# SANGOMA: Stochastic Assimilation for the Next Generation Ocean Model Applications EU FP7 SPACE-2011-1 project 283580

## Deliverable 4.4:
## Running small and medium case benchmarks to test data assimilation methods

Due date: 31/10/2015

Delivery date: 31/10/2015

Delivery type: Report , public


Jean-Marie Beckers Alexander Barth
Yajing Yan
University of Liège, BELGIUM

Peter Jan Van Leeuwen
University of Reading, UK

Lars Nerger
Alfred-Wegener-Institut, GERMANY

Arnold Heemink Nils van Velzen
Martin Verlaan
Delft University of Technology, NETHERLANDS

Pierre Brasseur Jean-Michel Brankart Guillem Candille
CNRS-LEGI, FRANCE

Pierre de Mey
CNRS-LEGOS, FRANCE

Laurent Bertino
NERSC, NORWAY

# Chapter 1

# Introduction

This deliverable illustrates the implementations of the small and medium case SANGOMA benchmarks to experiment a wide panel of data assimilation methods, including those developed in Work Package 3 (Innovative data assimilation methods). The medium case benchmark in particular has been extensively used, leading to 6 scientific articles (published, submitted, or in preparation). In particular, due to its flexibility and ease of implementation, this benchmark has been adopted by non-SANGOMA partners to test other data assimilation tools.

The WP3 developments lead to use a few other benchmarks not initially proposed in the first place, in order to test the methods in a wide range of challenging situations. The completely new method from SANGOMA (the Multivariate Rank Histogram Filter – MRHF –, Metref et al, 2014) was better assessed with the strongly nonlinear Lorenz 63 system (Lorenz, 1963). It also turned out that lightweight (1D vertical) ecosystem models were appropriate to challenge and test non-Gaussian assimilation methods. Methods specifically designed for coastal applications must be tested with coastal ocean models; this deliverable thus reports experiments performed with the SYMPHONIE coastal model of the Bay of Biscay and the SANGOMA data assimilation tools. This model is proposed as another SANGOMA benchmark.

The SANGOMA metrics have not systematically been used with the small and medium case benchmarks. This is essentially due to (1) the relative timing of the tasks dedicated to the metrics and to these benchmarks, and (2) the quite unexpected (and relative) straightforwardness of the use of the metrics directly with the large case benchmark.

In what follows, the implementation descriptions are sorted according to the classes suggested above: Lorenz systems, medium case benchmark, ecosystem models, and the coastal application.

# Chapter 2

# Small case benchmark

## 2.1 Histogram filters (CNRS/LGGE)

Though not specifically introduced as a SANGOMA benchmark, the Lorenz 63 system is used as a classical testbed for non Gaussian data assimilation methods. In the work of Metref et al (2014), five ensemble methods are implemented: the EnKF, the particle filter (With Sequential Importance Resampling – SIR), the Rank Histogram Filter (Anderson, 2010), the MRHF, and an approximated version of the MRHF. To assess the methods' ability to deal with strongly non-Gaussian distributions, an experiment is set up where the model dynamics create multimodal state distributions (illustrated on Figure 2.1, top), and the observations do not dispel ambiguities about the position of the true state among the modes. In practice, this is done by observing only the "Z" Lorenz variable every 40 time steps of 0.01 units. To quantify the performance of each method, the Kullback-Leibler divergence is specifically introduced to measure a distance between the probability density resulting from the ensemble of the tested method, and the density resulting from a particle filter run with a large number of particles (4096). The Kullback-Leibler divergence writes:

$$KL(P,Q) = \int \log\frac{P}{Q} dP.$$

Figure 2.1 (bottom) shows the Kullback-Leibler distances for the "X" variable, and for the different methods using different ensemble sizes. With large enough ensembles, the fully non-Gaussian methods (particle filter and MRHF, in blue and red) behave much better than the others that rely on linear observational updates. The MRHF, though not as good as the particle filter with large ensembles, performs much better than the particle filter with moderate ensemble size. This is because the method is based on the transformation of particles (like the EnKF) rather than the sampling (like the particle filter). This is one reason why the MRHF is promising, and is further explored below with an ecosystem model.

## 2.2 Square root filters (AWI)

The small case benchmark was used to study the influence of nonlinearity on ensemble smoothing. In smoothers the state is estimated using observations from

