

Interactive comment on “Recent development of the Met Office operational ocean forecasting system: an overview and assessment of the new Global FOAM forecasts” by E. W. Blockley et al.

E. W. Blockley et al.

ed.blockley@metoffice.gov.uk

Received and published: 8 April 2014

The authors would like to thank Gary for his review of this work and for taking the time to read through the paper in such detail.

Below is a response to Gary’s comments including some figures and the changes to the manuscript can be found in a GMD style paper uploaded as supplementary material.

Abstract

The FOAM system is described in the abstract and early section as inclusive of multiple models of varying resolutions. However the validation only focuses on the global system. As the focus is on the performance of the global system I suggest removing mention of the other systems from the abstract and retain in the introduction.

The paper itself is quite large and so we are conscious of the fact that to include everything would make it unmanageable. This is the reason that the regional models are not introduced fully and assessed in the paper. We agree that having mention of them in the abstract is misleading and have therefore removed mention of them from the abstract whilst keeping the brief mention in the introduction as you suggest.

The v12 system includes many significant changes. The control experiment chosen was to perform a free running model. It might have been more instructive to perform a data assimilative run with fewer of the changes to provide some comparison on what changes might be responsible for the positive and negative results.

Owing to the volume of information involved it has not been possible to present details of all of the RD work that has gone into the creation of this system. The following trials were performed in addition to those presented in this work:

1. bulk forcing vs. direct fluxes SBC using FOAM v12 system which show that the use of bulk formulae have a positive effect almost everywhere and in particular for the sea ice

- 2. NEMOVAR vs. OCNASM using FOAM v12 system which can be found detailed in Waters et al. (2014)*
- 3. LIM2 vs. CICE ice models using FOAM v11 system which show a significant improvement in the ice concentration fields.*

The results from these trials have been used to support some of the observations and discussion in Sections 4 and 5.

Section 2.1:

The instability is a numerical one....The problem is not just due to unrealistically high diffusion.

The comment about horizontal scaling of diffusion coefficient has been modified to explicitly mention 'numerical instability'.

Section 2.2:

Are only OGDR's used for altimetry, what data volume and coverage is assimilated on average from each platform in each of the two 24 hr analysis windows.

Yes only OGDRs are used for altimetry both operationally and for the reanalysis trials. At present, and after QC has been performed, the operational FOAM system is using approx. 60k observations per day during the (T-48h,T-24h] analyses and 40k observations per day during the (T-24h,T+0h] analyses on average. Coverage of tracks is global even during (T-24h,T+0h] (see example image for the observations used for the (T-24h,T+0h] analysis on 25th February this year - Figure 1) although, of course, there are gaps between tracks which are filled approximately every 12-15 days. Obviously data volume and coverage for SLA differ significantly depending on the state of the satellite observing system. However the statistics presented for the the reanalysis trials in Section 4.1 are calculated using an average of 47.5k observations per day (after QC filtering to a common subset).

Section 2.3:

Why is T-54h required for the atmospheric forcing?

The 'T-54h' in bullet 2. of Section 2.3 was a typo and has been changed to 'T-48h'.

What is the resolution of the atmospheric model used in the hindcasts? Could you add this to Table 1?

C2967

The NWP forcing fields used throughout the entire reanalysis trial period are from the UM Global system introduced in Section 2 which runs at approx. 25km resolution. Given that this is the case for all the reanalyses and the operational system that the paper is describing it would seem unnecessary to add it to the table. However the paper does not make it clear that the forcing for all 3 trials came from the same NWP model running at the same resolution throughout the trial period and so I have added a sentence to the start of Section 3 to explain this.

48-hour assimilation window

There seems to be some confusion relating to the 48-hour window that is used operationally and the 24-hour daily cycling used in the reanalyses. Given that Reviewer #2 also had some questions relating to this it is evident that the explanation provided needs to be made clearer.

