strategies and observations on seasonal forecast skill. *Geophysical Research Letters*, 36(1):L01701.

Bannister, R. N. (2017). A review of operational methods of variational and ensemble-variational data assimilation. *Quarterly Journal of the Royal Meteorological Society*, 143(703):607–633.

Battisti, D. S. and Hirst, A. C. (1989). Interannual Variability in a Tropical AtmosphereOcean Model: Influence of the Basic State, Ocean Geometry and Nonlinearity.

Behringer, D. (2007). The Global Ocean Data Assimilation System (GODAS) at NCEP.

Behringer, D. and Xue, Y. (2004). Evaluation of the Global Ocean Data Assimilation System at NCEP: The Pacific Ocean. In *Eighth Symposium on Integrated Observing and Assimilation Systems for Atmosphere, Oceans, and Land Surface, AMS 84th Annual Meeting*, number January, pages 11–15.

Berner, J., Shutts, G. J., Leutbecher, M., and Palmer, T. N. (2009). A Spectral Stochastic Kinetic Energy Backscatter Scheme and Its Impact on Flow-Dependent Predictability in the ECMWF Ensemble Prediction System. *Journal of the Atmospheric Sciences*, 66(3):603–626.

Bonavita, M., Isaksen, L., and Hólm, E. (2012). On the use of EDA background error variances in the ECMWF 4D-Var. *Quarterly Journal of the Royal Meteorological Society*, 138(667):1540–1559.

Bonjean, F., Lagerloef, G. S. E., Bonjean, F., and Lagerloef, G. S. E. (2002). Diagnostic Model and Analysis of the Surface Currents in the Tropical Pacific Ocean. *Journal of Physical Oceanography*, 32(10):2938–2954.

Boyer, T. P., Antonov, J. I., Baranova, O. K., Coleman, C., Garcia, H. E., Grodsky, A., Johnson, D. R., Locarnini, R. a., Mishonov, A. V., O'Brien, T. D., Paver, C. R., Reagan, J. R., Seidov, D., Smolyar, I. V., Zweng, M. M., Brien, T. D. O., Paver, C. R., Reagan, J. R., Seidov, D., Smolyar, I. V., and Zweng, M. M. (2013). World Ocean Database 2013.

Brasnett, B. (2008). The impact of satellite retrievals in a global sea-surface-temperature analysis. *Q.J.R. Meteorol. Soc.*, 134:1745–1760.

Casey, K. S., Brandon, T. B., Cornillon, P., and Evans, R. (2010). The Past, Present, and Future of the AVHRR Pathfinder SST Program. In Barale, V., Gower, J. F. R., and Alberotanza, L., editors, *Oceanography from Space: Revisited*, pages 273–287. Springer Netherlands, Dordrecht.

Cohn, S. E., da Silva, A., Guo, J., Sienkiewicz, M., and Lamich, D. (1998). Assessing the Effects of Data Selection with the DAO Physical-Space Statistical Analysis System*. *Monthly Weather Review*, 126(11):2913–2926.

Compo, G. P., Whitaker, J. S., and Sardeshmukh, P. D. (2006). Feasibility of a 100-year reanalysis using only surface pressure data. *Bulletin of the American Meteorological Society*, 87(2):175–190.

Conkright, M. E. and Coauthors (1999). World Ocean Database 1998, documentation and quality control version 2.0.

Cravatte, S., Ganachaud, A., Dewitte, B., and Hernandez, F. (2015). TPOS2020: Tropical Pacific observing system for 2020.

Cummings, J. A. and Smedstad, O. M. (2013). Variational Data Assimilation for the Global Ocean. In Park, S. K. and Xu, L., editors, *Data Assimilation for Atmospheric, Oceanic and Hydrologic Applications (Vol. II)*, volume II, pages 303–343. Springer Berlin Heidelberg, Berlin, Heidelberg.

Dai, A. (2016). Historical and Future Changes in Streamflow and Continental Runoff: A Review. *Terrestrial Water Cycle and Climate Change: Natural and Human-Induced Impacts*, pages 17–37.

Daley, R. and Barker, E. (2001). *NAVDAS Source Book 2001.* NRL.

Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M. A., Balsamo, G., Bauer, P., Bechtold, P., Beljaars, A. C., van de Berg, L., Bidlot, J., Bormann, N., Delsol, C., Dragani, R., Fuentes, M., Geer, A. J., Haimberger, L., Healy, S. B., Hersbach, H., Hólm, E. V., Isaksen, L., Kållberg, P., Köhler, M., Matricardi, M., Mcnally, A. P., Monge-Sanz, B. M., Morcrette, J. J., Park, B. K., Peubey, C., de Rosnay, P., Tavolato, C., Thépaut, J. N., and Vitart, F. (2011). The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the Royal Meteorological Society*, 137(656):553–597.

Derber, J. and Rosati, A. (1989). A Global Oceanic Data Assimilation System. *Journal of Physical Oceanography*, 19(9):1333–1347.

Dussin, R., Barnier, B., Brodeau, L., and Molines, J. M. (2016). The Making Of the DRAKKAR FORCING SET DFS5. *DRAKKAR/MyOcean Report 01-04-16*, April:1–34.

ECMWF (2016). ECMWF's strategy 20162025: The strength of a common goal. Technical report, ECMWF.

Ek, M. B. (2003). Implementation of Noah land surface model advances in the National Centers for Environmental Prediction operational mesoscale Eta model. *Journal of Geophysical Research*, 108(D22):8851.

Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99(C5):10143.

Feng, X., Haines, K., and de Boisséson, E. (2018). Coupling of surface air and sea surface temperatures in the CERA-20C reanalysis. *Quarterly Journal of the Royal Meteorological Society*.

Frolov, S., Bishop, C. H., Holt, T., Cummings, J., and Kuhl, D. (2016). Facilitating Strongly Coupled OceanAtmosphere Data Assimilation with an Interface Solver. *Monthly Weather Review*, 144(1):3–20.

Fu, X., Wang, B., Waliser, D. E., and Tao, L. (2007). Impact of AtmosphereOcean Coupling on the Predictability of Monsoon Intraseasonal Oscillations*. *Journal of the Atmospheric Sciences*, 64(1):157–174.

Gaspari, G. and Cohn, S. E. (1999). Construction of correlation functions in two and three dimensions. *Quarterly Journal of the Royal Meteorological Society*, 125(April 1998):723–757.

Gibson, J., Kallberg, P., Uppala, S., Hernandez, A., Nomura, A., and Serrano, E. (1999). ECMWF Re-Analysis Project Report Series, 1. ERA-15 Description, Version 2. Technical report, European Centre for Medium-Range Weather Forecasts, Reading, UK.

Good, S. A., Martin, M. J., and Rayner, N. A. (2013). EN4: Quality controlled ocean temperature and salinity profiles and monthly objective analyses with uncertainty estimates. *Journal of Geophysical Research: Oceans*, 118(12):6704–6716.

Greybush, S. J., Kalnay, E., Miyoshi, T., Ide, K., and Hunt, B. R. (2011). Balance and Ensemble Kalman Filter Localization Techniques. *Monthly Weather Review*, 139(2):511–522.

Grissom, K., Elliott, J. E., Singelton, T., Thompson, W., and Keene, J. (2017). Adding a Real-Time Current Profile to NDBC TAO Moorings for Enhanced Ocean BoundaryLayer Observations. In *Oceans 2017*, pages 1–6.

Han, G., Wu, X., Zhang, S., Liu, Z., and Li, W. (2013). Error Covariance Estimation for Coupled Data Assimilation Using a Lorenz Atmosphere and a Simple Pycnocline Ocean Model. *Journal of Climate*, 26(24):10218–10231.

Houtekamer, P. and Mitchell, H. L. (2005). Ensemble Kalman filtering. *Quarterly Journal of the Royal Meteorological Society*, 131(613):3269–3289.

Hunt, B. R., Kostelich, E. J., and Szunyogh, I. (2007). Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter. *Physica D: Nonlinear Phenomena*, 230(1-2):112–126.