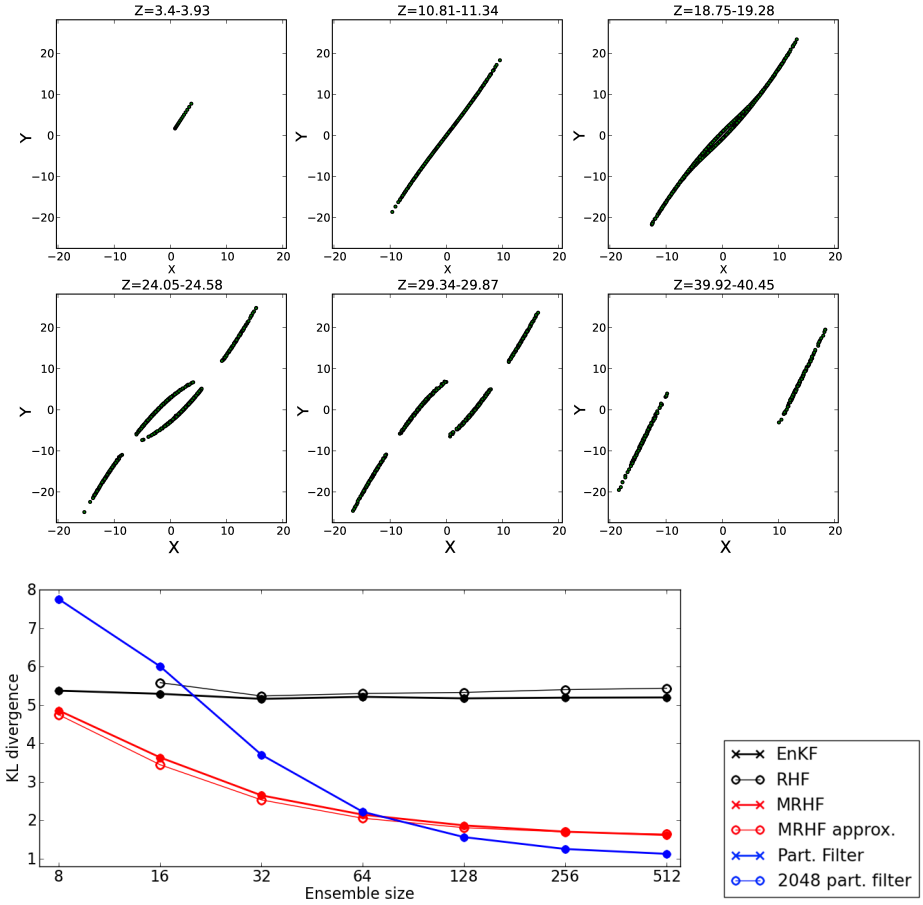


Figure 2.1: Top: Horizontal slices (in X-Y planes, Z intervals indicated on tops) through the Lorenz 63 attractor. Bottom: Kullback-Leibler divergence on X variable vs ensemble size, for the different methods indicated on the right side. The Z variable is observed every 40 time steps and the results shwon here are averaged over $10^5$ time steps. The particle filter with 2048 particles can be considered as the target score (perfect assimilation method).

both the past and the future, while a filter takes only past observations past up to the analysis time into account. The Lorenz-96 model of the small benchmark allows one to vary the nonlinearity of the dynamics by changing the value of the forcing parameter. Hence, it allows for a systematic study of the influence of nonlinearity on data assimilation methods. The results of this study have been published in Nerger et al (2014).

To examine a smoother based on a typical ensemble square-root Kalman filter, the smoother extension of the error-subspace transform Kalman filter (ES-TKF, Nerger et al, 2012) was used either without localization or with observation localization. For a global smoother without localization, experiments have been conducted using the smoother in which both the nonlinearity of the model dynamics and the ensemble size was varied. The performance of the smoother was assessed using the metric of the root-mean square (RMS) error. In particular, the optimal smoothing lag was examined, i.e. the number of analysis steps

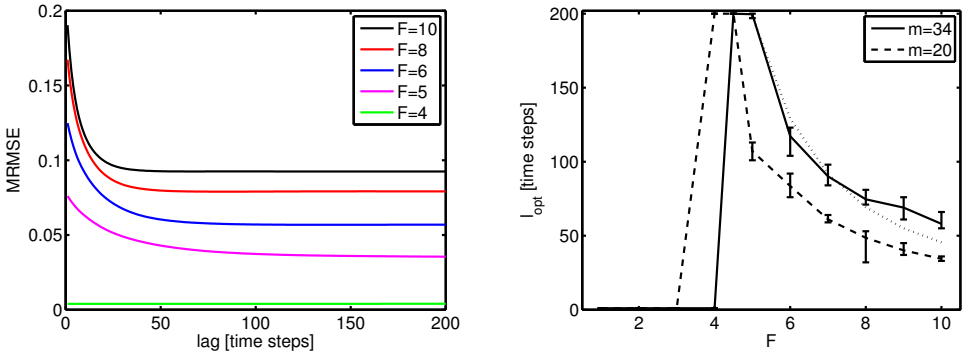


Figure 2.2: Left: time-average RMS error (MRMSE) as a function of the smoother lag for increasing values of the model forcing (i.e., increasing nonlinearity). Right: Optimal lag as a function of the model forcing for the two ensemble sizes $m = 20$ and $m = 34$. [From Nerger et al, 2014.]

into which the state estimate could be corrected into the past for obtaining the minimum RMS error. The typical behavior of the RMS error as a function of the lag is that for short lags the RMS error is strongly reduced compared to the state estimate obtained with the filter (i.e., a lag of 0 analysis steps). For large lags an asymptotic error level is reached. This typical behavior is visible in the left panel of Figure 2.2 for the RMS error averaged over a full experiment of 20000 analysis steps using the ESTKS without localization. Less visible is that before the asymptotic error level is reached a minimum of the RMS error can be reached if the model is nonlinear. The corresponding lag is the optimal lag $l_{opt}$. This optimal lag depends systematically on the nonlinearity of the model as is shown in the right panel of Fig. 2.2. Further more the optimal lag depends on the ensemble size – hence on the sampling quality of the ensemble – such that a larger ensemble results in a longer optimal lag.

If localization is used in the analysis of the filter and smoother, the localization length also influence the optimal lag. The upper left panel of Figure 2.3 shows the time-averaged RMSerror for a value of the model forcing parameter of 8 for different ensemble sizes. The corresponding optimal lag $l_{opt}$ is shown in the upper right panel. As for the global smoother, the optimal lag increases with the ensemble size. However, in addition, there is a localization radius at the optimal lag reaches a maximum. This lag corresponds to the localization radius at which the minimum RMS error of the filter is obtained. Thus, the localizations for optimal filtering and smoothing are the same. For the large ensemble size of $m = 34$ it is visible that even in this case, the localization increases the optimal lag. This increase is desirable, because the assimilation uses more observational information for a longer lag. Accordingly, the effect of the smoother get stronger if localization is used. This is visible in the lower right panel of Fig. **??**, which shows the differences between the RMS errors obtained with the filter and the smoother.