When we run FOAM operationally in near-real-time we start from T-48h and produce our 'best estimate' analysis for the (T-48h,T-24h] period. The model state at T-24h is then saved for the next day as our best estimate of the ocean state. An update run is then performed for the period (T-24h,T+00h] before we start the 7-day forecast. There is no doubling up or extra weight applied to any observations - we have simply moved our 'best estimate' back by 24-hours by running an analysis for the period (T-48h,T-24h]. The impact of running a second analysis period therefore is to allow more observations to be used in the generation of our best estimate initial conditions each day owing to the late arrival of some observations. As mentioned in Section 2.4.2 (now Appendix A) this led to a RMS error reduction of approx 5-6% globally in the operational near-real-time system when it was implemented. To make this clearer in the paper I have modified the explanation of the 48-hour window in Section 2.3 to better distinguish between the 'best estimate' analysis and the 'update run'.

C2968

We do not use this approach for the reanalyses because they are run in delayed time (typically > 6 months after real-time) and so have access to more observations than the operational system because late arrivals will be included. These reanalyses are not an identical copy of the operational suite and are used in a calibration sense to understand the potential errors in the operational system.

I have therefore expanded the last sentence at the end of the 1st paragraph of Section 3 to explain that we do not need to run 2 separate assimilation cycles in the reanalyses because the observations are extracted in delayed-time and are therefore more plentiful than those used operationally.

Does Waters et al., show whether FGAT matters over a 24h time window?

Waters et al. (2014) do not investigate what impact the use of FGAT has on the FOAM system.

“reanalysis observations are filtered. . .” do you mean sorted to a common subset?

Yes I do! The text has been changed to explicitly say ‘common subset’.

Bar-chart figures

The scale for the plots is being set by the free run.... I suggest that the x-axis be reduced and where the free model goes beyond the limit place a bracketed value above the line to indicate its value.

Figures 1 and 2 have been re-worked according to your suggestions which have improved the readability of the figures. In all plots where the maximum error for either of the v12/v22 runs is less than 2/3 of the free run, the x-axis scale is truncated so the finer detail can be seen for the assimilative runs. Where this happens the error value

C2969

for the free run is annotated to the plot above and at the end of the corresponding error bar. Those plots that have changed can be found in Figure 2 of this response and in the updated paper.

Profile error figures

It would be instructive to see the tropical Pacific as an additional 2 panels to shed further light on what the distribution of error is in this region.

Figure 3 has been extended to include temperature and salinity error profiles for the Tropical Pacific and the North Atlantic. The figure is now also larger to make the detail more visible.

This can be seen in the supplementary updated paper or in figure 3 below.

Salinity profile errors

The paper indicates that there is a general deterioration in the performance of v12 in the ocean interior. This requires some further discussion.

A comment has been added to Section 4.1.3 to explain the general deterioration in the salinity fields in the ocean interior at v12. This is related to the difficulties fitting the relatively sparse salinity observations with short correlation length-scales and it is hoped that the adoption of dual length-scales will improve things in the ocean interior.

Previous FOAM systems have converted altimetry into synthetic profiles. Is this still the case?

FOAM does not convert altimeter observations into synthetic profiles.

The large bias in the North Atlantic is associated with a large cool bias in the temperature. Whilst the Mediterranean is associated with a warm bias. Is this consistent with

C2970

the precipitation hypothesis?

The text has been changed from 'believed to be caused by excessively high precipitation in the surface forcing fields' to 'believed to be an artefact of the increased number of coastal observations in these areas'. The original text was a little misleading and suggested that precipitation biases are worse in these areas which is not necessarily the case. Rather these biases have more impact in these regions owing to the number of coastal salinity observations in these regions.

NEMO is a volume conserving model. The long term drift in the model needs to be explained by a change in volume rather than one in terms of steric expansion such as low density water. A volume conserving model does adjust sea level for steric effects but this is done without any net change to the global volume.

Yes this was misleading and has been changed to simply say that the SSH increase is caused by a mismatch between the precipitation and riverine freshwater inputs.

Near-surface velocities

It is worth emphasising that this verification is based on independent observations.

