

November 2007.

The Bluelink ocean data assimilation system (BODAS)

Peter R. Oke^{1*}, Gary B. Brassington², David A. Griffin¹ and Andreas Schiller¹

¹ *CSIRO Marine and Atmospheric Research and Wealth from Oceans Flagship
Program, Hobart, Tasmania, Australia*

² *Bureau of Meteorology Research Centre, Melbourne, Victoria, Australia*

* *Corresponding author. Email address: peter.oke@csiro.au*

Abstract

Bluelink is Australia's contribution to the Global Ocean Data Assimilation Experiment (GODAE). The goals of GODAE include the development and application of eddy-resolving, data assimilating ocean forecast systems. In this paper, we describe several aspects of the Bluelink effort that are motivated by this goal. One of the main innovations of Bluelink is the development of the Bluelink Ocean Data Assimilation System (BODAS). The initial test-bed for BODAS is the Bluelink ReANalysis (BRAN), a multi-year model integration with data assimilation. The Bluelink model is a global ocean general circulation model that is eddy-resolving in the Australian region. Observations that are assimilated into BRAN include satellite altimetry, sea surface temperature and in situ temperature and salinity data from Argo, XBT, TAO and other sources.

BODAS is an ensemble optimal interpolation system that uses an ensemble of in-

traseasonal anomalies from a free running model to estimate the background error covariances (BECs). The ensemble-based BECs are multivariate and inhomogeneous and are shown to reflect the length-scales, the anisotropy and the covariability of mesoscale oceanic processes. We evaluate the performance of BODAS and BRAN (version 1.5), spanning the period January 2003 to June 2006, by comparing reanalysed fields to a range of satellite-derived and in situ observations. Specifically, we demonstrate that BRAN realistically reproduces the mesoscale circulation around Australia, representing both the broad-scale circulation and, in many instances, the relatively small-scale mesoscale features. Quantitatively, we show that reanalysed fields in the region around Australia are typically within 6-12 cm of withheld altimetric observations, within 0.5-0.9° of observed sea-surface temperature and within 4-7 cm of observed coastal sea-level. Comparisons with Argo profiles and surface drifting buoys show that BRAN fields are within 1° of observed sub-surface temperature, within 0.15 psu of observed sub-surface salinity and within 0.2 m s⁻¹ of near-surface currents. We identify initialisation as a key area in which the Bluelink system could be improved.

Key words: Data Assimilation, Ensemble Optimal Interpolation, Ocean Reanalysis, GODAE, Operational oceanography, Multivariate assimilation

1 Introduction

Bluelink is an Australian partnership between the Commonwealth Scientific and Industrial Research Organisation, the Bureau of Meteorology (BoM) and the Royal Australian Navy. The primary objective of Bluelink is the development of a forecast system for the mesoscale ocean circulation in the Australian region. The Bluelink forecast system (Brassington et al. 2007) became operational at the BoM in August 2007 (www.bom.gov.au/oceanography/forecasts/).

Bluelink represents Australia's contribution to the Global Ocean Data Assimilation Experiment (GODAE). The goals of GODAE (Smith 2000; Le Traon et al. 2001) include the application of state-of-the art ocean models and assimilation methods for short-range open-ocean forecasts; and the provision of global ocean analyses and reanalyses for developing and improving understanding of the oceans, improving assessments of the predictability of ocean systems, and as a basis for improving the design and effectiveness of the global ocean observing system.

There are several data assimilating ocean modelling systems used for operational ocean forecasting or reanalyses on either regional or global scales that contribute to the goals of GODAE. These include Bluelink from Australia (Oke et al. 2005; Brassington et al. 2007) that is described here; FOAM from the United Kingdom (e.g., Bell et al. 2000; Martin et al. 2007); HYCOM and NLOM from the United States (e.g., Cummings et al. 2005; Smedstad et al. 2003); Mercator from France (Brasseur et al. 2006); TOPAZ from Norway (topaz.nersc.no); MOVE and COMPASS-K from Japan (Kamachi et al. 2004) and the ECCO group (www.ecco-group.org).

The purpose of this paper is to describe the main elements of the Bluelink system; and to evaluate its performance in a multi-year reanalysis. We briefly describe the model; and provide a detailed description of the data assimilation system and how it's applied. We also present results from the Bluelink ReANalysis (BRAN), a multi-year model integration with data assimilation, including a series of comparisons with observations. The Bluelink system has also been used to explicitly evaluate the impact of different observation types (altimetry, Argo and SST) on the performance of BRAN experiments (Oke and Schiller 2007).

The key elements of the Bluelink system are the Bluelink Ocean Data Assimilation System (BODAS; Oke et al. 2005) and the Ocean Forecasting Australia Model (OFAM), a global ocean general circulation model. One of the main innovations under Bluelink is the development of BODAS. Therefore, the focus of this paper is the development and evaluation of BODAS through its application to BRAN.

BODAS is an ensemble optimal interpolation (EnOI) system, similar to that described by Oke et al. (2002) and Evensen (2003; his Appendix B). BODAS uses model-based, multivariate background error covariances (BECs). The BECs used in BODAS are analogous to the Gaussian covariances used in more traditional optimal interpolation systems (e.g., Carton et al. 2000; Cummings 2005; Brasseur et al. 2006), and are the means by which an observation of some variable is projected onto the full model state at that time, including all model grid points and all model variables. There are many benefits in using the model-based covariances of EnOI. For example, the BECs reflect the length-scales and the anisotropy of the ocean circulation in different regions. They also quantify the covariances of different model variables in a dynamically consistent way (we use the term “dynamically consistent” to describe an ocean state that *can* be generated by the model). Additionally, EnOI does not involve the assumption of explicit balance constraints (e.g., Burgers et al. 2002) that are not valid everywhere. EnOI has the advantage that it can be used to readily assimilate different observation types in a single step. This contrasts to many other data assimilation schemes that require special treatment to assimilate both satellite derived sea-level anomalies (SLA) and in situ temperature (T) and salinity (S) (e.g., Cooper and Haines 1996; Troccoli and Haines 1999; Segsneider et al. 2000; Fox et al. 2002; Guinehut et al. 2004;

Cummings 2005; Martin et al. 2007; Chassignet et al. 2007). Despite these benefits, there are some limitations of EnOI. For example, a straight forward implementation of EnOI does not typically preserve water mass properties during the assimilation step. This makes EnOI inappropriate for some climate applications, where preservation of water masses may be regarded as essential. The estimates of the BECs for an EnOI scheme are only an approximation to the true BECs. Therefore EnOI is not an optimal method of assimilation. Also, EnOI does not provide time varying estimates of analysis errors, such as those derived from an Ensemble Kalman Filter (EnKF; e.g., Evensen 2003). Finally, the inversion step in EnOI requires some form of domain decomposition, in practice, while variational methods of assimilation do not require such approximations about the localised influence of observations.

The primary test-bed for BODAS is BRAN. BRAN is a multi-year integration of OFAM, where BODAS is used to sequentially assimilate observations once every 7 days. Observations assimilated in BRAN include SLA from satellite altimetry, satellite-derived sea-surface temperature (SST) and in situ T and S profiles. Conceptually, BRAN is a three-dimensional time-varying synthesis of oceanic observations that uses OFAM as a dynamic interpolator. The sequential nature of the assimilation used here means that BRAN can be regarded as a series of 7-day “forecasts”. Of course, the skill of BRAN should exceed that of an equivalent real-time forecast system with a 7-day update cycle, because of the use of real-time surface fluxes and observations; and particularly because of the latency of real-time altimetry in the operational system. Despite these differences, we assess the short-range predictive skill of the system for SLA and SST in BRAN as a way of measuring the potential performance of the operational system.

Results from the first BRAN experiment (BRAN1.0) are described by Oke et al. (2005), where it was shown that through the assimilation of SLA and in situ T and S observations, using BODAS, OFAM can produce three-dimensional, time-varying fields that are qualitatively consistent with the real ocean. BRAN1.5 is the next multi-year integration of the Bluelink system. The main differences between BRAN1.0 and BRAN1.5 are the inclusion of assimilation of SST, changes in the initialisation scheme and improvements to various aspects of BODAS and OFAM that are described below. In this paper we also provide a more comprehensive, quantitative assessment of BRAN than Oke et al. (2005). Results from the latest BRAN experiment (version 2.1) are described by Schiller et al. (2007), showing details of the seasonal circulation of the various current systems in the Asian-Australian region, with a strong focus on the circulation in the Indonesian Seas. BRAN2.1 is a 14-year reanalysis, using a configuration that is very similar to that of BRAN1.5.

This paper is organised as follows. The components of the reanalysis system are described in section 2, followed by an analysis of some characteristics of BODAS in section 3. Results from BRAN1.5, including comparisons with observations and an assessment of the predictive skill of the system are presented in section 4, followed by an analysis in section 5. Finally, a summary and the conclusions are presented in section 6.

2 Reanalysis System

2.1 OFAM

OFAM is based on version 4.0d of the Modular Ocean Model (Griffies et al. 2004), using the hybrid mixed layer model described by Chen et al. (1994). OFAM is intended to be used for reanalyses and short-range prediction. The horizontal grid has 1191 and 968 points in the zonal and meridional directions respectively; with $1/10^\circ$ horizontal resolution around Australia ($90\text{-}180^\circ\text{E}$, south of 17°N). Outside of this domain, the horizontal resolution decreases to 0.9° across the Pacific and Indian basins (to 10°E , 60°W and 40°N) and to 2° in the Atlantic Ocean. OFAM has 47 vertical levels, with 10 m resolution down to 200 m depth. The topography for OFAM is a composite of topography sources including *dbdb2* (www7320.nrlssc.navy.mil/DBDB2_WWW/) and *GEBCO* (www.ngdc.noaa.gov/mgg/gebco/). Horizontal diffusion is zero. Horizontal viscosity is resolution and state-dependent according to the Smagorinsky viscosity scheme (Griffies and Hallberg 2000).

OFAM is initialised with a blend of climatologies from Ridgway et al. (2002) and Levitus (2001); and is forced at the surface using 6-hourly fluxes of momentum, heat and freshwater from ERA-40 (www.ecmwf.int/research/era/) for 1992 to mid-2002; and using ECMWF 6-hourly forecasts from mid-2002 to 2006. A 13-year spin-up run, with no data assimilation, has been performed. During this spin-up OFAM forcing includes a flux correction, restoring the SST to a blend of Reynolds-SST (Reynolds and Smith 1994) and high-resolution satellite-derived observations over 30 days; and restoring surface salinity to monthly climatologies (Levitus 2001) over 30 days. An analysis of the mod-

elled trends of globally averaged sea-level indicate that the modelled upper ocean reaches a state of quasi-equilibrium after 3 years of integration. Interior T and S properties require much longer to properly equilibrate. The mean sea-level (MSL) during the spin-up run is shown in Figure 1. This field shows evidence of the major current systems in the Australian region including the East Australian Current (EAC), the Leeuwin Current and the Antarctic Circumpolar Current (ACC).

2.2 BODAS

Analyses of sea-level η , T, S, and horizontal currents (u,v), are computed by solving the analysis equations,

$$\mathbf{w}^a = \mathbf{w}^b + \mathbf{K} (\mathbf{w}^o - \mathbf{H}\mathbf{w}^b) \quad (1)$$

$$\mathbf{K} = (\rho \circ \mathbf{P})\mathbf{H}^* (\mathbf{H}(\rho \circ \mathbf{P})\mathbf{H}^* + \mathbf{R})^{-1}, \quad (2)$$

$$\text{where } \mathbf{w} = [\eta \ T \ S \ u \ v]^* \quad (3)$$

is the state vector; superscripts a , b , o and $*$ denote analysis, background, observed and matrix transpose respectively (here the background field refers to a model-generated estimate of the ocean state at the analysis time; sometimes called the first guess); \mathbf{K} is the gain matrix; \mathbf{H} is an operator that interpolates from the model grid to observation locations; ρ is a correlation function used to localise the ensemble-based BECs in \mathbf{P} ; \mathbf{R} is the observation error covariance matrix; and the open circles denote a Schur, or Hadamard, product (an element by element matrix multiplication). This formulation of the analysis equations is the same as that of Houtekamer and Mitchell (2001).

Estimates of the BECs in (2) are given by

$$\mathbf{P} = \mathbf{A}\mathbf{A}^T/(n - 1), \quad (4)$$

where n is the ensemble size and \mathbf{A} is an ensemble of model anomalies. Each anomaly field consists of all model variables included in (3). The calculation of the anomalies for an EnOI system is critical to the performance of the scheme. They should be computed in such a way that the scales of variability and features represented by the anomalies resemble the dominant errors of the model. For example, for application to OFAM, where we seek to correctly reproduce the mesoscale variability around Australia, we expect that the errors of an individual forecast will be dominated by the errors associated with mis-placement of eddies. Section 4 shows that this turns out to be true, although it is also true that errors at other scales exist. We hope to address this compromise in the future. Consequently, for application to OFAM, the ensemble of anomalies are generated by calculating intraseasonal anomalies derived from the spin-up run:

$$\mathbf{A} = \alpha [\mathbf{w}'_1 \quad \mathbf{w}'_2 \quad \cdots \quad \mathbf{w}'_n] \quad (5)$$

where α is a scalar that can tune the magnitude of the covariances for a particular application (for BRAN1.5, $\alpha = 1$); \mathbf{w}'_i is the i^{th} intraseasonal anomaly, defined here as the 3-day average of the model state minus the seasonal cycle of the model spin-up run. Using this approach, anomaly fields in the ensemble look like typical eddy fields (in the high resolution region). Because the ensemble in (5) is essentially a time series of anomalies, care is taken when constructing the ensemble to ensure that the time series is stationary, with zero mean and no trend. For both BRAN1.0 and BRAN1.5, BODAS uses an

ensemble size of 72, with one anomaly from every month of the last 6 years of a 13-year model spin-up.