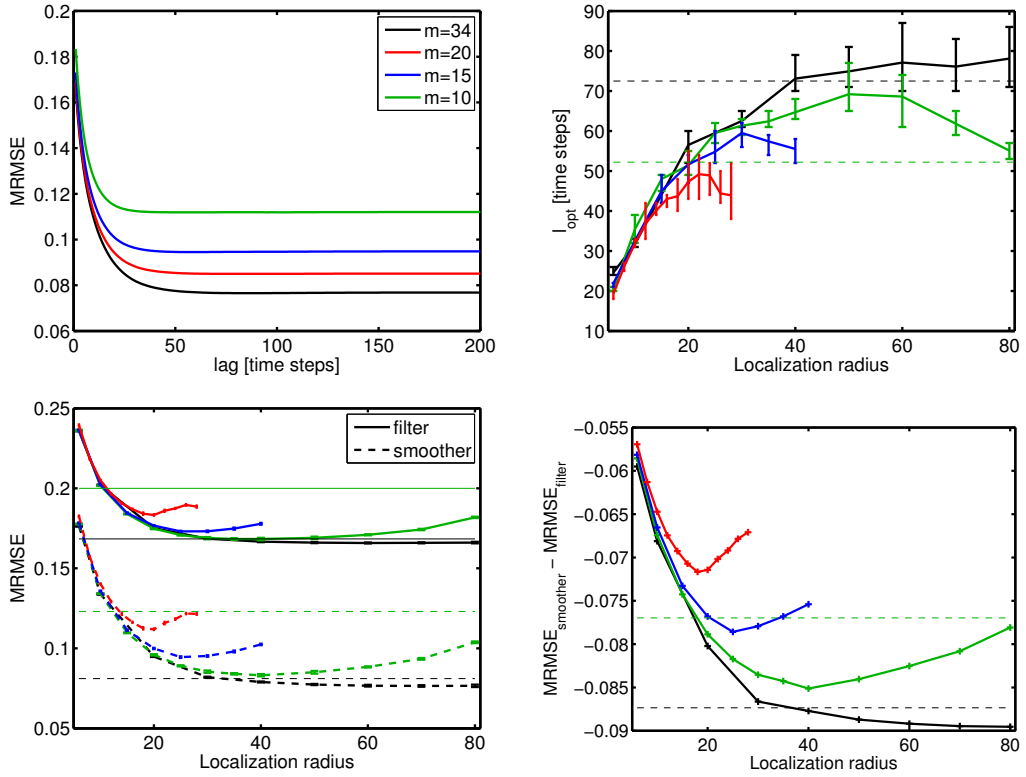


Figure 2.3: Behaviour of the smoothing in the case of localization: (top left) MRMSE as a function of the lag for the optimal localization radius, (top right) $l_{opt}$ as a function of the localization length, (bottom left) minimum MRMSE at $l_{opt}$ as a function of the localization length (dashed) and MRMSE for the filter (solid), and (bottom right) MRMSE of the smoother minus that of the filter. The thin dashed lines show the values for the global analysis for $m = 34$ and $m = 20$. For the smaller ensembles a global analysis is not possible. [From Nerger et al, 2014.]

## 2.3 Equivalent Weight Particle Filter (UREAD)

The small benchmark case was run to test the implementation of equal weights particle filter (EWPF) before releasing the code of EWPF to all partners of Sangoma for implementation in Deliverable 3.3.

We used the setup of the Lorenz 96 model as described in Sangoma specifications for small case benchmark, that is we used:

- 40 state variables
- the model is spun up from t=0 to t=150 (remind dt = 0.05)
- t0=150 as the initial time and t1=1150 as the final time.
- observation error is assumed Gaussian with zero mean and covariance R = $\sigma^2$ I, with $\sigma$=1.2.

However, for simplicity of testing the method, the model error covariance required in EWPF was set to be a diagonal matrix Q = $\sigma^2$ I, with $\sigma$=0.2. Equally,

observation operator was a simple linear operator selecting first 20 state variables, which were observed every 10 time steps. Our tests showed that EWPF was performing as expected for variables which were observed on both sides (see figure 2.4, two top panels) or on the boundary of observation operator (see figure 2.4, two bottom panels).

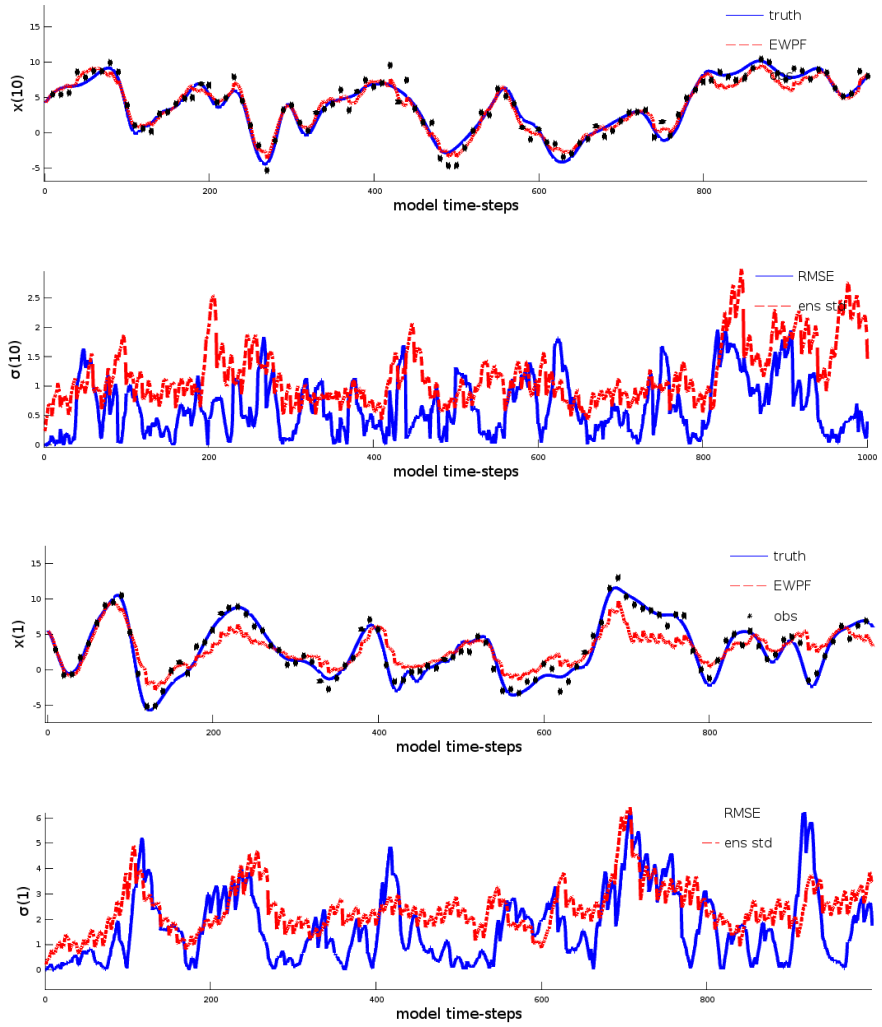


Figure 2.4: Small benchmark (Lorenz 96) run with Equal Weights Particle Filter with 20 particles observing first 20 (out of 40) state variables every 10 time steps. Left: State variable $x(1)$. Right: State variable $x(10)$.