Hurrell, J. W., Holland, M. M., Gent, P. R., Ghan, S., Kay, J. E., Kushner, P. J., Lamarque, J. F., Large, W. G., Lawrence, D., Lindsay, K., Lipscomb, W. H., Long, M. C., Mahowald, N., Marsh, D. R., Neale, R. B., Rasch, P., Vavrus, S., Vertenstein, M., Bader, D., Collins, W. D., Hack, J. J., Kiehl, J., and Marshall, S. (2013). The community earth system model: A framework for collaborative research. *Bulletin of the American Meteorological Society*, 94(9):1339–1360.

Ignatov, A., Zhou, X., Petrenko, B., Liang, X., Kihai, Y., Dash, P., Stroup, J., Sapper, J., and DiGiacomo, P. (2016). AVHRR GAC SST Reanalysis Version 1 (RAN1). *Remote Sensing*, 8(4):315.

Kalnay, E. (2003). *Atmospheric modeling, data assimilation and predictability*. Cambridge university press.

Kalnay, E., Li, H., Miyoshi, T., Yang, S.-C., and Ballabrera-Poy, J. (2007). 4-D-Var or ensemble Kalman filter? *Tellus A*, 59(5):758–773.

Kanamitsu, M., Ebisuzaki, W., Woollen, J., Yang, S.-K., Hnilo, J. J., Fiorino, M., and Potter, G. L. (2002). NCEP-DOE AMIP-II Reanalysis (R-2). *Bulletin of the American Meteorological Society*, 83(11):1631–1643.

Kang, J.-S., Kalnay, E., Liu, J., Fung, I., Miyoshi, T., and Ide, K. (2011). Variable localization in an ensemble Kalman filter: Application to the carbon cycle data assimilation. *Journal of Geophysical Research*, 116(D9):D09110.

Kleist, D. T. and Ide, K. (2015). An OSSE-Based Evaluation of Hybrid VariationalEnsemble Data Assimilation for the NCEP GFS. Part I: System Description and 3D-Hybrid Results. *Monthly Weather Review*, 143(2):433–451.

Klingaman, N. P., Inness, P. M., Weller, H., and Slingo, J. M. (2008). The importance of high-frequency sea surface temperature variability to the intraseasonal oscillation of Indian monsoon rainfall. *Journal of Climate*, 21(23):6119–6140.

Kretschmer, M., Hunt, B. R., and Ott, E. (2015). Data assimilation using a climatologically augmented local ensemble transform Kalman filter. *Tellus A*, 67:1–9.

Kröger, J. and Kucharski, F. (2011). Sensitivity of ENSO characteristics to a new interactive flux correction scheme in a coupled GCM. *Climate Dynamics*, 36(1-2):119–137.

Kucharski, F., Ikram, F., Molteni, F., Farneti, R., Kang, I.-S., No, H.-H., King, M. P., Giuliani, G., and Mogensen, K. (2015). Atlantic forcing of Pacific decadal variability. *Climate Dynamics*, pages 1–15.

Kucharski, F., Molteni, F., and Bracco, A. (2006). Decadal interactions between the western tropical Pacific and the North Atlantic Oscillation. *Climate Dynamics*, 26(1):79–91.

Kucharski, F., Molteni, F., King, M. P., Farneti, R., Kang, I.-S., and Feudale, L. (2013). On the Need of Intermediate Complexity General Circulation Models: A SPEEDY Example. *Bulletin of the American Meteorological Society*, 94(1):25–30.

Laloyaux, P., Balmaseda, M., Dee, D., Mogensen, K., and Janssen, P. (2015). A coupled data assimilation system for climate reanalysis. *Quarterly Journal of the Royal Meteorological Society*.

Large, W. G. and Yeager, S. G. (2004). Diurnal to decadal global forcing for ocean and sea-ice models: The data sets and flux climatologies. *NCAR Tech. Note*, TN-460+STR(May):105pp.

Lawless, A. (2012). Coupled Model Data Assimilation. In *Report from International Workshop on Coupled Data Assimilation, University of Reading, 10-12 Sept 2012*, page 25, Reading, England.

Lea, D. J., Mirouze, I., Martin, M. J., King, R. R., Hines, A., Walters, D., and 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.