The BECs in (4) are localised in the horizontal around each observation in (2) using the localising correlation function ρ . Elements of ρ are defined by the quasi-Gaussian function of Gaspari and Cohn (1992), after Houtekamer and Mitchell (2001). Localisation has been shown to reduce the effects of sampling error for applications of an EnKF (e.g., Hamill et al. 2001) and EnOI (Oke et al. 2006). The localising correlation function in ρ forces the BECs to reduce to exactly zero, over L° from an observation location. Note that in the application, we do not localise the covariances in the vertical direction. The present implementation of BODAS uses a uniform horizontal radial distance $L=8^\circ$, corresponding to an e -folding scale of about 3.5° . This has several ramifications for the system's performance. Firstly, the rank of the estimated BECs in \mathbf{P} is increased significantly. Using an ensemble size of n , the rank of \mathbf{P} without localisation is at most $n - 1$. By contrast, using an ensemble of only 72 with localisation with $L=8^\circ$ we estimate that the effective rank of $\rho \circ \mathbf{P}$ becomes $\mathcal{O}(10^4)$ (Oke et al. 2005). This enables the assimilation system to determine analysis increments that fit the background innovations, given by $(\mathbf{w}^o - \mathbf{H}\mathbf{w}^b)$ from (1), reasonably well. There are however, a few drawbacks to localisation. For example, analyses are not as dynamically balanced as they would be without localisation (Mitchell et al. 2002; Oke et al. 2006). Also, the inversion of the innovation covariance matrix $(\mathbf{H}(\rho \circ \mathbf{P})\mathbf{H}^T + \mathbf{R})$ becomes expensive, since the techniques for computational efficiency described by Evensen (2003) are not suitable when localisation is used. This makes the practical implementation of BODAS, described in the following sections, a technical challenge.

2.3 BRAN

Observations that can be assimilated for BRAN experiments include SLA from all altimeters (ERS 1 and 2, Topex/Poseidon, Geosat Follow-On, Jason and Envisat) and from a coastal tide gauge array around Australia. When SLA observations are assimilated, the background innovation ($\mathbf{w}^o - \mathbf{H}\mathbf{w}^b$) from (1), is constructed by first subtracting the MSL, shown in Figure 1, from the model sea-level. This field then represents a SLA that can be compared directly to altimetric SLA. BRAN also typically assimilates T and S profiles from a range of field surveys including Argo, the TAO array (McPhaden et al. 1998), XBTs and field surveys (e.g., WOCE); and satellite SST (currently Pathfinder and AMSR-E). Quality-controlled T and S profiles are obtained from the Bluelink Ocean Archive (Ridgway et al. 2002).

Results presented in section 4 of this paper are restricted to BRAN1.5, a 3.5-year experiment covering the period January 2003 to June 2006. The developments, compared to BRAN1.0 (Oke et al. 2005), include the assimilation of satellite SST, tuned observation error statistics and improvements to the ensemble fields that define the BECs. Problems with the neglect of the surface albedo in BRAN1.0 are also fixed for BRAN1.5. Additionally, Topex/Poseidon (T/P) and coastal sea-level observations are withheld from BRAN1.5 so they can be used for independent validation in section 4.

2.4 Treatment of Observations

2.4.1 Observation time window

The ensemble-based method used by BODAS is capable of assimilating different observation types in a single step. We have found that it is beneficial to use a relatively long time window (i.e., 7-11 days), in order to better constrain the analysis. In practice, this simply means that with more observations, our analysis is better constrained by observations and is therefore less likely to have unrealistic features that can result when a sparse network of observations are analysed. To this end, for a typical application of BODAS, we use an 11-day time window for SLA, yielding approximately global coverage, a 7-day time window for in situ T and S observations and a single day's AMSR-E SST observations. Because our focus here is on reanalyses, the time window is centered around the analysis time. So an 11-day time window for SLA means that we use SLA observations for the day of the analysis, plus SLA observations for 5 days before and after the analysis day.

Depending on the time, there might be around 10^6 individual observations available for assimilation on any given day. Using the long time windows discussed above, this can result in over 3×10^6 individual observations available for assimilation at any single analysis. Calculation of a global analysis using (1) involves the inversion of a $p \times p$ matrix, where p is the number of observations (here 3×10^6). It is not practical to even store this matrix, let alone explicitly invert it. This issue is addressed by BODAS in two ways; by reducing the number of observations to be assimilated directly, and by decomposing the global domain into many smaller sub-domains. These tasks are described

below in sections 2.4.3 and 2.5, respectively.

2.4.2 Observation error estimates

For assimilation, every observation is weighted according to its expected error. Given the relatively long time window of observations assimilated in each analysis the observations do not all correspond to the analysis time. We have not yet implemented the so-called first-guess at appropriate time (FGAT) method (Huang et al. 2002). Therefore, to be consistent, we should not give observations made several days before or after the same weight as observations made at the analysis time. This is achieved in practice by adjusting the estimated observation error variance (diagonals of \mathbf{R}) according to the relative “age” of each observation. The observation error variance ϵ_o^2 for an individual observation is here defined as:

$$\epsilon_o^2 = \epsilon_{instr}^2 + \epsilon_{RE}^2 + \epsilon_{age}^2 \quad (6)$$

where ϵ_{instr}^2 is the estimated variance of the instrument error, ϵ_{RE}^2 is the estimated variance of the representation error (RE), sometimes referred to as the error of representativeness, and ϵ_{age}^2 is the estimated variance of the error associated with the relative age of an observation. The estimates of ϵ_{instr} used in BRAN1.5 are listed in Table 1, along with the range of values for ϵ_{RE} and ϵ_{age} .

Estimates of ϵ_{age} are given by

$$\epsilon_{age} = RMS_{mod} \left(1 - e^{-0.5|t^a - t^o|/t_{ef}} \right), \quad (7)$$

where RMS_{mod} is the spatially dependent RMS of the model fields about a seasonal cycle during the spin-up run; t^a is the analysis time; t^o is the time

of the observation; and t_{ef} is an e -folding time scale (here 3 days), following Oke et al. (2005). Therefore, if an observation is made at the analysis time, $\epsilon_{age} = 0$; and as $|t^a - t^o|$ increases, ϵ_{age} approaches RMS_{mod} , so that the influence of the observation on the analysis decreases. At ± 4 days, for example, $\epsilon_{age} \approx 0.5 \times RMS_{mod}$.

Estimates of ϵ_{RE} are calculated using the method described by Oke and Sakov (2007). This method provides estimates of RE for T, S and η that reflect the variance of unresolved mesoscale variability in the ocean. Figure 2 shows a map of ϵ_{RE} for η , showing large (small) RE estimates in regions of coarse (high) resolution and energetic (quiescent) mesoscale variability.

2.4.3 Super-observations

As outlined above, a direct solution to (1) requires the storage and inversion of a very large matrix. In practice, this is not feasible. The first practical strategy to address this technical challenge is super-obing. Super-obing is commonly used in objective analysis (e.g., Ducet et al. 2000) and data assimilation (e.g., Cummings 2005), particularly in numerical weather prediction where many redundant (compared to the model resolution) observations are available. Super-obing is where a number of observations, with a known error estimate, are combined to produce a single super-observation with a smaller error. For a given application, suppose there is more than one observation of a certain type (e.g., SLA from altimetry) that fall within a model grid cell. These observations either contain the same information, making all but one observation redundant; or they contain different information, presumably representing some sub-grid-scale process that cannot be resolved by the model.

In either case, it is appropriate to combine these observations into a single super-observation. In practice, we identify all observations within a model grid cell and compute the average position of all observations. This position is ascribed to the super-observation. These observations are then combined, using a simple weighted average, yielding the super-observation. The errors of all observations, from (6), are used to compute an appropriate error for the super-observation using standard error propagation techniques. This process of super-obsing is used to significantly reduce the number of observations that are directly assimilated using (1).

The level of super-obsing for any application can be controlled according to the computational resources available. For application to OFAM, we have chosen to compute one super-observation for every 6×6 model grid cells. This yields a super-observation every $0.6^\circ \times 0.6^\circ$ around Australia and every 12° in the North Atlantic Ocean. In practice, we only apply this super-obsing to observations of altimetric SLA, and to satellite SST. For both observation types, super-obsing in the manner described above yields about 20,000 super-observations of each type.

We choose to treat observations from T and S profiles (e.g., Argo) differently to SLA and SST. Instead of computing super-observations, we simply subsample the available profiles out, rather than average them, so that there is no more than one profile for every 6×6 model grid cells. This is because the horizontal averaging of profiles is somewhat problematic, with observations at varying depths; and with adjacent profiles often covering different depth ranges. For a given region covered by a 6 by 6 cluster of model grid cells, we therefore select the T and S profile with the most *good* observations and with-hold the other profiles for validation. We define a *good* observation as one

that passes all of the automatic quality control checks. This sub-sampling of profiles eliminates many profiles in the coarse resolution region of the model and far fewer in the high resolution region. The nominal vertical resolution of Argo is typically finer than the model grid, so super-obing of each profile is done to ensure that no more than one super-observation falls within each model depth level. This is a significant saving for profiles that are well-resolved over depth.

2.5 Domain Decomposition

Even with the significant reduction in the number of observations assimilated by BODAS due to the super-obing described above, a global inversion of the innovation covariance matrix in (1) remains computationally expensive. There are a number of approaches to address this issue. One such approach that is used in BODAS involves a domain decomposition, where the model domain is decomposed into a number of smaller sub-domains that are here referred to as analysis domains. For application to OFAM, the global domain is divided into about 800 analysis domains. An analysis for each of these domains is computed independently, using observations from an observational domain. The observational domain spans each analysis domain and includes a surrounding halo region from which observations are used to compute each analysis (typically less than 3000 observations in each domain). Provided the size of the halo regions correspond to the localising length-scales, described in section 3, and provided the localising function reaches exactly zero within this halo (e.g., Gaspari and Cohn 1999), the analysis is seam-less, with no discontinuities between adjacent analysis domains; and is equivalent to the solution from a

global analysis, but without the prohibitive computational expense and with significantly more numerical accuracy. A benefit of this approach is that it is highly parallel; and can take advantage of computer platforms with multiple processors.

2.6 Initialisation

The problem of initialisation is common to all data assimilating geophysical applications. Briefly, when a non-linear model is updated with an unbalanced state, the model generates an artificial response that may involve, for example, the excitation of various types of waves. This response can often degrade the quality of a subsequent integration. Given that a statistically generated analysis field, such as that produced by BODAS, is always unbalanced to some degree when applied to a non-linear model, care must be taken when initialising, or updating, the model. For a comprehensive review of issues in initialisation, see Daley (1991). There are a number of options for dealing with initialisation that include digital filtering (e.g., Lynch and Huang 1992), normal mode initialisation (e.g., Moore 1990), balance constraints (e.g., Burgers et al. 2002), incremental analysis updating (Bloom et al. 1996; Ourmières et al. 2006) and nudging (e.g., Leslie et al. 1998). Here, we have opted for a very simple, inexpensive and conservative approach, and chosen to use nudging.

Details of the initialisation procedure we use are represented schematically in Figure 3. We integrate the model for 7-days with no data assimilation. After 6 days, we store the fields necessary for a model restart. On the 7th day, we compute the background field (BGF) that is the 24-hour mean. We then combine the BGF and the observations using BODAS and compute an analysis

of the full model state. We then restart the model, using the restart fields referred to above, and then nudge the model T, S and η towards the analysis using a time-scale of 1 day. The choice of the time-scale involves some tradeoffs. If it is too long, the model will not be corrected by the observations. If it is too short, however, the model will be continuously driven towards a stationary field for the whole day-long nudging period, suppressing all transients. Our choice of a 1-day nudging time-scale applied for 1 day means that the actual increment added to the model should only be about 70% of the total increment computed by BODAS. While this is not ideal, our preliminary experiments led us to the expectation that this is a reasonable compromise between forcing the model too hard and creating numerical instabilities; and forcing the model too gently and not adequately constraining the model.

For BRAN1.5, we do not explicitly update the velocities. We simply allow the currents to adjust according to the model equations during the period of nudging. However, we note that some centers have found a benefit in explicitly updating with geostrophically balanced currents (e.g., Martin et al. 2007). After 1 day of nudging, we again integrate the model forward in time for another 7-days and so the cycle repeats.

3 Characteristics of BODAS

As stated above, one of the advantages of EnOI is that the BECs are inhomogeneous and anisotropic, reflecting the variability and length-scales of the ocean circulation. An example of the anisotropy and inhomogeneity of the ensemble-based BECs used by BODAS is shown in Figure 4. The chosen examples show the localised ensemble-based cross-correlation between sea-level

at a reference location and sea-level in the surrounding region. These fields demonstrate the region of influence of an observation at the reference location. Where the correlation is positive, the increment that is given by the term $\mathbf{K}(\mathbf{w}^o - \mathbf{H}\mathbf{w}^b)$ in (1), due to each observation has the same sign as the background innovation $(\mathbf{w}^o - \mathbf{H}\mathbf{w}^b)$. So if the observed sea-level is higher than the background field at the reference location, in the absence of other observations in the same region, the solution to (1) will produce an increment that is also positive. The magnitude of the increment depends on the relative magnitudes of the estimated background error and observation error covariances (i.e., the relative magnitudes of $\mathbf{H}(\rho \circ \mathbf{P})\mathbf{H}^T$ and \mathbf{R}). The structure of the increment also depends on the structure of the localised background error covariance, $\rho \circ \mathbf{P}\mathbf{H}^T$.

The reference locations for the examples presented in Figure 4 are at 32.5°S , corresponding to the typical separation point of the EAC (Godfrey et al. 1980). These examples include correlations when the reference location is on the continental shelf at 115 m depth, where the correlation field has short decorrelation scales in the across-shore direction and long decorrelation scales in the alongshore direction (Figure 4a). The long length-scales in the alongshore direction are probably a reflection of the covariability associated with northward propagation of slow-moving coastal trapped waves and the along-shore advection of the EAC and wind-driven circulation. Also shown is a correlation map when the reference location is over the continental slope at 1800 m depth (Figure 4b), less than 50 km east of the shelf example (Figure 4a). This field shows a more isotropic correlation, but with a tendency to have higher correlations to the east and south-east, in the typical direction of the EAC as it separates from the coast. The final example considered here shows correlations when the

reference location is over the deep ocean at 4500 m depth (Figure 4c). This field shows a somewhat isotropic correlation field, but also shows areas of weak negative correlation to the north-east and south-west. These negative correlations may be related to variations in sea-level associated with the typical eddy field in this region. The correlation fields presented in Figure 4 highlight the anisotropy and inhomogeneity of the ensemble-based correlations, with relatively long quasi-isotropic correlations in the deep ocean, transitioning to very anisotropic correlations nearer the coast.

An example of the multivariate nature of BODAS is demonstrated in Figures 5 and 6, showing the increments from a single observation analysis (i.e., where an analysis is computed by combining a single observation with a background field). In this example, an observation of sea-level at the coast is presumed to be 20 cm lower than the background sea-level. BODAS is used to calculate the increments to the full model state (η , T , S , u and v) in the surrounding region. We choose to use sea-level from a coastal location off South Australia at 140.35°E and 37.85°S. This location is arguably the best region for observing wind-driven, coastal upwelling along the Australian coastline (e.g., Lewis 1981; Kampf et al. 2004). We therefore expect the increments to be consistent with a conceptual model of upwelling (i.e., consistent with the oceanic response to south-easterly winds).