# Chapter 3

# Medium case benchmark

## 3.1 Testing different ensemble filters (AWI)

The medium case benchmark is run with three different data assimilation schemes: the LETKF, the EWPF (both initially planned for SANGOMA) and the NETF, a recently developed second-order ensemble square root filter (Tödter and Ahrens, 2015). All of these methods are implemented in PDAF. The experiments are run with 120 ensemble members and localization is used for the LETKF and the NETF.

Since the EWPF duplicates some ensemble members at the analysis step, it is necessary to introduce some stochastic process in the model to differentiate ensemble members issued from the same sampled member.The NEMO model is then extended with a stochastic wind forcing. In addition, several improvements in the implementation of the algorithm are made to achieve a fast and efficient runtime.The EWPF successfully runs on the medium case benchmark. However, the results are still not as good as obtained with current EnKF methods. It turns out that the definition of the model error is a major challenge for ocean models. The stochastic forcing must be strong enough to have an impact in the nudging step, and weak enough not to dominate the change in the states between consecutive model steps. Further developments on this aspect are further required, possibly in connection with the developments performed for the large scale benchmark (Task 4.6 and Deliverable 4.5).

Figure 3.1 shows a comparison of the performance of the LETKF and the NETF in terms of RMSE and CRPS (Tödter et al, 2015). Both methods successfully reduce the errors in the analysis fields through time. The scores converge to very similar values, although the NETF is slower. This fact is arguably a sign that the medium case benchmark represents weakly nonlinear dynamics only.

## 3.2 Impact of Incremental Analysis Updating with the EnKF (GHER)

Intermittent data assimilation, as it is implemented in standard ensemble Kalman filters, consists in providing an increment to each ensemble members at observation times. Such intermittent adjustments sometimes have detrimental effects,

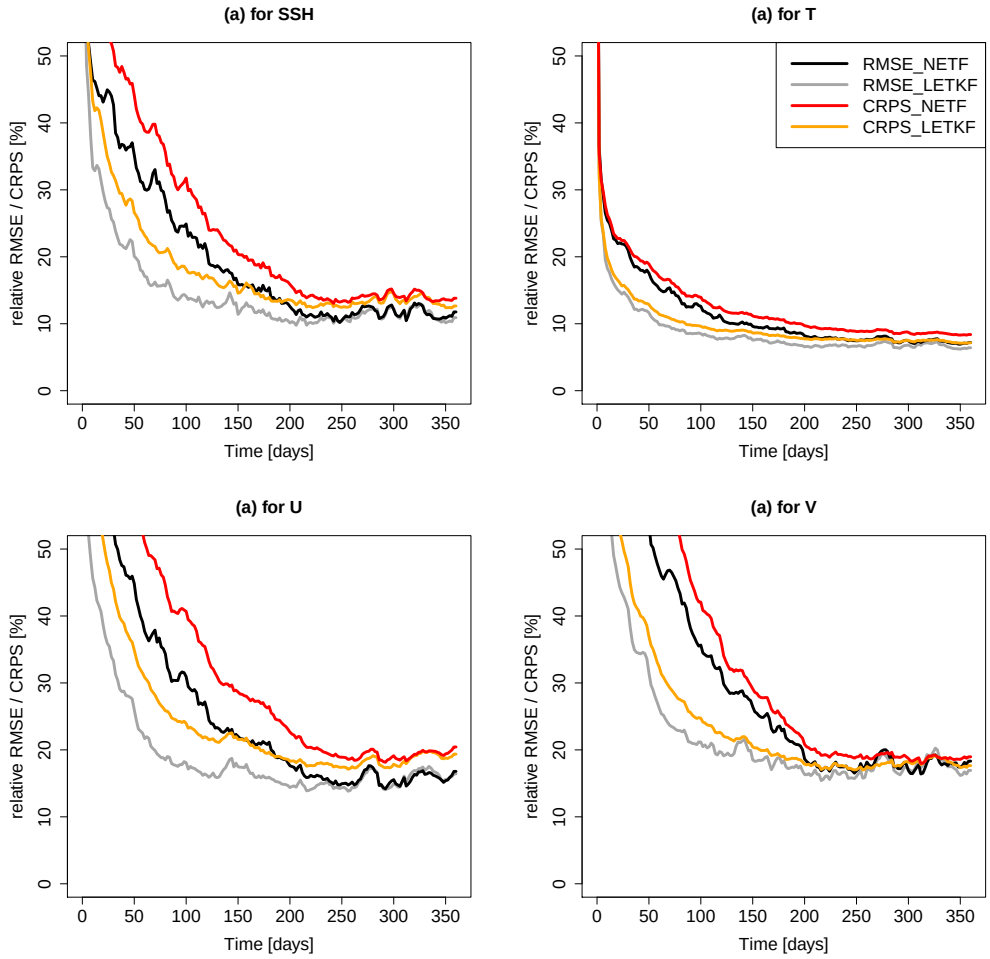


Figure 3.1: Comparison of NETF and LETKF in terms of RMSE (black/gray) and CRPS (red/orange). The lines represent the field-averaged relative RMSE and CRPS, respectively, for all prognostic variables, i.e., (a) $SSH$, (b) $T$, (c) $U$, and (d) $V$. The legend in (b) is valid for all panels.

typically the creation of spurious gravity waves that are difficult to damp. *Incremental Analysis Updating* (IAU) aims at avoiding the drawbacks of intermittent analysis by introducing the analysis corrections in a "homeopathic" way: it is introduced progressively during the model integration of the ensemble.