A sentence has been added to the 1st paragraph of Section 4.1.6 to emphasise that the velocity validation is based on independent data.

How many observations are being used in each basin?

Approximately 725 drifter observations per day are used globally for the assessments in Section 4.1.6. The 1st paragraph of Section 4.1.6 has been modified to include this number.

For your information the regional breakdown is roughly 300 per day in each of the Atlantic and Pacific Oceans, 65 per day in the Indian Ocean and 150 per day in the

C2971

Southern Ocean. (NB. these figures sum to more than 725 because the Southern Ocean domain overlaps with each of the other 3 to some extent.)

Why is the Indian Ocean excluded?

I am not sure I follow your comment re. the exclusion of the Indian Ocean. The Indian Ocean is included in the velocity validation and is mentioned in the text (Section 4.1.6). In particular it is noted that v12 velocities are better in the Indian Ocean in contrast to the Tropical Pacific and Tropical Atlantic where v11 is better.

If the question is why is it not included in the Taylor plots in Figure 5 then this is because I think that 4 regions (with 3 points per region) is the most that can fit on these Taylor plots without making them too complicated. As described in the text, the results were generally consistent across the regions – being better in all extra-tropical regions and worse in the Tropical Pacific Tropical Atlantic. It was therefore decided that we would only show results from the global ocean and 1 region for each of the tropical, mid-latitude, and high-latitude regimes (namely Global, Tropical Pacific, North Atlantic Southern Ocean) to illustrate this. Results for all the 8 regions (apart from global) can be found in Figure 7 below for your interest.

For the Tropical Pacific both the meridional and zonal components seem to be inferior to v11 which is not consistent with the text.

This is explicitly mentioned in the text of Section 4.1.6 as follows: 'Although better in the Indian Ocean the v12 system is worse elsewhere in the tropics; in particular in the Tropical Pacific. Further comparisons with currents measured by the TAO/TRITON (McPhaden et al., 1998) and PIRATA (Servain et al., 1998) tropical moorings (not shown) confirm the findings of the drifter regional results that the skill of current predictions is reduced in the Tropical Pacific and Tropical Atlantic.'

C2972

Is this the DA u,v fields or the initialisation shock of the T/S/eta state? Do you track KE during the initialisation?

This is the assimilation of T/S/SLA not the velocity balancing. The latter is geostrophic and therefore effectively zero near the equator. In the text I have replaced 'increments' with 'tracer increments' to make this clearer.

We do not track KE through the IAU as standard but have performed a few quick runs to check and can see no evidence of shock.

Increase fontsize – difficult to read in this scaled down version. Also use the full column width.

I have increased the size of the Taylor plots in Figure 5 as per your suggestion. Please see the end of this response regarding font sizes.

It is instructive to see errors in each component. However, it is also useful to perform the analysis on the total vector. For example, Kundu, 1976, JPO using a complex correlation.

Although the Kundu and Allen (1976) technique looks potentially interesting we would not be comfortable including this at this stage. This will be considered as a useful extension for future validation exercises.

Forecast validation

It would be instructive to compare the power spectrum of the analysis and the forecasts. Is there larger power in the high wavenumbers as speculated which is subsequently dissipated through the forecast period?

Yes we agree that this would be instructive and we plan to perform this sort of analysis in the future to guide development.

C2973

It is worth noting that the hypothesis of over mixing for the temperature biases does not seem to be present in the salinity results.

A comment has been added to Section 4.2.3 to explain that, although global salinity profiles do not show evidence of excessive mixing, the mixing bias is present in mid-latitude regions. An example for the North Atlantic can be found in Figure 4 of this response.

Comparisons with gridded observations (Section 4.3)

The example shown would benefit from the calculation of spatial correlations and included in the text to quantify the improvement of v12.

Yes we agree! Anomaly correlation have been calculated for the 2D fields shown in Fig. 10 over the Agulhas retroflection region which can be found in the new Table 4. These results support the conclusions of the qualitative assessments described in Section 4.3 that the v12 system provides a better representation of surface mesoscale fields. Subsequent calculation of spatial correlations for other case studies confirm the statement that, in general, the v12 fields show a better agreement with the observations.