Decreased sea-level at the coast may be due to a number of factors that include, but are not limited to, wind-driven upwelling. However, the dominance of locally wind-driven circulation in this region leads us to the expectation that it will also dominate the statistical properties of the ensemble fields here. We find that the increments to sea-level and surface currents in Figure 5 are consistent with wind-driven upwelling, with reduced sea-level along the coast,

a relatively strong coastal jet and weak south-westerly flow in the deep ocean. Figure 6 shows the impact of the observation along a shore-normal section offshore of the observation location. This figure shows that the negative sea-level increment is strongest at the coast as we expect, and approaches zero moving offshore in a quasi-exponential fashion. Increments to the across-shore currents show an offshore flow of up to 5 cm s^{-1} over the top 30 m, that is consistent with an offshore wind-driven Ekman layer, and a weak, shoreward return flow at depth. Increments to the along-shore currents are consistent with a baroclinic, wind-driven coastal jet, with the strongest currents of 0.5 m s^{-1} at the surface near the coast. Figure 6d-e shows the background and analysed T and S fields. Both of these fields show an uplift of isotherms and isohalines, with an implied vertical excursion of about 25 m near the shelf break and 40-50 m over the upper slope. From this analysis, we conclude that the statistical properties of the ensemble are consistent with a conceptual model of wind-driven upwelling in the region considered.

The analysis described above indicates that if the model does not produce an upwelling event, due to incorrect surface forcing for example, and we assimilate a single observation of sea-level at the coast that reflects the upwelling through reduced sea-level, then BODAS will compute increments that more closely match the observation at the coast, and in a manner that is consistent with wind-driven upwelling. While this feature is presented here as a benefit of EnOI, there is also a down-side. Suppose the model-observation mismatch is due to a mis-represented process other than wind-driven upwelling. Say, the encroachment of an eddy, or the propagation of a coastal trapped wave. Then the EnOI-derived increments will still be consistent with upwelling as in Figures 5 and 6. This may not be desirable, and may result in a reanalysed

state that is somewhat inconsistent with reality. However, we note that the example described here is very idealised. In practice, we typically assimilate multiple observations of various types (e.g., sea-level, T and S) and in the presence of additional observations that provide a more complete picture of the ocean state, BODAS will compute increments that more appropriately represent the true circulation.

Another limitation of EnOI that can be seen from the example described by Figures 5 and 6 is that it implies a symmetry in the increments for an observation-model mis-match of the opposite sign. Suppose that the observation-model difference considered above is reversed. That is, the observation is 20 cm higher than the background sea-level. In this case, the increments simply have the opposite sign to those presented in Figures 5 and 6. This might occur if a modelled upwelling is too strong, or perhaps if the model fails to represent a downwelling event. However, we note that studies into the dynamics of idealised wind-driven upwelling (e.g., Allen et al. 1995) and wind-driven downwelling (e.g., Allen and Newberger 1996) on the continental shelf have shown that they are very different, and do not simply result in symmetric anomalies about some mean field as the EnOI-based increments would imply. These problems are the result of assuming ergodicity in the BECs. A better approach is to employ some flavour of the EnKF (e.g., Evensen 2003; Sakov and Oke 2007), where the anomalies in (5) evolve in time and more accurately reflect the dominant dynamical processes for a particular time. However, these types of filters require an ensemble of states to be evolved, which is currently too computationally expensive for applications as large as OFAM.

4 Results

In this section, we present a qualitative assessment of BRAN, a quantitative assessment of the predictive skill for SLA and SST in BRAN; comparisons with T and S fields from Argo observations, and comparisons of near-surface currents with surface drifting buoys. For each comparison presented here, we use daily mean fields from BRAN and instantaneous observations, unless otherwise stated. We also typically refer to BRAN fields as a forecast. This is because BRAN can be regarded as a series of 7-day forecasts (Figure 3) for the purpose of assessing predictive skill. However, the use of real time surface fluxes and observations; and the latency of altimetry for the operational system, means that results from BRAN should exceed the performance of a true forecast system.

4.1 Qualitative assessment

A qualitative assessment of the ability of BRAN to realistically reproduce the mesoscale variability in the Tasman Sea is presented in Figure 7. We show a series of comparisons, from the 15th day of every month during 2004, between 6-day composite SST images from AVHRR and 5-day averaged SST from BRAN. These SST observations are not assimilated into BRAN (AMSR-E data are the only SST data assimilated). Also included in Figure 7 are Lagrangian trajectories, computed from the time-varying surface velocities in BRAN. This region is dominated by the EAC and a complex field of eddies (e.g., Nilsson and Cresswell 1981). These comparisons are representative of the rest of the BRAN period and show good agreement between the observed and

reanalysed mesoscale features. The broad-scale features, such as the southward penetration and position of the warm water associated with the EAC, are well reproduced in BRAN. Moreover, there are also many examples in Figure 7, where even relatively small-scale eddies are realistically represented in BRAN in a qualitative sense. Examples include the cold-core feature around 156°E, 31°S in September 2004 and the structure of the filaments between the warm EAC water ($> 19^{\circ}\text{C}$) and the cold Tasman Sea water ($< 19^{\circ}\text{C}$) around 32-36°S in June-August 2004.

The comparisons in Figure 7 are repeated for the region off south-west Western Australia in Figure 8. This region is dominated by the southward flowing Leeuwin current and a complex field of mesoscale eddies and meanders that typically spawn from instabilities along the coast (e.g., Ridgway and Condie 2004; Feng et al. 2005). Most of the features evident in the observations are well reproduced in BRAN. For example, the changes in the southward penetration of the warm Leeuwin current water along the coast is well modelled. Similarly, much of the mesoscale variability in the observations is also evident in BRAN.

A comparison between reanalysed, observed and climatological T along a section off eastern Australia, at 30°S, in August 2003 is shown in Figure 9b-d. Also shown in Figure 9a is the SLA field for the same period. The SLA field is included to give the context of these comparisons; and show the presence of a large warm-core eddy at around 156°E and a series of smaller, cold-core eddies to the east. Both the BRAN and XBT fields show the sub-surface T structure that is consistent with these features. The qualitative agreement between the observed and reanalysed T fields is good. By contrast, the T climatology (Ridgway et al. 2002) shows broadscale agreement with the observations, but does not represent any of the mesoscale features as we expect.

A comparison between reanalysed, observed and climatological T along a section off Western Australia, along the IX15 line, in July 2004 is shown in Figure 10b-d. The corresponding SLA field, from BRAN, is also shown in Figure 10a to provide the context of the circulation. The SLA field shows a series of warm- and cold-core features along the XBT section. The BRAN T field shows good qualitative agreement with the XBT data. Again, the T climatology shows the same broad-scale structure as the observations, but does not show the mesoscale features as we expect.

Based on the comparisons described above, it is clear that BRAN qualitatively reproduces many aspects of the complex mesoscale circulation around Australia. However, in the examples shown, BRAN does not always get the mesoscale features in precisely the right locations or with the correct strength or intensity. This has implications for the interpretation of the statistical comparisons that we present below.

4.2 Assessment of predictive skill

As noted above, BRAN can be regarded as a series of 7-day forecasts. We wish to determine whether BRAN has any skill over a typical 7-day cycle. This is quantified here by calculating the root-mean-squared error (RMSE), here defined as the RMS difference between observed fields and reanalysed fields from BRAN. We assess the predictive skill, by showing the RMSE as a function of lead time, averaged over all 7-day cycles; and comparing it to the RMSE of the persistence of both the analysed and initialised fields, hereafter the “analysed persistence” and the “initialised persistence” respectively. Persistence is best described as a forecast of no change from the analysed or initialised state.

The RMSE of the analysed (initialised) persistence is computed by comparing the analysed (initialised) fields, valid for the analysis time (0 lead time), to observed fields that are valid for each day of the subsequent 7-day cycle. The analysed fields are simply those computed by BODAS. The initialised fields are daily means computed on the “nudging” day (Figure 3). We consider the initialised fields to be a dynamically filtered version of the analysis, where the model rejects features in the analysis, during the nudging period, that are not dynamically consistent. As a result, the RMSEs of the initialised fields are inevitably greater than the RMSEs of the analysed fields. However, we expect that over the course of the nudging day, the difference between the model field and the analysis gradually reduces. Therefore a daily mean field over this period doesn’t necessarily represent the final model state at the end of the nudging period.

We include statistics of the persistence of the analysis to enable us to assess which product provides better agreement with observations over a 7-day cycle; an N -day old analysis, or an N -day forecast? Similarly, we include statistics of the persistence of the initialised state because it enables us to assess whether, given the initial conditions of the initialised field, the model adds any skill during the subsequent 7-day integration. When the RMSE of the forecast is less than the RMSE of the initialised persistence, we conclude that the ocean model itself has some skill. When the RMSE of the forecast is less than the RMSE of persistence of the analysis, we conclude that the whole system (assimilation and model) has skill.

We note the use of a centered observation time window for assimilation (Figure 3) means that we should have some artificial skill. However, recall that we only assimilate SST from 1 day in 7; so data from 6 days in 7 are independent.

Similarly, we do not assimilate all T and S profiles available. We simply thin out the profiles so that there is no more than one profile for every 6-grid points. Also, recall that the observations made before or after the analysis time are down-weighted by ϵ_{age} as indicated in (6-7) and in Table 1. Therefore as the “age” of an observation increases, its impact on an analysis decreases.

The RMSE of SLA from BRAN forecasts and analysed and initialised persistence, using SLA observations from T/P (tracks are interleaved with Jason), are plotted in Figure 11. Corresponding calculations for SST are plotted in Figure 12, using SST observations from AMSR-E. Statistics are shown for the north-west (NW), north-east (NE), south-west (SW) and south-east (SE) quadrants of the Australian region. The RMSEs shown in Figures 11 and 12 indicate that the errors in SLA typically range from 6-12 cm; and that the errors in SST typically range from 0.5-0.9°C. As noted in the captions of Figures 11 and 12, the errors in BRAN SLA and SST are typically less than the errors in the spin-up run, with no data assimilation. Similarly, we note in the caption of Figure 12 that the errors in BRAN and also less than the differences between observations and seasonal climatology (Ridgway et al. 2002).

We find that the RMSE of the analysed persistence for SLA is consistently less than the RMSE of the SLA forecasts (Figure 11). Thus, we conclude that the analysis typically provides the best quantitative estimate of an SLA field over a 7-day forecast period.

We find that the RMSE of the initialised persistence for SLA is consistently greater than the RMSE of the SLA forecasts. This indicates that the model typically adds some skill to the initialised state over a 7-day forecast.

Figure 11 shows that there is a big discrepancy between the RMSE of the

analysed and initialised fields at day-0. This indicates that the nudging approach that we use here is ineffective in some regions; particularly in the southern quadrants. That is, the initialised fields do not usually match the analysed fields towards which they are being nudged. We expect that this is due to a number of factors. Recall that we simply nudge the model T, S and η towards the analysed T, S and η for 1 day, with a nudging time-scale of 1 day. We anticipate that if we update the velocities as well as T, S and η , using incremental analysis updating (Bloom et al. 1996) for example, then the subsequent forecast should retain more of the features introduced during the initialisation process. Indeed, this has been the experience of the UK Met Office (e.g., Martin et al. 2007); and recently conducted experiments using the Bluelink system confirm this. We also think that part of these problems may relate to the presence of systematic errors in some regions of the model. This suggestion is explored below, in section 5.

We expect the RMSE should increase monotonically over the course of a 7-day forecast cycle. However, this is not always the case here. In some regions, the RMSEs for SLA and SST typically decrease for some period over the course of a 7-day forecast (e.g., Figures 11a,c and 12a,b). In part, this may be due to sampling error; these statistics are computed over 182 cycles. However, it may also be due to problems with initialisation, where the model field is temporarily degraded by the shocks introduced when the model is initialised; and then may subsequently recover (e.g., Figure 12a). In some cases it may also be due to the definition of the initialised state. Recall that the initialised state (0-day forecast) is the daily mean field during the day of nudging (Figure 3).

The RMSE statistics for SST, shown in Figure 12, indicate that the forecast typically has a small amount of skill compared to both analysed and initialised

persistence. This indicates that the forecast SST is typically the best quantitative estimate of SST over a 7-day forecast; and that the system has predictive skill over 7-days. The RMSE for SST is generally less in the NE and NW, compared to the SE and SW. This is probably a reflection of the errors associated with the mixed layer scheme used in OFAM or errors in the surface fluxes. Both of these factors may result in errors in the depth of the ocean surface mixed layer that are likely to be most pronounced under strong wind conditions and weak stratification that frequently occur at high latitudes.

Analysis of the RMSEs for SST show a clear seasonal dependence. This is demonstrated in Figure 13 showing the RMSEs of BRAN forecasts for each season. The seasonal dependence is most pronounced in the southern quadrants, where the RMSE is typically about 0.15° greater in winter than it is in summer. Again, we suspect that this is due to errors in the mixed layer model or surface fluxes, and the poorer performance at high southern latitudes in Austral winter is due to poorly represented mixed layer variability during winter storms. By contrast, the RMSEs are largest during summer in the NW quadrant. This is probably associated with the representation of the NW monsoon that occurs in Austral summer. The NW monsoon results in strong surface heating and a large, intermittent injection of freshwater that results in barrier layers (Sprintall and Tomczak 1992). Barrier layers are particularly difficult to model and may be poorly represented in BRAN, especially if there are significant errors in the surface fluxes associated with the NW monsoon.

BRAN1.5 does not assimilate coastal sea-level from tide gauge stations, or altimetric sea-level in water less than 200 m deep. To assess the predictive skill of BRAN along the Australian coast, we compare SLA from BRAN to daily mean observations at 58 tide gauge stations along the Australian coastline

(locations are denoted in Figure 1). We present the RMSE of SLA, compared to SLA at tide gauge stations, as a function of lead time in Figure 14. Specifically, we show the RMSE of BRAN forecasts; persistence of the analysed fields, the initialised fields and the observed fields; and the RMSE of the spin-up run with no data assimilation. We present results that are averaged over all forecast cycles in each Australian state (states boundaries are denoted in Figure 1). We also show the corresponding average of the observed standard deviations in Figure 14.

For all Australian states, both the modelled SLA in BRAN and the spin-up run with no data assimilation generally have smaller RMSEs than any of the persistence fields, including persistence of the observed field (Figure 14), for lead times of greater than 2-days. This indicates that the decorrelation time-scale of coastal sea-level is quite short; and therefore that any persisted field does not provide a good indication of the SLA beyond 2-3 days into the future.