In the work published in Yan et al (2014), three differents flavors of IAU are compared with each other and with the intermittent data assimilation. The difference between different IAU schemes lies on the increment update window position with respect to the assimilation window under consideration. In the scheme IAU 0, at the end of each assimilation window, the analysis is done using observations around the analysis time. An increment is calculated from the difference between the analyzed and the forecast model states. This increment is then added to the model integration for the subsequent assimilation window. Therefore, the model integration is always forward, there is no model integration repeat for each assimilation window. For the scheme IAU 50, the increment update time

window is located at half of the assimilation window length before and after the analysis time step. After the increment update, the model integration continues for a period of half of the assimilation window length, the model state obtained at the end of this model integration is used for the analysis at the subsequent step. This scheme is then 50% more expensive than IAU 0. With the scheme IAU 100, the increment is calculated in the same way as in the scheme IAU 0, but this increment is added to the model integration in the current assimilation window. In other word, in each assimilation window, the model integration is run twice: the first time for the forecast state and the second time for the increment update. The model state obtained after the increment update is used to initialize the model integration for the subsequent assimilation window. IAU 100 is therefore twice more expensive than IAU 0.

Every 6 days, the ensemble simulations assimilate altimetric observations located along satellite tracks that mimics the trajectory of ENVISAT. To allocate the same computation time to the four data assimilation methods, the intermittent (INT) and IAU 0 are run with 100 ensemble members; IAU 50 with 67 members; IAU 100 with 50 members. Figure 3.2 (left) shows that, quite surprisingly, none of the IAU improves the RMSE scores of the intermittent EnKF. But IAU strongly reduces errors on the vertical velocity. Vertical velocity being a proxy for the system noise in (hydrostatic) circulation models, this result indicates that IAU efficiently damps the spurious gravity waves that INT generates. Figure 3.3 is another way to the same result over a short time interval, and shows that IAU does not rule out the spurious gravity waves, but only reduces then significantly.

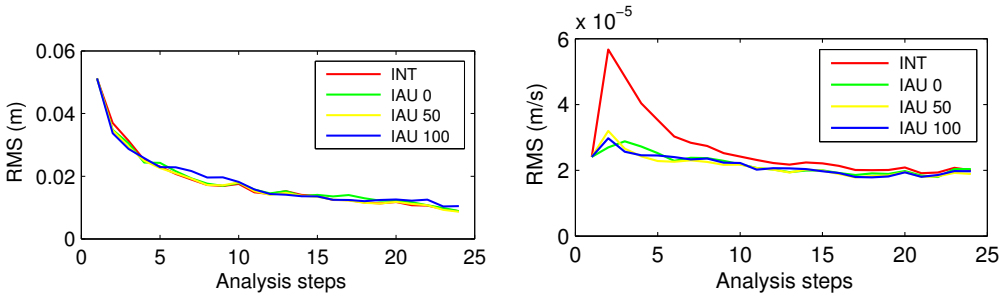


Figure 3.2: Temporal evolution of the spatially averaged RMS error for SSH (left) and vertical velocity (right) in different assimilation schemes with analysis every 6 days.

## 3.3 Testing automatic localization with the EnKF (TUDelft)

The medium case benchmark is implemented with the OpenDA toolbox and the automatic localization strategy is tested for the first time with NEMO. This work is described in details in a paper submitted (van Velzen et al, submitted).The implementation follows a black box coupling strategy, meaning that it requires only the NEMO executable and a wrapper code for reading and writing the input and output files of the model. It is briefly recalled here that localization is meant to remove the effects, on the analysis corrections, of spurious, long-range error covariances

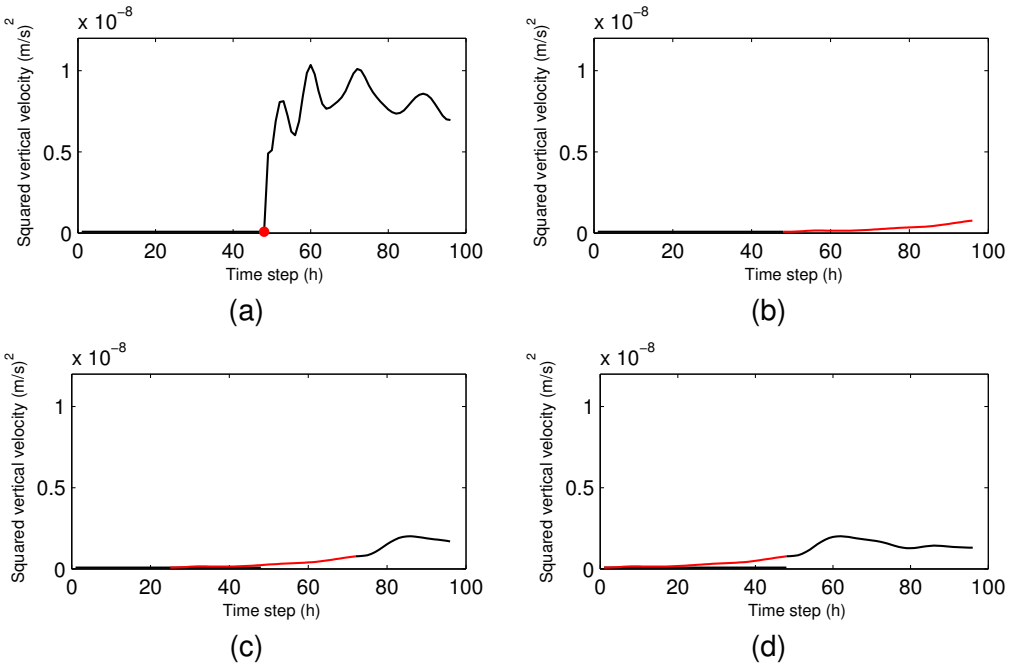


Figure 3.3: Hourly evolution of the squared vertical velocity averaged over the whole domain before and after the first assimilation cycle with the scheme (a) INT (b) IAU 0 (c) IAU 50 (d) IAU 100. The black line corresponds to free model integration, and the red dot and line correspond to instantaneous correction and incremental correction of the model state respectively.