There is a lot of material introduced for a single case study example. If supplementary material is permitted it would be desirable to show at least two other examples.

The Agulhas case study was chosen because there were some fairly obvious and interesting features present at this particular time.

I have included a figure and table analogous to Fig 10 Table 4 for a second case study covering the East Australia Current region. These can be found in Figures 5 and 6 of this document.

(This could be added to the paper as supplementary material although I cannot see a way to do that from here.)

C2974

A further extension to this (as suggested by Reviewer 2) has been to quantify the relative improvement to the mesoscale eddy fields in the new system by assessing the SLA and near-surface current fields separately for areas of high and low mesoscale variability. These high/low variability regions were defined using the spatial variability distribution of SLA observations over the 2 year assessment period by using a threshold standard deviation. These results can be seen in the new Table 3 and are discussed in Section 4.1.4. Results show that the improvement at v12 in areas of high variability is considerably more pronounced than for areas of low variability which is consistent with the findings presented in Section 4.3 and discussed throughout the paper.

Summary

The opening statement of the 2nd paragraph must state that the results are mixed. There are clear advantages when the observation density is high but for regions with sparse observations the performance has deteriorated.

Paragraph 2 of the Summary (Section 5) has been modified to say that results are mixed with considerable improvement where observations density is high but with some deterioration in areas of sparse observations.

There is no information/diagnostics presented on the initialisation shock such as global KE. Given the rapid deterioration in the T/S profiles in the forecast compared with v11 these diagnostics would be instructive.

We have not included any kinetic energy diagnostics in the paper but have performed a few short tests to output global KE at each model time step. These tests show no sign of any shock and no suggestion that the situation is different at v12 than at v11. So we are confident that the IAU is performing as expected. What is shown however is that the data assimilation is increasing the KE in the system (or rather preventing the global KE from decreasing). This is expected given that the 1/4 degree model is not eddy resolving and relies on the DA to provide some of the eddy variability.

C2975

The rapid deterioration in profile error is a perhaps a little exaggerated by the comparisons between forecasts and analyses. By comparing with the analysis daily-mean fields we are essentially comparing against a field where the observation in question has already been assimilated (certainly approx. half of the increments have been applied). Whereas for forecast day 1 the system will not have assimilated data from this instrument for a number of days (10 for Argo) and so the deterioration will have happened over the 10 days not 1 day and in that respect the lead-time plots. So this is partly owing to over-fitting but also to the sparsity of sub-surface observations (in time and space).

Figure font sizes

Reviewer #2 also commented on the relatively small font size used for the figures and so we acknowledge that this will need to be increased. The font sizes used in the

figures are either 18 or 20pt but it is the figure scaling employed by the typesetting that causes them to be smaller than the journal text. Ideally these should align with the font-size used for the figure caption. When the final typesetting is done we shall make sure, in conjunction with the journal, that the fonts used in the figures are clear and in keeping with GMD guidelines. We have not done so as part of this response because

the figure scaling used for this (single column) GMDD discussion document will most likely be different from that used in the (dual column) final version – meaning that any changes we made now may very well need to be redone at the typesetting stage.

Please also note the supplement to this comment:

<http://www.geosci-model-dev-discuss.net/6/C2965/2014/gmdd-6-C2965-2014-supplement.pdf>

Interactive comment on Geosci. Model Dev. Discuss., 6, 6219, 2013.