Interestingly, there is generally no significant difference in the RMSEs of BRAN and the spin-up run with no data assimilation. This indicates that the model skill is not due to the assimilation, but is likely to be simply due to the accuracy of the local wind forcing and the model's response to this forcing. Recall that we use delayed-model, quality controlled surface fluxes in BRAN. We do not therefore expect the quality of these fluxes to degrade over each forecast cycle, as we would if the fluxes were from a true weather forecast. This aspect of the reanalysis, together with the dominance of local wind forcing around much of Australia's coastline, explains why there is no significant increase in the RMSE of the forecast over the course of a typical 7-day forecast.

The statistics of the SLA at the stations off Tasmania (Figure 14c), two of which are off north-west Tasmania and one is off the east coast (see Figure 1), are different to those for other states. Firstly, the magnitude of the observed standard deviation is relatively small off Tasmania. Secondly, the RMSEs of the forecast SLA are significantly less than that of the spin-up run with no data assimilation. This is because the coastal circulation off Tasmania depends more on the circulation in the open ocean than on the local wind forcing. In this case, the influence of the offshore SLA that is constrained by the assimilation of altimeter and SST data has a positive impact on the skill of coastal sea-level forecasts.

The statistics off South Australia (Figure 14a) also represent somewhat of an anomaly compared to the other Australian states. Specifically, the magnitude of the observed variance is very large, owing to the strong wind-driven circulation, and the RMSE of the persistence of observed SLA is greater than that of the modelled SLA after 1-day. There is also a dip in the RMSE of the forecast after 1 day, possibly indicating that the assimilation does have a positive impact on the skill of SLA, but that the dominance of the local wind forcing means that this impact is quickly “forgotten” by the model.

We note that for most Australian states the RMSE of the analysed field on day 0 is typically greater than the RMSE of the forecast on day 7 (i.e., the background field used to compute the analysis). This result indicates that the extrapolation of the other observation types (including altimetric SLA, SST and sub-surface T and S) often produces analysed SLA near the coast that has larger errors than the background field. Recall that SLA from these tide gauges are not assimilated into this version of BRAN. We also note that the RMSE of the analysed field on day 0 is typically greater than the RMSE of the

initialised field on day 0. This result indicates that during the nudging process, the model appropriately filters the analysis, resulting in an initialised state that is in better agreement with the withheld observations. It also suggests that assimilation of coastal SLA would be worthwhile.

4.3 Comparisons with Argo

Profiles of the RMSE of T and S in BRAN1.5, the spin-up run (with no data assimilation) and monthly climatology (Ridgway et al. 2002) are presented in Figures 15 and 16, using both assimilated and withheld Argo profiles separately for comparison. Assimilated and withheld Argo profiles made between January 2003 and December 2005 (the overlapping period for BRAN1.5 and the spin-up run) are binned into 30° by 30° boxes (comprising about 9500 assimilated and 3000 withheld profiles in the Australian region) and RMSEs are computed for each region.

The RMSE of BRAN T and S is typically less than the RMSE of T and S in the spin-up run for both assimilated and withheld profiles (Figures 15 and 16). This demonstrates the benefit of data assimilation. By contrast the errors in BRAN T and S are typically greater than errors in climatology.

There are some places where the RMSE of BRAN T and S are greater than the RMSEs of the spin-up run. Examples include the withheld S-profiles off north-eastern Australia; and the withheld T-profiles, below 1000 m depth, south and south-west of Australia (Figures 15 and 16). There may be a number of reasons for this result. The BECs used by BODAS are only an approximation of the true BECs. This may result in a degradation of the BRAN state in some

regions. Also, the configuration of BODAS used here does not localise the BECs in the vertical spatial direction. As a result, a T or S observation near the surface, for example, may degrade the analysis at depth if the BECs are not well defined by the ensemble.

In general, the S-field is not very well constrained in BRAN. This may be because of the short length-scales of S. This means that the observing system required to properly observe S variability is very dense; and that the Argo profiles that are assimilated here may be insufficient to properly constrain the modelled S-field using the chosen assimilation method. In most regions, the RMSEs of BRAN S near the surface are greater than the RMSEs of S in the spin-up run. This is, in part, because the surface S in the spin-up run is restored to climatological S and surface S in BRAN is not.

Clearly, the exploitation of Argo observations in BRAN is not optimal. However, despite this, in general, we find that the RMSEs of T and S in BRAN are typically less than 1° and 0.15-0.2 psu respectively (Figure 15 and 16). The T field is generally quite well constrained considering the eddy-resolving nature of the model. However, the S field does not appear to be well constrained as discussed above.

4.4 Comparison with drifting buoys

Surface drifting buoys are an important observation platform that is not assimilated in BRAN1.5. We compare the near surface currents and T from BRAN with all available quality controlled and interpolated drifter data from MEDS with reanalysed fields from the second model level at 15 m depth. We

calculate the complex cross-correlation (Kundu 1976) between the observed and reanalysed currents and present the magnitude ρ , and phase angles θ , in Table 2, along with the RMSE of the speed of the currents and T. We find that the correlation between observed and reanalysed currents is 0.29-0.6, the phase angle (i.e., the error in the direction of the currents) is less than 6° in the northern quadrants and is 15° and 23° in the SW and SE quadrants respectively. The RMSE of the drifter-derived current speed is about 0.2 m s^{-1} and for T is about $0.35\text{-}0.81^\circ$.

5 Analysis

In the comparisons presented in Figures 11 and 12 we find that the RMSEs of the initialised fields are consistently greater than the RMSEs of the analysed fields. This is most pronounced in the southern quadrants of the Australian region (Figure 11c,d). To further investigate this we present geographically resolved estimates of the RMSE for SLA, compared to assimilated and unassimilated SLA observations, in Figure 17. Specifically, we show the RMSE of BODAS analyses, forecasts and, for comparison, $1/3^\circ$ SSALTO SLA maps produced by Aviso (accessed in September 2006; Ducet et al. 2006). Note that for BODAS and SSALTO, we only use observations from the day of each analysis (i.e., every 7th day), which by design corresponds to the day of each Aviso map.

Figure 17 shows that compared to the assimilated altimeter data, the errors of BODAS analyses are slightly less than the RMSEs of Aviso. This is likely to be due to differences in the assumed errors, length-scales and the time-windows used in the different analysis systems. T/P data are not assimilated

by BODAS, but is included in the SSALTO analyses, so it is no surprise that the Aviso RMSE is slightly less than the BODAS RMSE for these data. Most striking in Figure 17 are the differences in the RMSEs of the forecast, compared to the BODAS analyses. This is consistent with the results presented in Figure 11, showing a significant discrepancy between the RMSEs of the analysed and initialised fields for SLA. The RMSEs of the forecast are largest in the Tasman Sea and along the path of the Antarctica Circumpolar Current (ACC). Both of these regions correspond to areas of energetic mesoscale variability. The larger errors in the forecast may therefore be due to a poor initialisation, poor evolution of mesoscale features, or systematic errors in the model.

To determine whether there is any systematic errors in the model, such as a bias in MSL, we consider the statistics of the increments to sea-level (Figure 18). Specifically, we compute the mean sea-level increments, the RMS of the increments and the ratio of the absolute value of the mean to the RMS of the sea-level increments. In these calculations, we use the actual increments applied to OFAM that are the time integrals of the nudging term on the analysis day. We note that these increments are generally smaller than the increments computed by BODAS (typically 70-90%) because the nudging period is the same as the time-scale of the nudging, as discussed earlier.

The analysis equations (1-2), solved by BODAS, assume that the model is unbiased. This should result in mean increments of zero. We find that while the mean sea-level increments (Figure 18a) are close to zero in most of the Australian region, there are some regions where the mean increments are relatively large. A comparison between the magnitude of the mean and RMS increments provides an indication of the significance of these biases. This is

quantified in Figure 18c by the ratio of the absolute value of the mean to the RMS of the increments. Where the mean and RMS of the increments are of comparable magnitude (i.e., where the ratio in Figure 18c are large; near 1), we conclude that the model is consistently being adjusted in the “same direction” at each assimilation step. This is an indication that the model quickly “forgets” the introduced changes and reverts to its state before it is updated. This is symptomatic of a problem either with 1) the initialisation of the model, 2) a systematic error in the model, or 3) the observations. As a result, BODAS computes similar increments, with the same sign, at each assimilation step. Regions where this is the case (Figure 18c) include the regions around 120°E and 160°E off Antarctica. This region is not well observed by the global ocean observing system because it is covered by sea-ice for much of the year. Also this may be due to the lack of a sea-ice model in OFAM. Other regions where the ratio of the absolute value of the mean to the RMS of the increments is large are south of Western Australia (120°E), west of New Zealand and in the region of the south equatorial current. This may be due, for example, to errors in the MSL (Figure 1), used to convert model sea-level to SLA, or due to systematic errors in the surface fluxes.

Regions where the mean increment is close to zero but where the RMS of the increments is large, for example off SE Australia in the Tasman Sea, probably correspond to areas where the circulation is poorly constrained. This may be due to problems with the initialisation, where the introduced changes are quickly “forgotten” by the model, or where the circulation is highly non-linear and therefore quite unpredictable and even chaotic. Certainly, the circulation in the Tasman Sea is very complex (Figure 7), with a vast array of warm- and cold-core eddies and meanders. It is therefore reasonable to expect that

the circulation there might be difficult to constrain with a 7-day assimilation cycle.

During the first BRAN experiment (BRAN1.0), described by Oke et al. (2005), we simply updated the full model state (including currents) in a single time step. This resulted in a large artificial response that degraded the solution. This result provided part of the motivation for adopting a more moderate nudging for the initialisation of BRAN1.5. In section 4, we showed that the errors of SST and SLA in BRAN1.5 are half as large as those in BRAN1.0. We attribute part of this improvement to the changes in initialisation. However, as noted above, we still regard initialisation as a key area in which the system could be improved.

Statistical comparisons, such as those presented in Figures 15 and 16, are particularly sensitive to the precise locations of mesoscale features like eddies, meanders and fronts. If, for example, an eddy is misplaced by its diameter, then the magnitude of the error can be up to twice the magnitude of the signal associated with the eddy. For this reason, a climatological estimate that contains no eddies at all, might be expected to have smaller RMSEs than an estimate that contains eddies that are not in the right place. The results in Figures 15 and 16 show clear evidence of this. This issue was recognised in the numerical weather prediction literature and resulted in the introduction of alternative measures of skill to assess categorical forecasts (e.g., Gandin and Murphy 1992). An example of a categorical forecast provides a yes or no answer to a question like: is there a cold-core eddy in region X on day Y? The comparisons in Figures 9 and 10 demonstrate that climatology is not useful for making a categorical forecast. By contrast, the comparisons in Figures 7, 8, 9 and 10 showing qualitative agreement between mesoscale features in BRAN

and independent observations, lead us to believe that the BRAN system may be useful for categorical forecasts, although we do not assess this here.

6 Summary and Conclusions

The current status of the Bluelink data assimilation system, BODAS, is described. BODAS uses an EnOI scheme to assimilate observations of various types in a single analysis step. This contrasts to many other data assimilation schemes that require special treatment to assimilate both satellite derived SLA and in situ T and S (e.g., Cooper and Haines 1996; Troccoli and Haines 1999; Segschneider et al. 2000; Fox et al. 2002; Guinehut et al. 2004; Cummings 2005; Martin et al. 2007; Chassignet et al. 2007). We provide a comprehensive description of BODAS, describe several positive characteristics of the assimilation method and identify the components of the system that are presently limiting the system's performance.

We provide a comprehensive assessment of BRAN. We demonstrate that, together with the main components of the global ocean observing system (altimetry, SST, Argo, TAO and XBT), BODAS can constrain an eddy-resolving ocean model. As far as we are aware, this is the first quantitative demonstration that an EnOI-based system can realistically constrain such a model. Although we note that there are several examples, using the SEEK filter (Pham et al. 1998), which is similar to EnOI, that provide a similar demonstration (e.g., Penduff et al. 2002; Brankart et al. 2003; Testut et al. 2003).

We present a series of qualitative and quantitative comparisons between BRAN and satellite and in situ observations. Through these comparisons we demon-

strate that BRAN represents both the broad-scale ocean circulation and the mesoscale ocean circulation, including features like eddies, meanders and fronts. Such a demonstration is one of the primary goals of the international GODAE community (Le Traon et al. 2001).

Based on the model-data comparisons we present in this study, we conclude that in the region around Australia, BRAN SLA is within 6-12 cm of observations over the open ocean and within 4-7 cm along the coast; and that BRAN SST is within 0.4-0.9° of observed SST. We also conclude that BRAN subsurface T and S are typically within 1° and 0.15 psu of observed fields in most of the Australian region; and that near-surface currents are typically within 0.2 m s⁻¹ of observations.

We also present a comprehensive evaluation of the predictive skill of the Bluelink system for SLA and SST. We demonstrate that the model has positive skill for each of these variables over a typical 7-day cycle. Specifically, we show that the RMSEs of a 7-day “forecast” are generally less than the RMSEs of the persistence of the initialised fields. This means that given some initial condition, the model integration over 7 days yields a better state than simply allowing the initial conditions to persist over time. However, we find that the RMSEs of the initialised fields are significantly greater than the RMSEs of the analysed fields for SLA over the open ocean. This indicates that some of the details in the analyses are not “getting through” to the model during the initialisation step. We therefore conclude that the simple nudging approach we use to initialise the model can be ineffective in some regions. As a result, we identify initialisation as one of the key areas in which we need to improve.

Our analysis of the predictive skill of SST reveals a significant seasonal de-

pendence. We attribute this to limitations of the mixed layer model that we use here, but note that it may also be related to seasonal biases in the surface fluxes. Another explanation could be that there is a seasonal dependence of the non-linearity of the ocean, with more energetic eddies at different times of the year resulting in a seasonal dependence to the predictability of the mesoscale ocean circulation.

We demonstrate that OFAM has positive skill for coastal sea level compared to analysed, initialised and observed persistence. We attribute the skill for coastal sea level to the accuracy of the wind stress and the model's response to the forcing.

We analyse the statistics of increments to sea-level and find that there are some regions where the model exhibits a significant bias. The reasons for these biases are unclear, though we suspect they are related either to limitations of the model parameterisations, inconsistency between the model's MSL and the reference MSL used for altimetric SLA, errors in surface fluxes or sub-optimal initialisation.

The Bluelink forecast system, of which BODAS is a key component, became operational at the BoM in August 2007. It now produces two forecasts each week, using operational fluxes from the Bureau's global numerical weather prediction system and assimilation of near-real-time observations. The results presented in this paper demonstrate the current potential of this system if we continue to run for many years in its present configuration. Analyses of BRAN, however, indicate that the performance could be significantly improved if the initialisation step could be made to constrain the model fields closer to the observations in a dynamically consistent way.