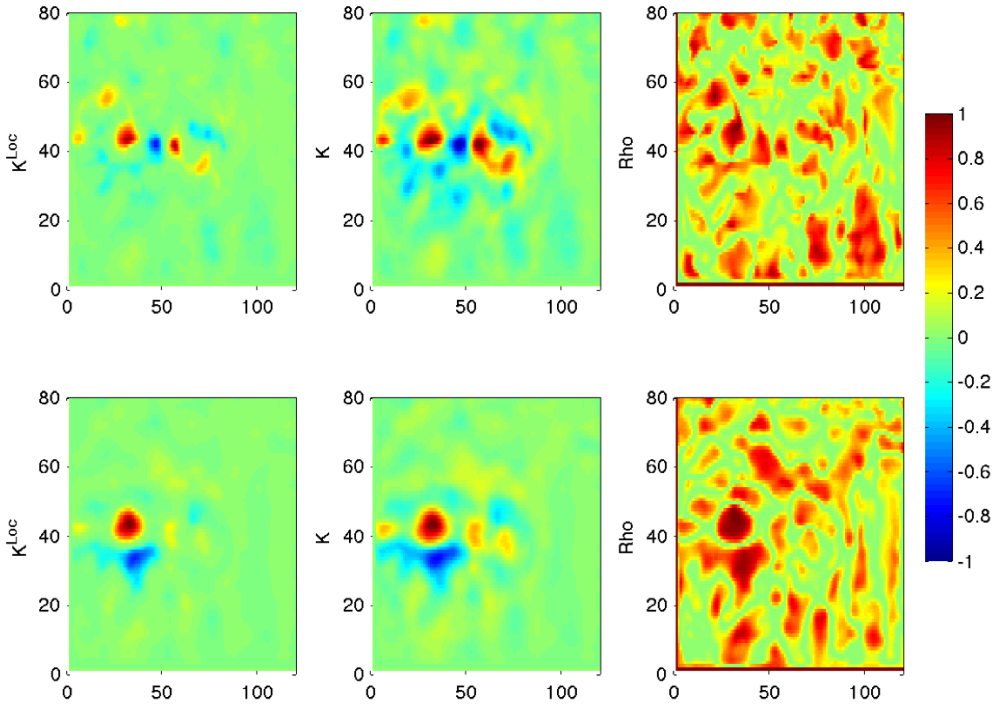


Figure 3.4: Impact of automatic localization on the gain matrix for SSH for various ensemble sizes. The top row is created for an ensemble size of 15 and the lower row with an ensemble size of 100. The gain matrices are generated for the same observation, on the same analysis step of the assimilation run. The first column shows the localized gain, the second column the gain without localization and the last column shows the automatically computed entries of the localization matrix. The method removes structures in the gain matrix which are expected to spurious, sharpening the important structures in the gain matrix.[From van Velzen et al, submitted.]

due to the small size of the ensemble. Localization is often implemented by Schur multiplying the covariance matrix with a sparse matrix that contains zeros at the entries of distant model grid points. The Schur product can also be straightforwardly applied to the gain matrix. This is the strategy chosen in OpenDA. In most applications, the coefficients of the localization matrix are prescribed as a function of the distance between grid points. On the other way round, automatic localization (Zhang and Oliver, 2011) identifies spurious correlations in the gain matrix and then generate the weights in the localization matrix accordingly. The variances of the uncertainty in the gain matrix entries are estimated using a bootstrap technique applied on the ensemble; Entries of the localization matrix are then prescribed as a function of these variances (they decrease with increasing variances).

## 3.4 Iterative smoothing methods (Univ. Nice Sophia-Antipolis)

The well-tested medium case benchmark reveals to be appealing beyond SANGOMA, to implement and test data assimilation methods other than those developed in SANGOMA. In their recent study, Ruggiero et al (2015) implements the diffusive back-and-forth nudging (DBFN), a data assimilation method based on the well-known nudging method. The DBFN consists of a sequence of forward and backward model integrations, within a given time window, both of them using a feedback term to the observations. The backward integration is carried out with the backward model, i.e. the forward model with reversed time step sign. To ensure numeral stability, the diffusion terms also have their sign reversed, giving a diffusive character to the algorithm. In the same study, and for comparison with the DBFN, 4DVar is also implemented with the benchmark. In a following-on study (Ruggiero et al, in preparation), the same benchmark is used to experiment iterative formulations of the SEEK filter.

## 3.5 A new scheme to account for correlated observation errors (CNRS/LGGE)

The medium case benchmark has also been used to test a newly developed scheme to account for correlated observation errors, particularly in the perspective of the SWOT mission. As this aspect is described in details in Deliverable 5.4, it is not detailed here.

# Chapter 4

# Testing non-Gaussian methods with lightweight ecosystem models

Two projects groups (NERSC and CNRS/LGGE) have used lightweight ecosystems models (Simplified 1D vertical dynamics coupled with an biogiochemistry model) to investigate non-Gaussian data assimilation methods.

At NERSC, a one-dimensional configuration of the General Ocean Turbulence Model coupled with the Norwegian Ecological Model (GOTM-NORWECOM) is implemented to simulate the evolution of two classes of phytoplankton (diatom, flagellate), two classes of zooplankton (micro-, meso-), three types of nutrients (inorganic nitrogen, silicon, phosphorous), detritus, biogenic silica and oxygen at Mike weather station (66N, 2E) in the North Sea. Real in-situ observations of nutrients profiles (including Nitrate, Silicate, Phosphate) and oxygen, obtained over the course of 4 years (2001-2004), are assimilated to improve the accuracy of the predicted biological variables and further tune nine biogeochemical parameters. The EnKF and the EnKF with Gaussian anamorphosis are tested and compared.

Figure 4.1 shows the time-evolving estimation of 4 biological parameters: temperature dependence of diatoms, mortality rate of flagellates, sinking rate of detritus and micro-zooplankton loss rate. The trends given by the two assimilation methods are rather similar for all parameters, but there also are significant differences in the estimations, both in terms of mean and standard deviation. This is in spite of similar improvements in the estimation of the observed variables for the two methods. The results for the parameters cannot be easily checked (these are real observation experiments) but they clearly show the influence of the assimilation method in the (non-Gaussian) context of biogeochemistry.