C2976

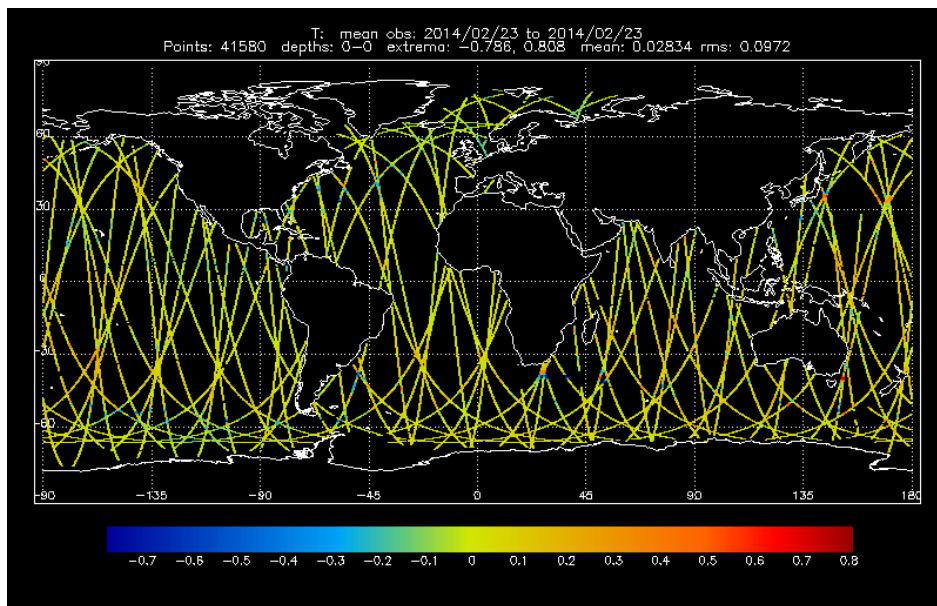


Fig. 1. Example daily altimeter observation coverage for the (T-24h, T+00h) period

C2977

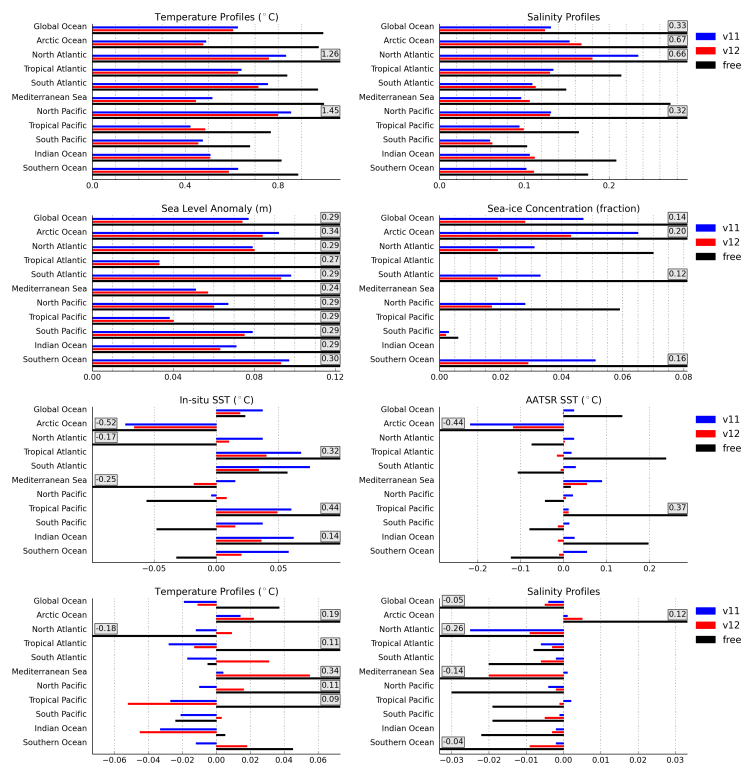


Fig. 2. The barchart plots that have been replaced in Figures 1 and 2.

C2978




Fig. 3. Mean error profiles against EN3 data for temperature (top row : °C) and salinity (bottom row : measured on the practical salinity scale) plotted against model depth (m) on a log scale for the Global Ocean (left), North Atlantic (centre) and Tropical Pacific (right) regions. Solid lines denote RMS errors and dashed lines denote mean errors for the v12 (red), v11 (blue) and free (black) trials. Mean errors are plotted as modelled-observed meaning that positive temperature (salinity) values indicate that the model is too warm (salty).