Acknowledgments. Financial support for this research is provided by CSIRO, the Bureau of Meteorology and the Royal Australian Navy as part of the Bluelink project, and the US Office of Naval Research (grant no. N00014-07-1-0422) . The authors also acknowledge the technical assistance of Russell Fiedler, Jim Mansbridge, Pavel Sakov and the other members of the Bluelink team. Satellite altimetry is provided by NASA, NOAA and CNES. Drifter data are provided by NOAA-AOML and SST observations are provided by NOAA and Remote Sensing Systems. Coastal sea-level around Australia is provided by the National Tidal Centre and state authorities.

References

- [] Allen, J. S., P. A. Newberger and J. Federiuk (1995) Upwelling circulation on the Oregon continental shelf. Part I: Response to idealized forcing. *J. Phys. Oceanogr.* **25**, 1843-1866.
- [] Allen, J. S., and P. A. Newberger (1996) Downwelling circulation on the Oregon continental shelf. Part I: Response to idealized forcing. *J. Phys. Oceanogr.* **26**, 2011-2035.
- [] Bell, M. J., R. M. Forbes, and A. Hines (2000) Assessment of the FOAM global data assimilation system for real-time operational ocean forecasting. *J. Mar. Syst.*, **25**, 1-22.
- [] Bloom, S. C., L. L. Takas, A. M. Da Silva and D. Ledvina (1996) Data assimilation using incremental analysis updates. *Mon. Weath. Rev.*, **124**, 1256-1271.
- [] Brankart, J. M., C. E. Testut, P. Brasseur and J. Verron (2003) Implementation

of a multivariate data assimilation scheme for isopycnic coordinate ocean models: Application to a 1993-96 hindcast of the North Atlantic Ocean circulation. *J. Geophys. Res.*, **108**, 3074, doi :10.1029/2001JC001198.

[] Brasseur, P., P. Bahurel, L. Bertino, F. Birol, J.-M. Brankart, N. Ferry, S. Losa, E. Remy, J. Schrter, S. Skachko, C.-E. Testut, B. Tranchant, P.-J. van Leeuwen, J. Verron (2006) Data assimilation in operational ocean forecasting systems: the MERCATOR and MERSEA developments. *Q. J. Roy. Meteor. Soc.*, **131**, 3561-3582.

[] Brassington, G. B., T. Pugh, C. Spillman, E. Schulz, H. Beggs, A. Schiller and P. R. Oke (2007) BLUElink: Development of operational oceanography and servicing in Australia, *J. Res. Pract. Inf. Tech.*, **39**, 151-164.

[] Burgers, G., M. A. Balmaseda, F. C. Vossepol, G. J. van Oldenborough, P. J. van Leeuwen (2002) Balanced ocean-data assimilation near the equator. *J. Phys. Oceanogr.*, **32**, 2509-2519.

[] Carton, J. A., G. Chepurin and X. Cao (2000) A simple ocean data assimilation analysis of the global upper ocean 1950-95. Part I: Methodology. *J. Phys. Oceanogr.*, **30**, 294-309.

[] Chassignet, E. P., H. E. Hurlburt, O. M. Smedstad, G. R. Halliwell, P. J. Hogan, A. J. Wallcraft, R. Baraille, and R. Bleck (2007) The HYCOM (HYbrid Coordinate Ocean Model) data assimilative system. *J. Marine Syst.*, **65**, 60-83.

[] Chen, D., L. M. Rothstein and A. J. Busalacchi (1994) A hybrid vertical mixing scheme and its application to tropical ocean models. *J. Phys. Oceanogr.*, **24**, 2156-2179.

[] Cooper, M. and K. Haines (1996) Altimetric assimilation with water property conservation. *J. Geophys Res*, **101**, 10591077.

- Cummings, J. A., (2005) Operational multivariate ocean data assimilation. *Q. J. Roy. Met. Soc.*, **131**, 3583-3604.
- Daley, R., (1991) *Atmospheric Data Analysis*, Cambridge University Press, 457pp.
- Ducet, N., P.-Y. Le Traon and G. Reverdin (2000) Global high-resolution mapping of ocean circulation from TOPEX/Poseidon and ERS-1 and -2. *J. Geophys. Res.*, **105**, 19477-19498.
- Etienne, H., and M. Benkiran (2007) Multivariate assimilation in MERCATOR project: New statistical parameters from forecast error statistics. *J. Marine Syst.*, **65**, 430-449.
- Evensen, G., (2003) The Ensemble Kalman Filter: theoretical formulation and practical implementation. *Ocean Dyn.*, **53**, 343-367.
- Feng, M., S. Wijffels, S. Godfrey and G. Meyers (2005) Do eddies play a role in the momentum balance of the Leeuwin current? *J. Phys. Oceanogr.*, **35**, 964-975.
- Fox, D. N., W. J. Teague, C. N. Barron, M. R. Carnes and C. M. Lee (2002) The Modular Ocean Data Assimilation system (MODAS). *J. Atmos. Oceanic Technol.*, **19**, 240-252.
- Gandin, L. S., and A. H. Murphy (1992) Equitable skill scores for categorical forecasts. *Mon. Weath. Rev.*, **120**, 361-370.
- Gaspari, G., and S. E. Cohn (1999) Construction of correlation functions in two and three dimensions. *Quart. J. Roy. Meteor. Soc.*, **125**, 723-757.
- Godfrey, S. J., G. R. Cresswell, T. J. Golding and A. F. Pearce (1980) The separation of the east Australian current, *J. Phys. Oceanogr.*, **10**, 430-440.
- Griffies, S. M., and R. W. Hallberg (2000) Biharmonic friction with a Smagorinsky viscosity for use in large-scale eddy-permitting ocean models. *Mon. Weath. Rev.*, **128**, 2935-2946.

- Griffies, S. M., R. C. Pacanowski, and A. Rosati (2004) A technical guide to MOM4. GFDL Ocean Group Technical Report No. 5. NOAA/Geophysical Fluid Dynamics Laboratory, 371 pp.
- Guinehut, S., P. Y. Le Traon, G. Larnicol, S. Philipps (2004) Combining Argo and remote-sensing data to estimate the ocean three-dimensional temperature fields - a first approach based on simulated observations. *J. Mar. Syst.*, **46**, 85-98.
- Hamill, T. M., J. S. Whitaker, and C. Snyder (2001) Distance-dependent filtering of background error covariances in an ensemble Kalman Filter. *Mon. Weath. Rev.*, **129**, 2776-2790.
- Houtekamer, P. L., and H. L. Mitchell (2001) A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Weath. Rev.*, **129**, 123-137.
- Huang, X.-Y., K. S. Mogensen, and X. Yang (2002) First-guess at the appropriate time: The HIRLAM implementation and experiments. *Proc. HIRLAM Workshop on Variational Data Assimilation and Remote Sensing*, Helsinki, Finland, Finnish Meteorological Institute, 2843.
- Kamachi, M., T. Kurangano, H. Ichikawa, H. Nakamura, A. Nishina, A. Isobe, D. Ambe, M. Arai, N. Gohda (2004) Operational Data Assimilation System for the Kuroshio South of Japan: Reanalysis and Validation. *J. Oceanogr.*, **60**, 303-312.
- Kampf, J., M. Doubell, D. A. Griffin, R. L. Matthews and T. M. Ward (2004) Evidence of a large seasonal coastal upwelling system along the southern shelf of Australia. *Geophys. Res. Lett.*, 31: L09310, doi:10.1029/2003GL019221.
- Kundu, P. K., (1976) Ekman veering observed near the ocean bottom, *J. Phys. Oceanogr.*, **6**, 238-242.
- Leslie, L. M., J. F. LeMarshall, R. P. Morison, C. Spinoso, R. J. Purser, N. Pescod, and R. Seecamp (1998) Improved hurricane track forecasting from a continuous

assimilation of high quality satellite wind data. *Mon. Weath. Rev.*, **126**, 1248-1257.

- Le Traon, P. Y., M. Reinecker, N. R. Smith, P. Bahurel, M. Bell, H. Hurlburt and P. Dandin (2001) Operational oceanography and prediction: A GODAE perspective. In: Koblinsky, C. J. and Smith, N. R., Editors, 2001. Observing the Oceans in the 21st Century, GODAE Project Office, Bureau of Meteorology, Melbourne, pp. 529-545.
- Levitus, S., (2001) World Ocean Database. NOAA Professional Paper, 13.
- Lewis, R. K., (1981) Seasonal upwelling along the south-eastern coastline of south Australia. *Aust. J. Mar. Freshwater Res.*, **32**, 843-854.
- Lynch, P., and X.-Y. Huang (1992) Initialization of the HIRLAM model using a digital filter. *Mon. Weath. Rev.*, **120**, 1019-1034.
- Martin, M. J., A. Hines and M. J. Bell (2007) Data assimilation in the FOAM operational short-range ocean forecasting system: a description of the scheme and its impact. *Q. J. Roy. Met. Soc.*, **133**, 981-995.
- McPhaden, M. J., and co-authors (1998) The Tropical Ocean Global Atmosphere (TOGA) observing system: a decade of progress. *J. Geophys. Res.*, **103**, 14,169-14,240.
- Mitchell, H. L., P. L. Houtekamer and G. Pellerin (2002) Ensemble size, balance, and model-error representation in an ensemble Kalman filter. *Mon. Weath. Rev.*, **130**, 2791-2808.
- Moore, A. M., (1990) Linear equatorial wave mode initialization in a model of the tropical Pacific Ocean: An initialization scheme for tropical ocean models. *J. Phys. Oceanogr.*, **20**, 423-445.
- Nilsson, C. S., and G. R. Cresswell (1981) The formation and evolution of East Australian Current warm-core eddies. *Prog. Oceanogr.*, **9**, 133-183.

- Oke P. R., J. S. Allen, R. N. Miller, G. D. Egbert, and P. M. Kosro (2002) Assimilation of surface velocity data into a primitive equation coastal ocean model. *J. Geophys. Res.*, **107**, 3122, doi:10.1029/2000JC000511.
- Oke P. R., A. Schiller, D. A. Griffin and G. B. Brassington (2005) Ensemble data assimilation for an eddy-resolving ocean model of the Australian Region. *Q. J. Roy. Met. Soc.*, **131**, 3301-3311.
- Oke P. R., P. Sakov and S. P. Corney (2007) Impacts of localisation in the EnKF and EnOI: experiments with a small model. *Ocean Dyn.*, **57**, 32-45.
- Oke, P. R., and P. Sakov (2007) Representation error of oceanic observations for data assimilation. *J. Ocean. Atmos. Technol.*, in press.
- Oke, P. R., and A. Schiller (2007) Impact of Argo, SST and altimeter data on an eddy-resolving ocean reanalysis. *Geophys. Res. Lett.*, **34**, L19601, doi:10.1029/2007GL031549.
- Ourmieres, Y., J.-M. Brankart, L. Berline, P. Brasseur and J. Verron (2006) Incremental analysis update implementation into a sequential ocean data assimilation system. *J. Atmos. Ocean. Tech.*, **23**, 1729-1744.
- Penduff, T., P. Brasseur, C. E. Testut, B. Barnier and J. Verron (2002) Assimilation of sea-surface temperature and altimetric data in the South Atlantic Ocean : impact on basin-scale properties. *J. Mar. Res.*, **60**, 805-833.
- Pham, D. T., J. Verron and M. C. Roubaud (1998) A singular evolutive extended Kalman filter for data assimilation in oceanography. *J. Mar. Syst.*, **16**, 323-340.
- Reynolds, R. W., and T. M. Smith (1994) Improved global sea surface temperature analyses using optimum interpolation. *J. Climate*, **7**, 929-948.
- Ridgway, K. R., J. R. Dunn and J. L. Wilkin (2002) Ocean interpolation by four-dimensional weighted least squares- Application to the waters around Australasia. *J. Atmos. and Oceanic Tech.*, **19**, 1357-1375.

- Ridgway, K. R., and S. A. Condie (2004) The 5500-km-long boundary flow off western and southern Australia. *J. Geophys. Res.*, **109**, C04017, doi:10.1029/2003JC001921.
- Sakov, P., and P. R. Oke (2007) Implications of the form of the ensemble transformations in ensemble square root filters. *Mon. Weath. Rev.*, in press.
- Schiller, A., P. R. Oke, G. Brassington, M. Entel, R. A. S. Fiedler, D. A. Griffin, J. V. Mansbridge and K. R. Ridgway (2007) Eddy-resolving ocean circulation in the Asia-Australian region inferred from an ocean reanalysis effort. *Prog. Oceanogr.*, submitted.
- Segschneider, J. D. L. T. Anderson, J. Vialard, M. Balmaseda, T. N. Stockdale, A. Troccoli, and K. Haines (2000) Initialization of seasonal forecasts assimilating sea level and temperature observations. *J. Climate*, **14**, 4292-4307.
- Smedstad, O. M., H. E. Hurlburt, E. J. Metzger, R. C. Rhodes, J. F. Shriver, A. J. Wallcraft, A. B. Kara (2003) An operational Eddy resolving 1/16° global ocean nowcast/forecast system *J. Mar. Syst.*, **40-41**, 341-361.
- Smith, N. R., (2000) The global ocean data assimilation experiment. *Adv. Space Res.*, **25**, 1089-1098.
- Sprintall, J., M. Tomczak, 1992: Evidence of the Barrier Layer in the Surface Layer of the Tropics. *J. Geophys. Res.*, **97**, 7305-7316.
- Testut, C. E., P. Brasseur, J. M. Brankart, and J. Verron (2003) Assimilation of sea-surface temperature and altimetric observations during 1992-1993 into an eddy-permitting primitive equation model of the North Atlantic Ocean. *J. Mar. Syst.*, **40-41**, 291-316.
- Troccoli, A. and K. Haines (1999) Use of the Temperature-Salinity Relation in a Data Assimilation Context. *J. Atmos. Oceanic Tech.*, **16**, 2011-2025.

Table 1

Estimates of the instrument error of different observation platforms; and the range of values for ϵ_{RE} and ϵ_{age} used in BRAN1.5. The lower range of ϵ_{RE} is for observations in the high resolution region; and the upper range of ϵ_{RE} is generally for observations in the coarse resolution region. The lower range of ϵ_{age} is for observations that occur at the analysis time; and the upper range of ϵ_{age} is generally for observations made 5-days before or after the analysis time in regions where the model's variability is greatest (e.g., Tasman Sea).