At CNRS/LGGE, the ModECOGel model of the biogeochemistry of the Ligurian sea is used to test a variation of the MRHF. This variation relies on the high flexibility of the method, that computes the corrections on observed and non-observed sequentially, as it can be implemented in the EnKF (Metref et al, 2014): Corrections are computed with the EnKF for the nearly Gaussian variables, and with the MRHF for the non-Gaussian ones. It turns out that the physical variables (temperature, salinity) are close to Gaussian in this system; On the other way round, biological variables are not. We report here experiments where tempera-

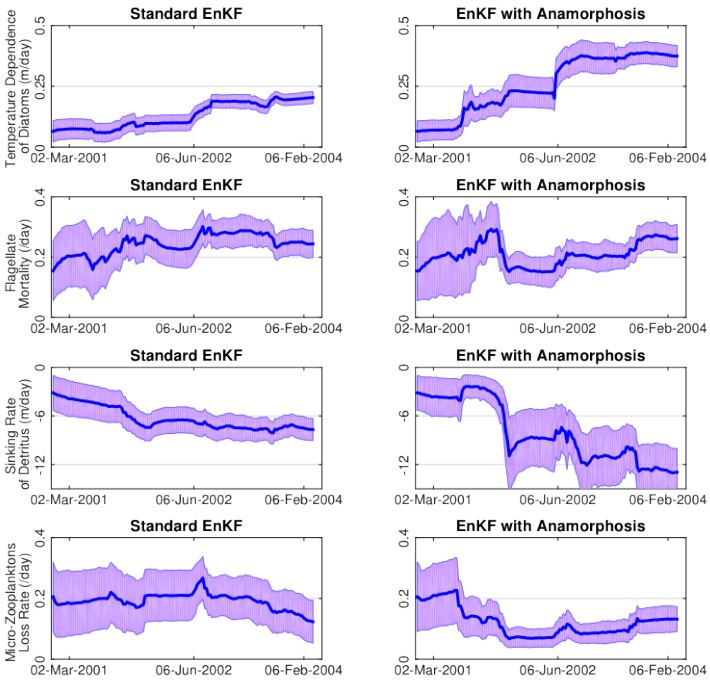


Figure 4.1: Time evolution of the ensemble mean (thick blue) and the averaged mean plus/minus two standard deviations (shaded area) of four different biological parameters. The left panel plots are obtained using the standard EnKF, whereas the right panel-ones are given using anamorphosis.

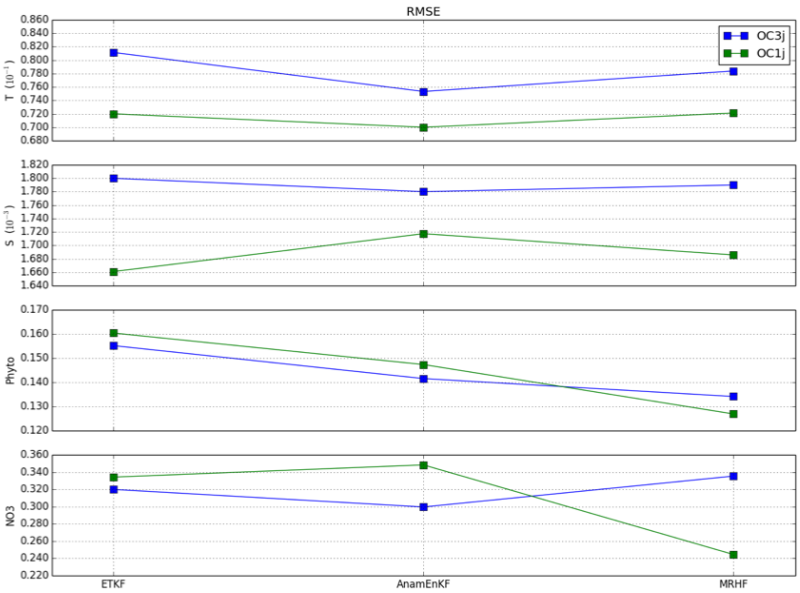


Figure 4.2: RMS errors for temperature, salinity, phytoplankton and nitrate from twin assimilation experiments using the ETKF, the EnKF with Gaussian anamorphosis, and the MRHF. Surface phytoplankton is assimilated every 3 days (in blue) or every day (in green).

ture and salinity profiles are assimilated with the ETKF, while surface phytoplankton data (mimicking ocean colour at this stage) are assimilated with the ETKF, the ETKF with Gaussian anamorphosis, and the MRHF. The experiments are twin, and performed during the 2006 April phytoplankton bloom.

Figure 4.2 reports RMSE scores from the three assimilation methods and for four variables. Experiments are performed assimilating surface phytoplankton every three days and every day. The temporal density of observations is changed to assess the impact of a potential future geostationary, ocean color observing satellite mission (GEO-OCAPI). The MRHF is the only method that brings benefits, for all variables, from an increase of the observation frequency. The Gaussian anamorphosis and the MRHF are perhaps not satisfactory on all aspects, but the results indicate that non-Gaussian methods need further investigation in the perspective of high observation frequencies.

# Chapter 5

# A coastal application

The EnKF is run with the SANGOMA data assimilation tool SDAP and the coastal ocean model SYMPHONIE designed to simulate the circulation in the Bay of Biscay. This configuration of SDAP and Symphonie has not been developed within the SANGOMA project. It is included here since it has been recommended as an additional SANGOMA benchmark in a recent project review. The configuration is based on the SYMPHONIE ocean code with 3km horizontal resolution and 43 generalized sigma levels. The model is forced at the open-boundary by daily outputs from the PSY2 MERCATOR-Ocean operational system and by 9 tidal components from the FES2004 atlas. The atmospheric forcing fields are 3-hour ALADIN products from Météo-France. Symphonie is coupled to SDAP with the objective to study the constraint brought by surface observations on the 3D circulation at daily to monthly time scales. The ensemble is generated by perturbing the wind velocity, assuming that the main model errors sources in terms of surface current and upper stratification over the slope and shelf are due to wind uncertainties. An illustration of the model sensitivity to wind perturbations is given in Kourafalou et al. (2015).