Lien, G.-y., Kalnay, E., and Miyoshi, T. (2013). Effective assimilation of global precipitation: simulation experiments. *Tellus A*, 65:1–16.

Liu, Z., Wu, S., Zhang, S., Liu, Y., and Rong, X. (2013). Ensemble data assimilation in a simple coupled climate model: The role of ocean-atmosphere interaction. *Advances in Atmospheric Sciences*, 30(5):1235–1248.

Locarnini, R. A., Mishonov, A. V., Antonov, J., Boyer, T. P., Garcia, H. E., Baranova, O. K., Zweng, M. M., Paver, C. R., Reagan, J. R., Johnson, D. R., Hamilton, M., and Seidov, D. (2013). World Ocean Atlas 2013, Volume 1: Temperature. Technical report, NOAA Atlas NESDIS 73.

Lu, F., Liu, Z., Zhang, S., and Liu, Y. (2015a). Strongly Coupled Data Assimilation Using Leading Averaged Coupled Covariance (LACC). Part I: Simple Model Study*. *Monthly Weather Review*, 143(9):3823–3837.

Lu, F., Liu, Z., Zhang, S., Liu, Y., and Jacob, R. (2015b). Strongly Coupled Data Assimilation Using Leading Averaged Coupled Covariance (LACC). Part II: CGCM Experiments*. *Monthly Weather Review*, 143(11):4645–4659.

Luo, X. and Hoteit, I. (2014). Ensemble Kalman Filtering with a Divided State-Space Strategy for Coupled Data Assimilation Problems. *Monthly Weather Review*, 142(12):4542–4558.

Maclachlan, C., Arribas, A., Peterson, K. A., Maidens, A., Fereday, D., Scaife, A. A., Gordon, M., Vellinga, M., Williams, A., Comer, R. E., Camp, J., Xavier, P., and Madec, G. (2015). Global Seasonal forecast system version 5 (GloSea5): A high-resolution seasonal forecast system. *Quarterly Journal of the Royal Meteorological Society*, 141(689):1072–1084.

Madden, R. a. and Julian, P. R. (1972). Description of Global-Scale Circulation Cells in the Tropics with a 4050 Day Period.

Madec, G. (2008). NEMO ocean engine.

Marshall, A. G., Hudson, D., Wheeler, M. C., Hendon, H. H., and Alves, O. (2011). Assessing the simulation and prediction of rainfall associated with the MJO in the POAMA seasonal forecast system. *Climate Dynamics*, 37(11-12):2129–2141.

Miyoshi, T. (2005). *Ensemble Kalman Filter Experiments with a Primitive-Equation Global Model*. Ph.d. thesis, University of Maryland.

Miyoshi, T. (2011). The Gaussian Approach to Adaptive Covariance Inflation and Its Implementation with the Local Ensemble Transform Kalman Filter. *Monthly Weather Review*, 139(5):1519–1535.

Mogensen, K., Alonso, M., and Weaver, A. (2012). The NEMOVAR ocean data assimilation system as implemented in the ECMWF ocean analysis for System 4. Technical report, CERFACS, Toulouse, France.

Molod, A., Takacs, L., Suarez, M., Bacmeister, J., Song, I.-S., and Eichmann, A. (2012). The GEOS-5 atmospheric general circulation model: Mean climate and development from MERRA to Fortuna.

Molteni, F. (2003). Atmospheric simulations using a GCM with simplified physical parametrizations. I: Model climatology and variability in multi-decadal experiments. *Climate Dynamics*, 20(2):175–191.

Mulholland, D. P., Laloyaux, P., Haines, K., and Balmaseda, M. A. (2015). Origin and Impact of Initialization Shocks in Coupled Atmosphere-Ocean Forecasts. *Monthly Weather Review*, 143(11):4631–4644.

National Academies of Sciences Engineering and Medicine (2016). Next Generation Earth System Prediction: Strategies for Subseasonal to Seasonal Forecasts. Technical report, National Academies of Sciences, Engineering, and Medicine., Washington, DC: The National Academies Press.