Platform	ϵ_{instr}	ϵ_{RE}	ϵ_{age}
GFO	5 cm	2-18 cm	0-20 cm
T/P, Envisat, Jason	3 cm	2-18 cm	0-20 cm
CTG	3 cm	1-3 cm	0-4 cm
AMSR-E SST	0.25°	0-2°C	0-2°C
CTD/Argo/TAO - S	0.05 psu	0-1 psu	0-1 psu
CTD/Argo/TAO - T	0.1°	0-3°C	0-4°C
XBT	0.2°	0-3°C	0-4°C

Table 2

Statistical comparisons between BRAN and observations from surface drifting buoys for each quadrant around Australia for the period January 2003-June 2006. Statistics include the number of observations (n_{obs}) the number of drifters (n_{dr}), the magnitude of the complex cross-correlation $\rho(u, v)$ and the phase angles $\theta(u, v)$ (Kundu, 1976) between the observed and reanalysed currents, the RMSE of current speed ($|u|$, in m s^{-1}) and T (in Deg C).

	n_{obs}	n_{dr}	$\rho(u, v)$	$\theta(u, v)$	RMSE $ u $	RMSE T
NE	279715	481	0.53	0.7°	0.19	0.35
NW	98098	463	0.60	5.9°	0.23	0.81
SE	86025	155	0.29	-23.9°	0.20	0.63
SW	111249	158	0.37	-14.7°	0.19	0.63

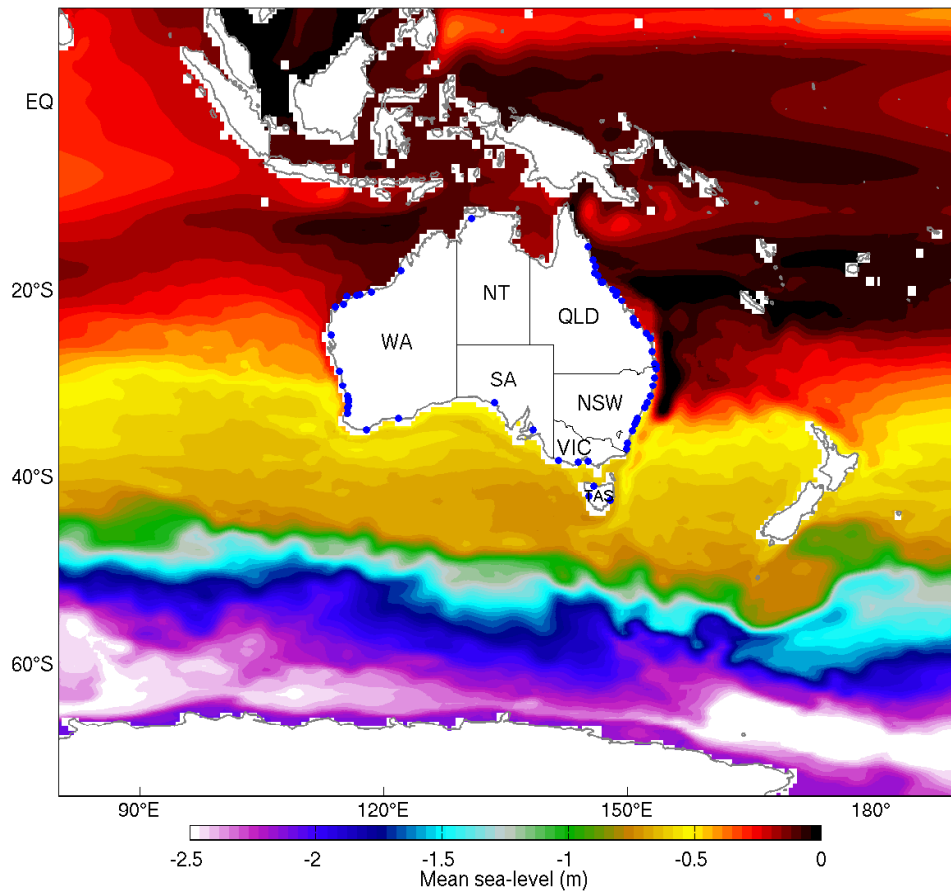


Fig. 1. Mean sea-level from the spin-up run that is used to convert model sea-level into SLA. Also shown are the state political boundaries and tide gauge locations (blue dots) that are relevant to Figure 14.

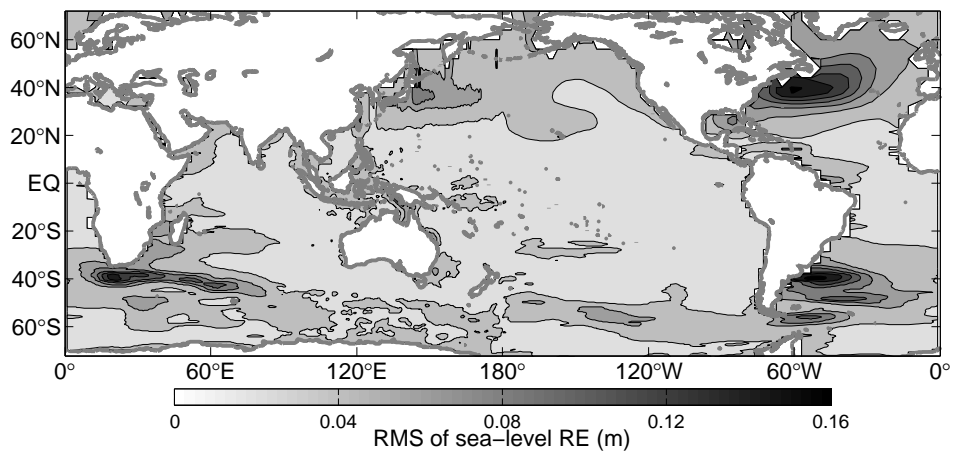


Fig. 2. Estimated RE for sea-level using the atSLA-based method described by Oke and Sakov (2007).

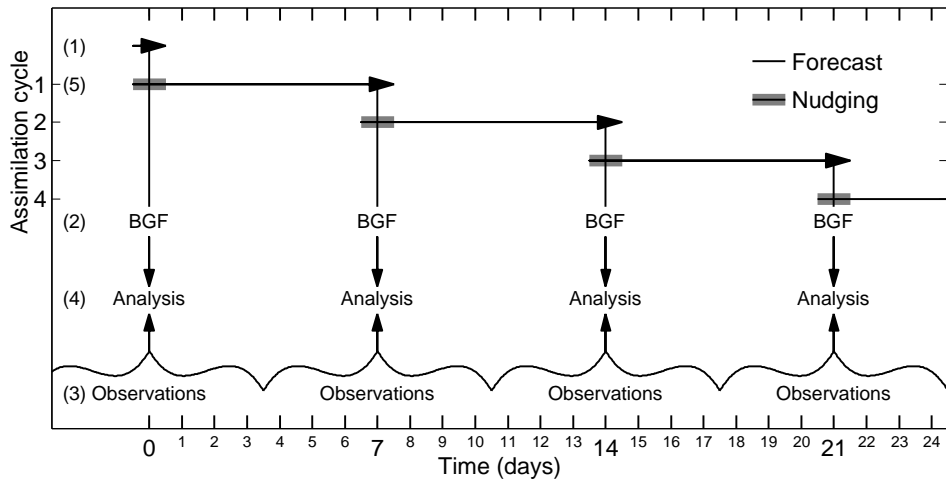


Fig. 3. Schematic of the 7-day assimilation cycle used during BRAN. The sequence of steps are denoted on the left of the figure showing (1) integrate previous forecast; (2) calculate background field (BGF; here a daily mean); (3) identify observations that fall in the observation time window (shown here as 7-days); (4) compute an analysis using BODAS; (5) nudge the model towards the analysis for one day and then integrate the next 7-day forecast.

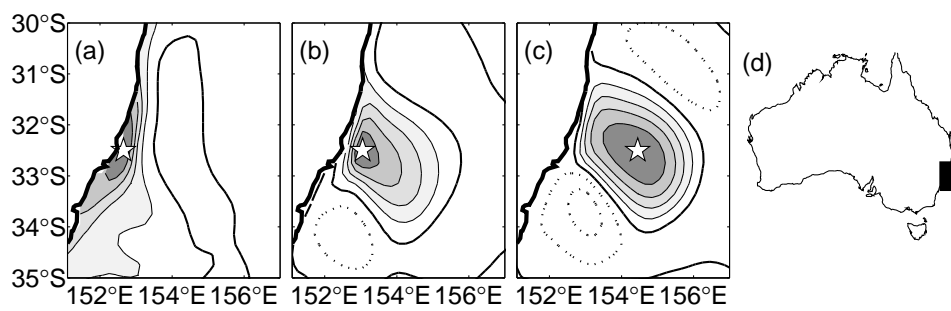


Fig. 4. Examples of the ensemble-based cross-correlations between sea-level at a reference location, denoted by the star, and sea-level in the surrounding region for a reference location (a) on the continental shelf, (b) over the continental slope and (c) over the deep ocean off eastern Australia (panel (d)). Contour intervals are 0.2; zero is bold, dotted is negative, correlations above 0.6 are shaded.

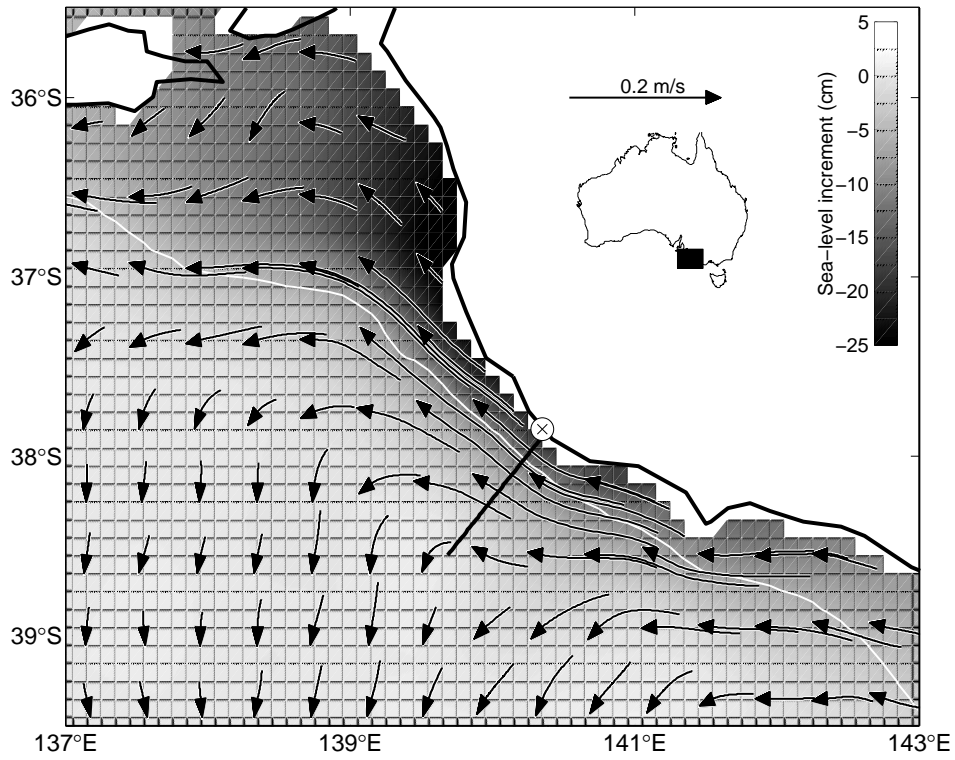


Fig. 5. Increments to sea-level (grey scale) and surface currents (vectors) based on a single observation of sea-level, denoted by the \otimes , such that the observed sea-level is 20 cm lower than the background sea-level. The inset shows the region of interest off south Australia.




Fig. 6. Increments to (a) sea-level, (b) across-shore currents (dashed contours are offshore; bold contour is zero) and (c) along-shore currents (solid contours are into the page; bold contour is zero); and the background field (solid contours) and analysis (dashed contours) for (d) T and (e) S, based on a single observation of sea-level at the coast as in Figure 5.

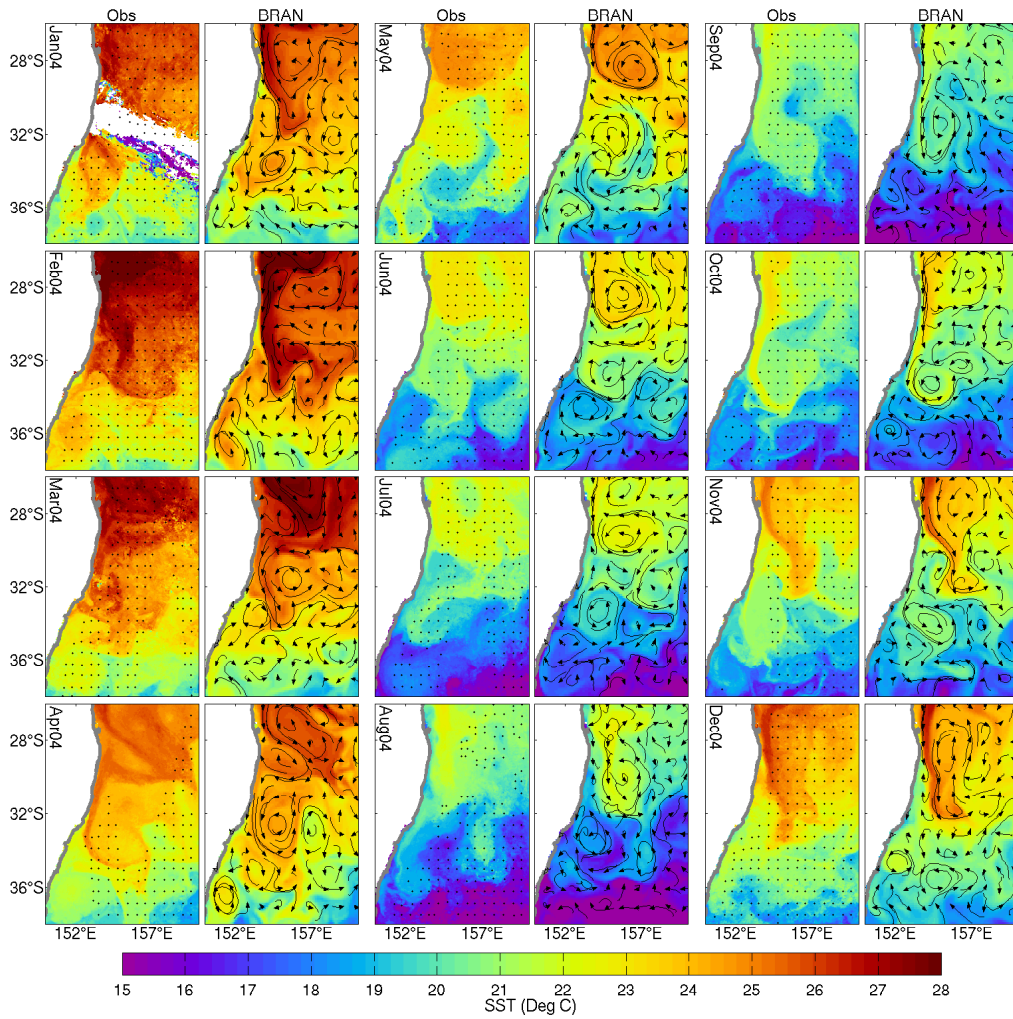


Fig. 7. Columns 1, 3 and 5 show 6-day composite SST from AVHRR measurements. Column 2, 4 and 6 show the corresponding 5-day averaged SST from BRAN. Comparisons are shown for the 15th day of each month of 2004, as labelled in the panels showing observations. Overlaying the BRAN SST are Lagrangian trajectories, computed from the time-dependent surface velocities from BRAN over the 5-day averaging period. Overlaying the observed AVHRR SST fields are dots denoting the locations of the assimilated, super-obsed AMSR-E SST data.