In the experiment presented below, the assimilation cycle is 2 days; daily synthetic observations are generated from a perturbed run and are subsampled every 5 grid points. SDAP is designed to handle full matrices of model errors but in this test case, observational errors are assumed uncorrelated with a centered Gaussian distribution; the error standard deviation is 0.2°C for SST. EnKF in SDAP is asynchroneous as innovations are computed at observational times. Localization is implemented in order to remove spurious long-range error covariances. Experiments with assimilation of SST synthetic observations have been performed over a 1.5-month period (feb-mar 2008) with an ensemble of 54 members; results are illustrated by figure 5.1 (a paper is in preparation). Figure 5.1 shows the impact of constraining the model by SST data on sea surface salinity over the Aquitain shelf. It illustrates the multivariate benefit of the method in an area where dynamical processes at work lead to a strong covariance between both variables at daily time scales. Both the model sensitivity to wind perturbations and the data assimilation impact are different in the deep plain of the Bay of Biscay (Kourafalou et al., 2015; Ayoub, De Mey et al., in prep.).

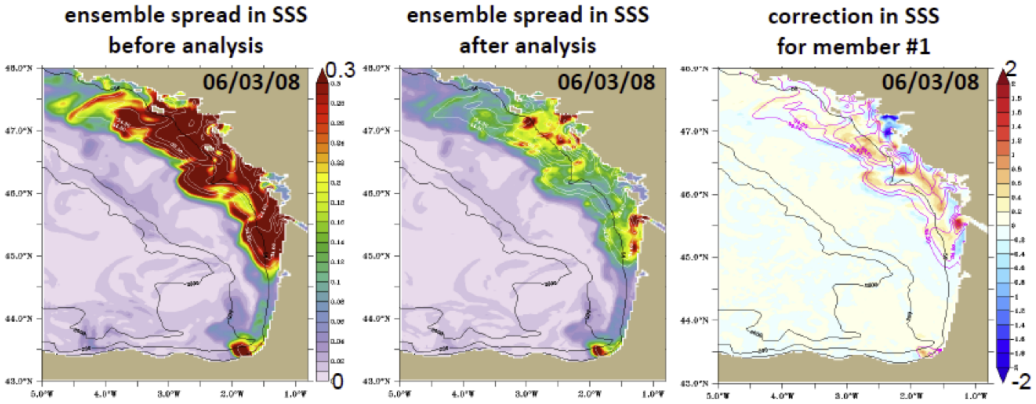


Figure 5.1: Illustration of an EnKF analysis with the coastal case benchmark: standard deviation of the ensemble in SSS before (left) and after (middle) the analysis. Right: correction in SSS for member 1. From Ayoub, De Mey et al., in preparation.

# Chapter 6

# References

Anderson J., 2010. A non-Gaussian ensemble filter update for data assimilation. *Mon. Weather Rev.*, **138**, 4186–4198.

Kourafalou, V. H., P. De Mey, J. Staneva, N. Ayoub, A. Barth, Y. Chao, M. Cirano, J. Fiechter, M. Herzfeld, A. Kurapov, A.M. Moore, P. Oddo, J. Pullen, A.J. van der Westhuysen and R.H. Weisberg, 2015. Coastal Ocean Forecasting: science foundation and user benefits. *Journal of Operational Oceanography*, doi:10.1080/1755876X.2015.1022348.

Lorenz E., 1963. Deterministic nonperiodic flow. *J. Atmos. Sci.*, **20**, 130–141.

Metref S., E. Cosme, C. Snyder and P. Brasseur, 2014. A non-Gaussian analysis scheme using rank histograms for ensemble data assimilation. *Nonlin. Processes Geophys.*, **21**, 869–885.

Nerger L., T. Janjic, J. Schröter, and W. Hiller, 2012. A unification of ensemble square root Kalman filters. *Mon. Wea. Rev.*, **140**, 2335–2345.

Nerger L., S. Schulte, and A. Bunse-Gerstner, 2014. On the influence of model nonlinearity and localization on ensemble Kalman smoothing. *Q. J. Roy. Meteor. Soc.*, **140**, 224–2259.

Ourmières Y., J.-M. Brankart, L. Berline, P. Brasseur and J. Verron, 2015. Incremental analysis update implementation into a sequential ocean data assimilation system. *Nonlin. Processes Geophys.*, **22**, 233–248.

Ruggiero G., E. Cosme, J.-M. Brankart, J. Le Sommer, C. Ubelmann, in preparation for *J. Atmos. Oceanic Technol.*. An efficient way to account for observation error correlations in the assimilation of data from the future SWOT high-resolution altimeter mission.

Ruggiero G., E. Cosme, Y. Ourmières, D. Auroux, J. Blum and J. Verron, in preparation for *Tellus*. Iterative smoothing using the Back and Forth Kalman Filter: applications for the NEMO ocean model.

Ruggiero G., Y. Ourmières, E. Cosme, J. Blum, D. Auroux and J. Verron, 2015. Data assimilation experiments using diffusive back-and-forth nudging for the NEMO ocean model. *Nonlin. Processes Geophys.*, **22**, 233–248.

Tödter J.and B. Ahrens, 2015. A second-order exact ensemble square roor filter for nonlinear data assimilation. *Mon. Weather Rev.*, **143**, 1347–1367.

Tödter J., P. Kirchgessner, L. Nerger and B. Ahrens, in press. Assessment of a nonlinear ensemble transform filter for high-dimensional data assimilation. *Mon. Weather Rev.*

van Velzen N., M. U. Altaf, M. Verlaan, submitted. OpenDA-NEMO framework for Ocean Data Assimilation. Ocean Dynamics.

Yan Y., Barth, A., and Beckers, J., 2014. Comparison of different assimilation schemes in a sequential Kalman filter assimilation system. *Ocean Modelling*, **73**, 123–137.

Zhang Y, D. Oliver, 2011. Evaluation and error analysis: Kalman gain regularisation versus covariance regularisation. *Comput Geosciences*, **15**, 489–508.