Fig. 3. New Figure 3 for paper including extra panels for the North Atlantic and Tropical Pacific

C2979




Fig. 4. North Atlantic salinity forecast error profiles analogous to those in Fig 7 of the paper.

C2980




Fig. 1. An array of monthly-mean gridded contour plots over the East Australia Current region (longitude: 148° E – 165° E; latitude: 22° S – 38° S) for December 2011. Sea surface height (left column : m) and temperature (centre column : °C) are plotted as coloured contours and overlaid with black contour lines. Surface currents (right column : m/s) are displayed as coloured contours of current intensity (speed) with white arrows overlaid to show direction. Output from the v12, v11 and free trials are plotted in the 1st, 3rd and 4th rows respectively whilst the 2nd row plots show the gridded observational products: AVISO SSH, OSTIA SST and OSCAR surface currents. Model currents shown (i.e. v12, v11 and free) are total integrated velocity over the top 15m. Anomaly correlations for each of the modelled fields against the corresponding gridded observations can be found in Table 1.

Fig. 5. Supplementary case study figure for the East Australia Current analogous to Fig 10 of the paper

C2981

Table 1. Anomaly correlations for each of the modelled SSH, SST and the magnitude of near-surface velocity fields (speed) against the corresponding gridded observational products for all the 2D spatial maps shown in Fig. 1. Correlations are calculated over all ocean points and anomalies are calculated relative to the WOA2001 1/4 degree analysis (Boyer et al., 2005) for SST, the CNES09 MDT (Rio et al., 2011) for SSH and the GDP drifter climatology (Lumpkin and Garraffo, 2005) for the speed.

	SSH	SST	CRT
v12	0.89	0.92	0.92
v11	0.86	0.87	0.82
free	0.13	0.61	0.35

Fig. 6. Anomaly correlations table for the EAC case study in Figure 5 analogous to Table 4 of the paper

C2982

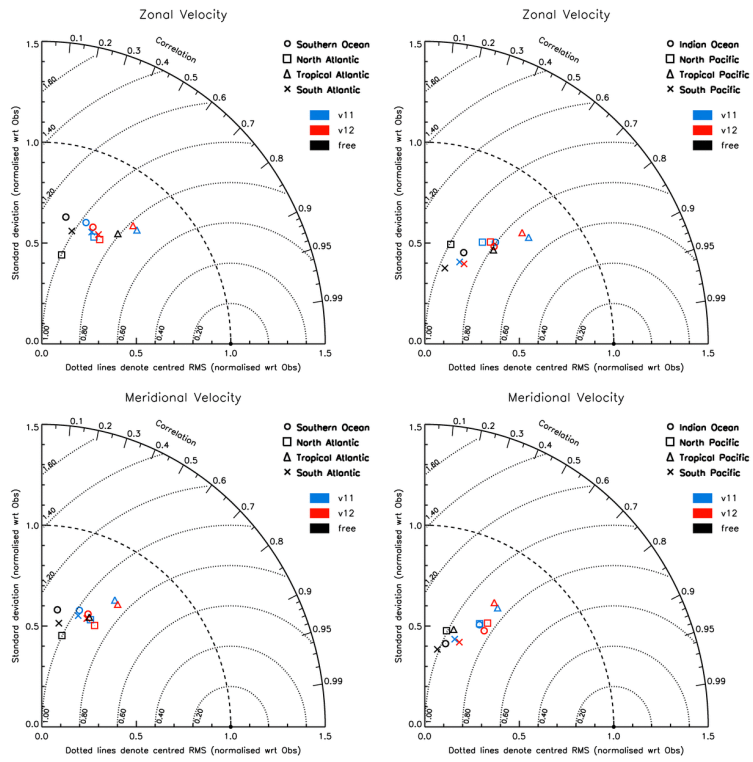


Fig. 7. Taylor plots analogous to those in Figure 5 of the paper but containing all regions apart from global