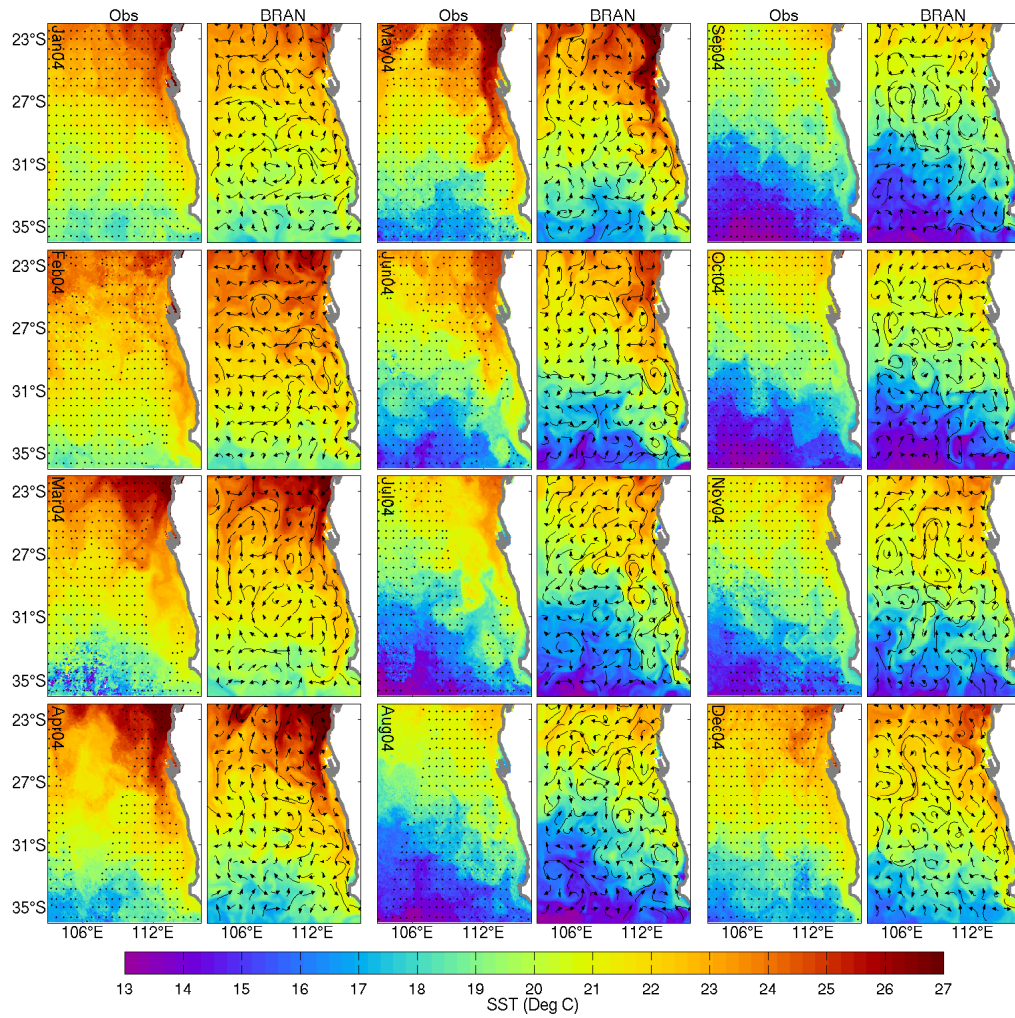


Fig. 8. As for Figure 7, except off south-west Western Australia.

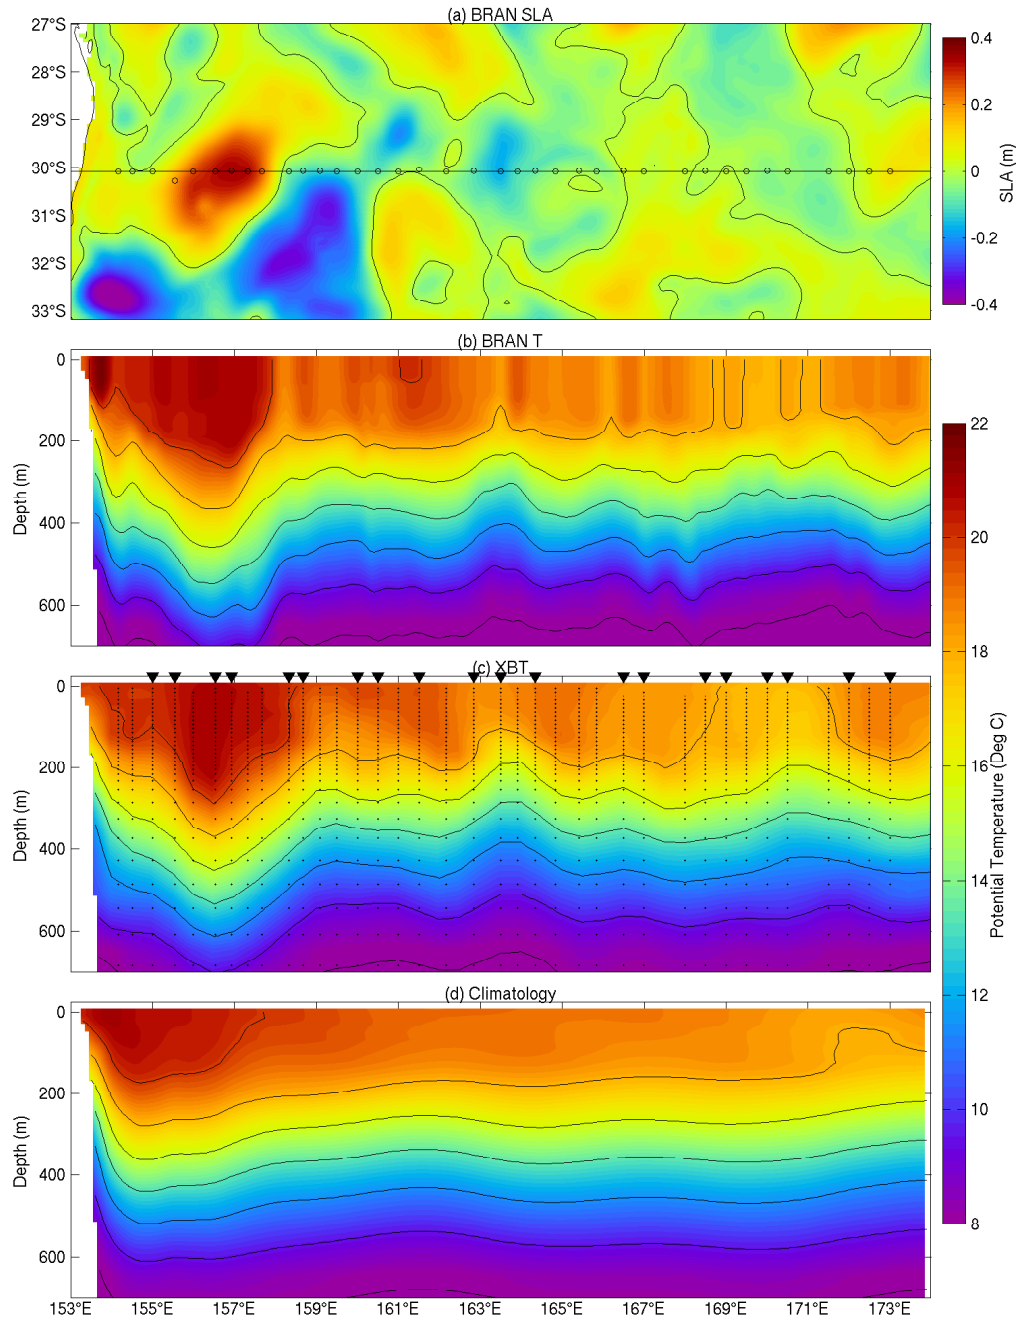


Fig. 9. Comparison between a observed, reanalysed and climatological T section off eastern Australia at 30°S. Panel (a) shows the daily mean SLA field in BRAN on August 8 2003; the solid line denotes the BRAN section and the circles denote the profile locations. Longitude versus depth sections are shown in panels (b-d) showing (b) daily mean T in BRAN on August 8 2003; (c) objectively analysed T from XBT observations between August 4-13 2003; and (d) T climatology (Ridgway et al. 2002). The dots in panel (b) denote the observation locations and the triangles indicate the profiles that are assimilated into BRAN. Contour intervals in panels (b-d) are 2°C; and the contour in panel (a) denotes zero SLA.




Fig. 10. As for Figure 9, except for a section off Western Australia along the IX15 line. BRAN fields are for July 7 2004; and XBT data are from July 6-8 2004.




Fig. 11. RMSE of SLA, compared to along-track Topex/Poseidon data (tracks are interleaved with Jason), plotted as a function of lead time for each quadrant of the Australian region as labelled in each panel. Results are shown for the forecast fields and for persistence of the analysed and initialised fields. These results are based on the period January 2003 to September 2005. For comparison, the RMSE of the spin-up run, with no data assimilation is off-scale at 0.12, 0.13, 0.14 and 0.17 m for the NW, NE, SW and SE regions respectively.




Fig. 12. As for Figure 11, except compared to AMSR-E data in the period January 2003 to June 2006. For comparison, the RMSE of the spin-up run, with no data assimilation, is off-scale at 1.2, 1.6, 1.3 and 1.3°C for the NW, NE, SW and SE regions respectively. Similarly, the RMS difference between the observed SST and SST climatology is 0.86, 0.91, 0.92 and 1.1° for the NW, NE, SW and SE regions respectively.




Fig. 13. As for Figure 12, except showing the RMSE of the forecast in each season; namely Austral summer (DJF), autumn (MAM), winter (JJA) and spring (SON).

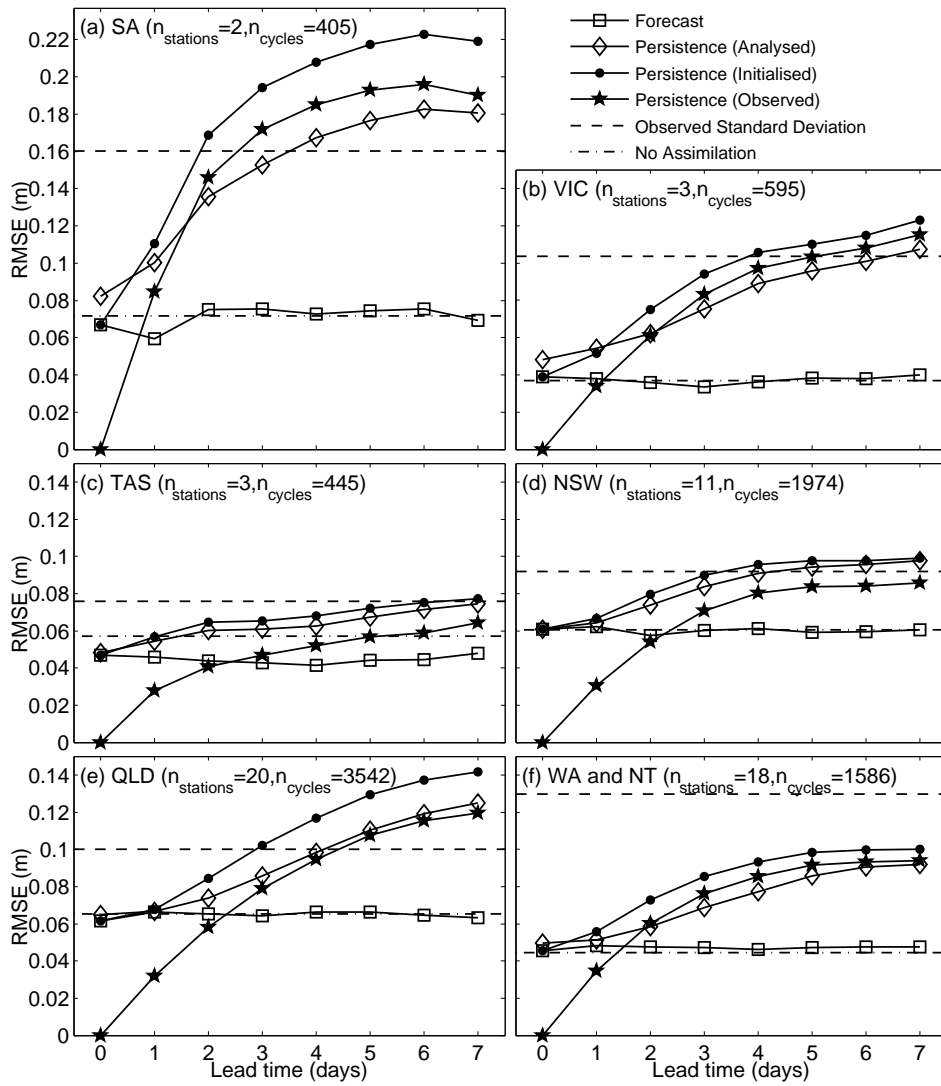


Fig. 14. As for Figure 11, except for SLA at coastal tide gauge stations around Australia for each state (see Figure 1 for political boundaries). Most tide gauge stations are not continuous for the entire BRAN1.5 period. Therefore the number of tide gauge stations (n_{stations}) in each state and the number of assimilation cycles (n_{cycles}) included in each calculation is shown in each panel.