Ott, E., Hunt, B., Szunyogh, I., ZIMIN, A. V., KOSTELICH, E. J., CORAZZA, M., KALNAY, E., PATIL, D. J., and YORKE, J. A. (2004). A local ensemble Kalman filter for atmospheric data assimilation. *Tellus A*, 56(5):415–428.

Pena, M., Kalnay, E., and Cai, M. (2003). Statistics of locally coupled ocean and atmosphere intraseasonal anomalies in Reanalysis and AMIP data. *Nonlinear Processes in ...*, 10:245–251.

Penny, S. G. (2014). The Hybrid Local Ensemble Transform Kalman Filter. *Monthly Weather Review*, 142(6):2139–2149.

Penny, S. G., Behringer, D. W., Carton, J. a., and Kalnay, E. (2015). A Hybrid Global Ocean Data Assimilation System at NCEP. *Monthly Weather Review*, 143(11):4660–4677.

Penny, S. G. and Hamill, T. M. (2017). Coupled Data Assimilation for Integrated Earth System Analysis and Prediction. *Bulletin of the American Meteorological Society*, 98(7):ES169–ES172.

Penny, S. G., Kalnay, E., Carton, J. a., Hunt, B. R., Ide, K., Miyoshi, T., and Chepurin, G. a. (2013). The local ensemble transform Kalman filter and the running-in-place algorithm applied to a global ocean general circulation model. *Nonlinear Processes in Geophysics*, 20(6):1031–1046.

Reynolds, R. W., Rayner, N. A., Smith, T. M., Stokes, D. C., and Wang, W. (2002). An improved in situ and satellite SST analysis for climate. *Journal of Climate*, 15(13):1609–1625.

Reynolds, R. W., Smith, T. M., Liu, C., Chelton, D. B., Casey, K. S., and Schlax, M. G. (2007). Daily High-Resolution-Blended Analyses for Sea Surface Temperature. *Journal of Climate*, 20(22):5473–5496.

Roemmich, D., Johnson, G., Riser, S., Davis, R., Gilson, J., Owens, W. B., Garzoli, S., Schmid, C., and Ignaszewski, M. (2009). The Argo Program: Observing the Global Oceans with Profiling Floats. *Oceanography*, 22(2):34–43.

Ruiz-Barradas, A., Kalnay, E., Peña, M., BozorgMagham, A. E., and Motesharrei, S. (2017). Finding the driver of local oceanatmosphere coupling in reanalyses and CMIP5 climate models. *Climate Dynamics*, 48(7-8):2153–2172.

Saha, S., Moorthi, S., Pan, H.-L., 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.-T., Chuang, H.-Y., Juang, H.-M. H., Sela, J., Iredell, M.,

Treadon, R., Kleist, D., Van Delst, P., Keyser, D., Derber, J., Ek, M., Meng, J., Wei, H., Yang, R., Lord, S., Van Den Dool, H., Kumar, A., Wang, W., Long, C., Chelliah, M., Xue, Y., Huang, B., Schemm, J.-K., Ebisuzaki, W., Lin, R., Xie, P., Chen, M., Zhou, S., Higgins, W., Zou, C.-Z., Liu, Q., Chen, Y., Han, Y., Cucurull, L., Reynolds, R. W., Rutledge, G., and Goldberg, M. (2010). The NCEP Climate Forecast System Reanalysis. *Bulletin of the American Meteorological Society*, 91(8):1015–1057.

Saha, S., Moorthi, S., Wu, X., Wang, J., Nadiga, S., Tripp, P., Behringer, D., Hou, Y.-T., Chuang, H.-y., Iredell, M., Ek, M., Meng, J., Yang, R., Mendez, M. P., van den Dool, H., Zhang, Q., Wang, W., Chen, M., and Becker, E. (2014). The NCEP Climate Forecast System Version 2. *Journal of Climate*, 27(6):2185–2208.

Saha, S., Nadiga, S., and Thiaw, C. (2006). The NCEP climate forecast system. *Journal of Climate*, pages 3483–3517.

Shutts, G. (2005). A kinetic energy backscatter algorithm for use in ensemble prediction systems. *Quarterly Journal of the Royal Meteorological Society*, 131(612):3079–3102.