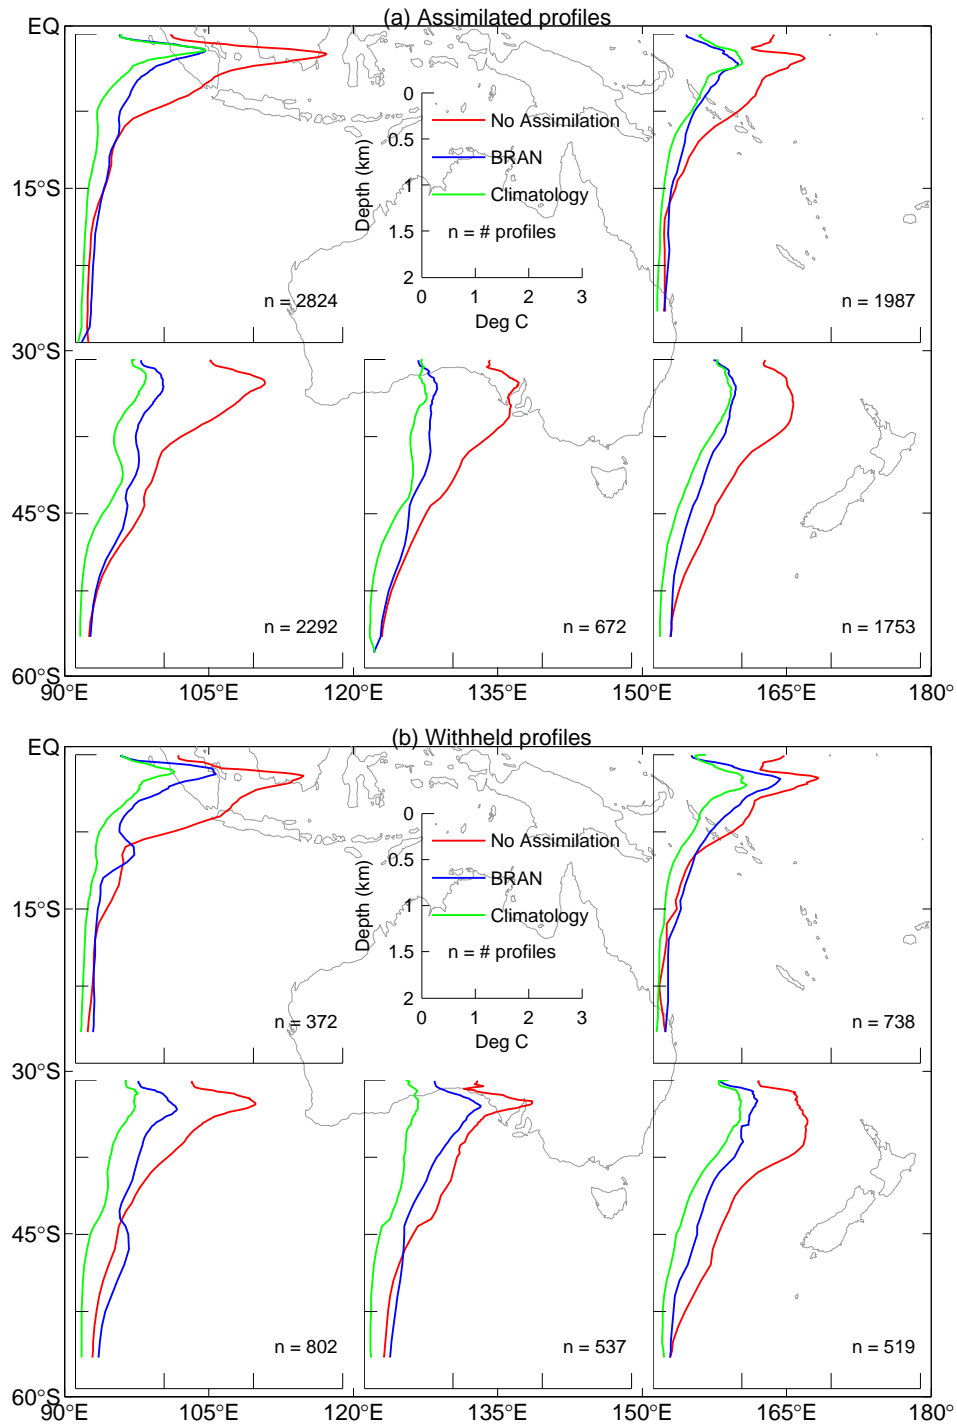


Fig. 15. Depth profiles of the RMSE between observed T profiles from Argo and the spin-up run (red; with no data assimilation), BRAN1.5 (blue) and climatology (green), using (a) assimilated profiles and (b) withheld profiles for the period January 2003 to December 2005 (the overlapping period for BRAN1.5 and the spin-up run). Plots are shown for each 30° by 30° region around Australia. The central panel over Australia shows the legend and the axis dimensions for all plots.

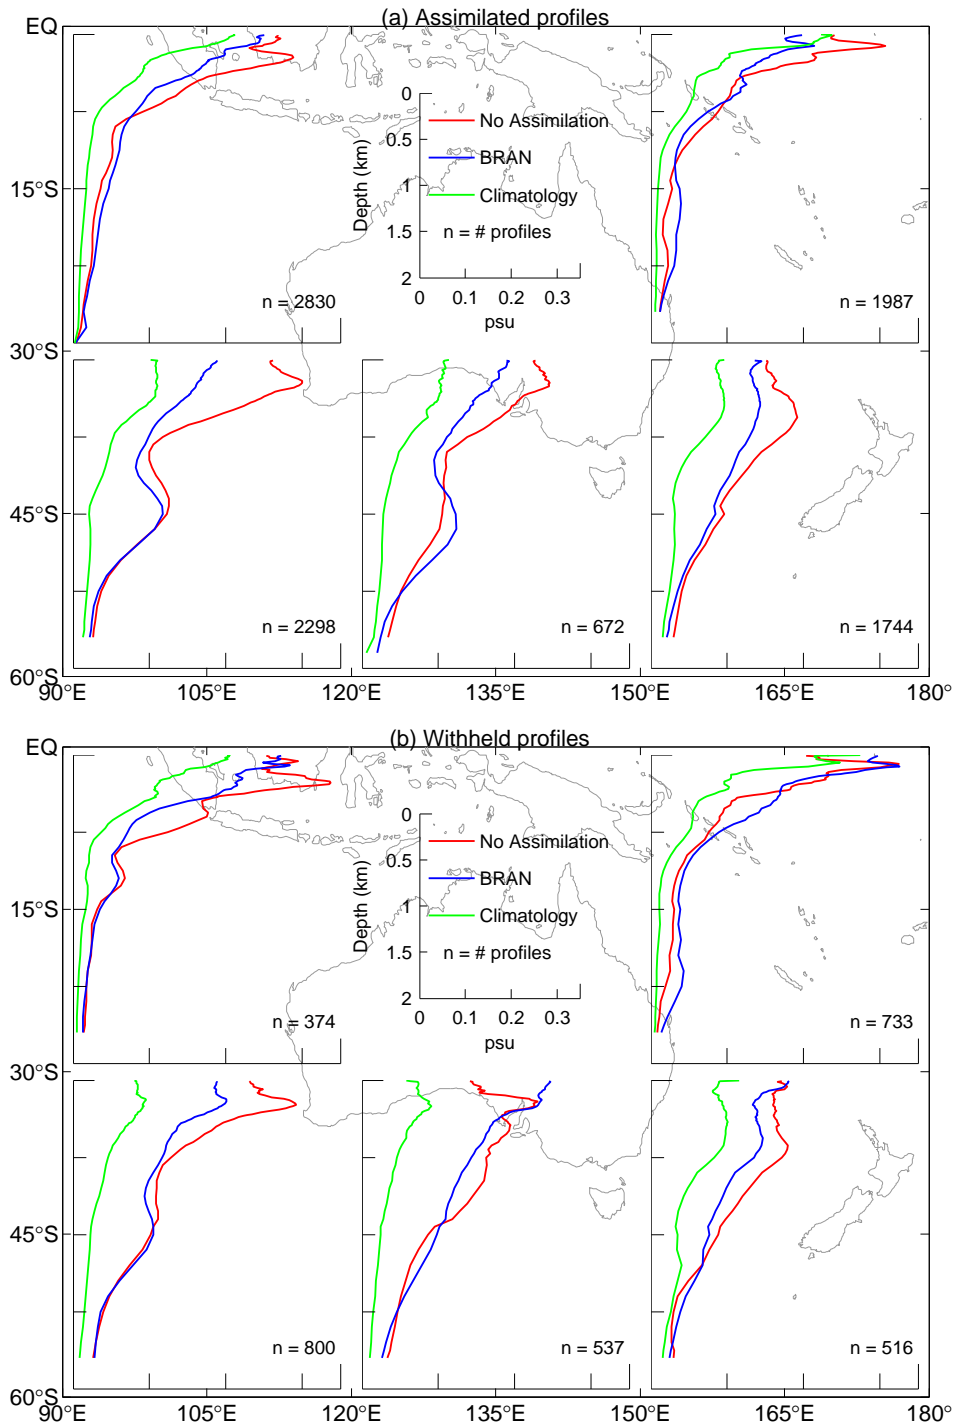


Fig. 16. As for Figure 15, except for S profiles from Argo.

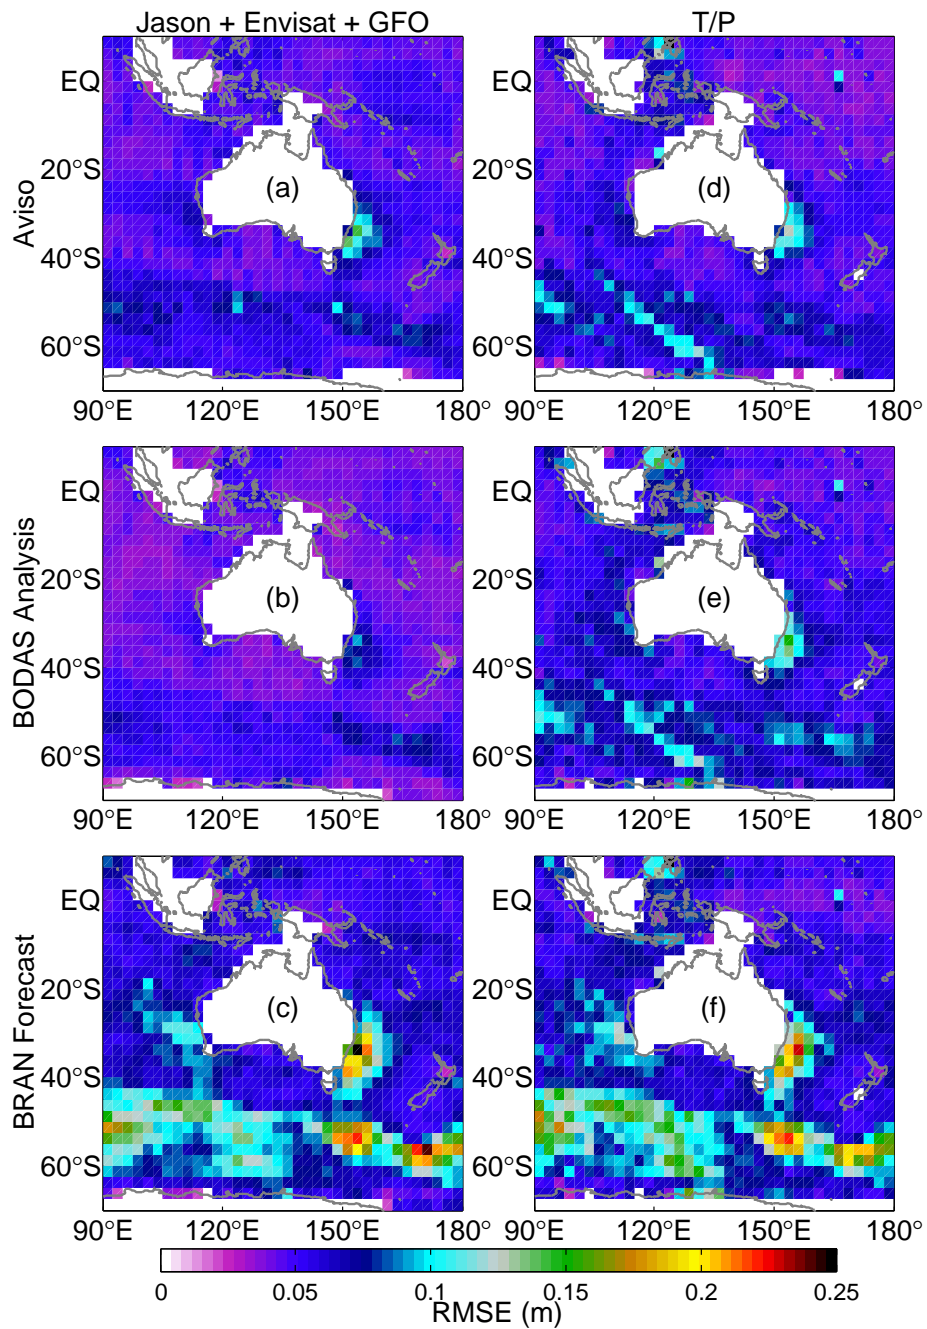


Fig. 17. RMSE between BRAN SLA and atSLA from (a-c) Jason, Envisat and GFO; and (d-f) T/P for (top-bottom) Aviso, BODAS, BRAN forecasts, as labelled to the left of each row. Calculations are made using all available atSLA observations for the period January 2003 to June 2006 (note that T/P ended in October 2005).

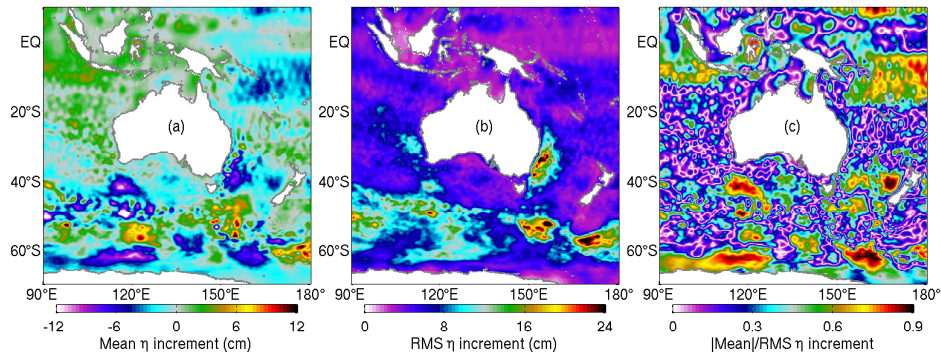


Fig. 18. Statistics of the increments to sea-level showing the (a) time-mean, (b) RMS and (c) the ratio of the absolute value of the time-mean to the RMS of the sea-level increments for the last 3-years of BRAN1.5 (June 2003 to June 2006).