Singleton, T. (2011). *Data Assimilation Experiments with a Simple Coupled Ocean-Atmosphere Model*. Ph.d. thesis, University of Maryland, College Park.

Sluka, T. (2016–2018a). Hybrid global ocean data assimilation system (Hybrid-GODAS). http://www.github.com/umd-aosc/hybrid-godas.

Sluka, T. (2016–2018b). Universal multi-domain local ensemble transform kalman filter (UMD-LETKF). http://www.github.com/travissluka/umd-letkf.

Sluka, T. C., Penny, S. G., Kalnay, E., and 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., and Lawless, A. S. (2015). Exploring strategies for coupled 4D-Var data assimilation using an idealised atmosphere-ocean model. *Tellus A*, 67:1–25.

Smith, P. J., Lawless, A. S., and Nichols, N. K. (2018). Treating Sample Covariances for Use in Strongly Coupled Atmosphere-Ocean Data Assimilation. *Geophysical Research Letters*, 45(1):445–454.

Tardif, R., Hakim, G. J., and Snyder, C. (2014). Coupled atmosphere-ocean data assimilation experiments with a low-order climate model. *Climate Dynamics*, 43(5-6):1631–1643.

Thomas, C. (2017). *Multivariate Correlations: Balance Operators and Variable Localization in Ensemble Data Assimilation*. phd dissertation, University of Maryland.

Waters, J., Lea, D. J., Martin, M. J., Mirouze, I., Weaver, A., and While, J. (2015). Implementing a variational data assimilation system in an operational 1/4 degree global ocean model. *Quarterly Journal of the Royal Meteorological Society*, 141(687):333–349.

Whitaker, J. S. and Hamill, T. M. (2012). Evaluating Methods to Account for System Errors in Ensemble Data Assimilation. *Monthly Weather Review*, 140(9):3078–3089.

Xue, Y., Wen, C., Kumar, A., Balmaseda, M., Fujii, Y., Alves, O., Martin, M., Yang, X., Vernieres, G., Desportes, C., Lee, T., Ascione, I., Gudgel, R., and Ishikawa, I. (2017). A real-time ocean reanalyses intercomparison project in the context of tropical pacific observing system and ENSO monitoring. *Climate Dynamics*, 49(11-12):3647–3672.

Yang, S.-C., Kalnay, E., Hunt, B., and E. Bowler, N. (2009). Weight interpolation for efficient data assimilation with the Local Ensemble Transform Kalman Filter. *Quarterly Journal of the Royal Meteorological Society*, 135(638):251–262.

Yoshida, T. and Kalnay, E. (2018). Correlation-cutoff method for Covariance Localization in Strongly Coupled Data Assimilation. *Monthly Weather Review*.

Zebiak, S. E. (1989). On the 3060 Day Oscillation and the Prediction of El Niño. *Journal of Climate*, 2(11):1381–1387.

Zhang, F., Snyder, C., and Sun, J. (2004). Impacts of Initial Estimate and Observation Availability on Convective-Scale Data Assimilation with an Ensemble Kalman Filter. *Monthly Weather Review*, 132(5):1238–1253.

Zhang, L., Kumar, A., and Wang, W. (2012). Influence of changes in observations on precipitation: A case study for the Climate Forecast System Reanalysis (CFSR). *Journal of Geophysical Research Atmospheres*, 117(8):1–14.

Zhang, S., Harrison, M. J., Rosati, A., and Wittenberg, A. (2007). System Design and Evaluation of Coupled Ensemble Data Assimilation for Global Oceanic Climate Studies. *Monthly Weather Review*, 135(10):3541–3564.

Zhou, Y. (2014). *Minimizing reanalysis jumps due to new observing systems*. PhD thesis, University of Maryland.

Zweng, M., Reagan, J., Antonov, J., Locarnini, R., Mishonov, A., Boyer, T., Garcia, H., Baranova, O., Johnson, D., Seidov, D., and Biddle, M. (2013). World Ocean Atlas 2013, Volume 2: Salinity. Technical report, NOAA Atlas NESDIS 74.