Journal of Computational Science 27 (2018) 46–56

Contents lists available at ScienceDirect

Journal of Computational Science

journal homepage: www.elsevier.com/locate/jocs

A fault-tolerant HPC scheduler extension for large and operational ensemble data assimilation: Application to the Red Sea

Habib Toye [a], Samuel Kortas [b], Peng Zhan [c], Ibrahim Hoteit [a,c,*]

[a] Division of Computer, Electrical and Mathematical Sciences & Engineering, King Abdullah University of Science and Technology, Thuwal, Jeddah 23955-6900, Saudi Arabia
[b] KAUST Supercomputing Laboratory (KSL), King Abdullah University of Science and Technology, Thuwal, Jeddah 23955-6900, Saudi Arabia
[c] Division of Physical Science & Engineering, King Abdullah University of Science and Technology, Thuwal, Jeddah 23955-6900, Saudi Arabia

ARTICLE INFO

Article history:
Received 15 November 2017
Received in revised form 7 March 2018
Accepted 22 April 2018
Available online 26 April 2018

Keywords:
High performance computing
Ensemble data assimilation
Bayesian filtering
Operational oceanography
Red Sea

ABSTRACT

A fully parallel ensemble data assimilation and forecasting system has been developed for the Red Sea based on the MIT general circulation model (MITgcm) to simulate the Red Sea circulation and the Data Assimilation Research Testbed (DART) ensemble assimilation software. An important limitation of operational ensemble assimilation systems is the risk of ensemble members' collapse. This could happen in those situations when the filter update step imposes large corrections on one, or more, of the forecasted ensemble members that are not fully consistent with the model physics. Increasing the ensemble size is expected to improve the assimilation system performances, but obviously increases the risk of members' collapse. Hardware failure or slow numerical convergence encountered for some members should also occur more frequently. In this context, the manual steering of the whole process appears as a real challenge and makes the implementation of the ensemble assimilation procedure uneasy and extremely time consuming.

This paper presents our efforts to build an efficient and fault-tolerant MITgcm-DART ensemble assimilation system capable of operationally running thousands of members. Built on top of *Decimate*, a scheduler extension developed to ease the submission, monitoring and dynamic steering of workflow of dependent jobs in a fault-tolerant environment, we describe the assimilation system implementation and discuss in detail its coupling strategies. Within *Decimate*, only a few additional lines of Python is needed to define flexible convergence criteria and to implement any necessary actions to the forecast ensemble members, as for instance (i) restarting faulty job in case of job failure, (ii) changing the random seed in case of poor convergence or numerical instability, (iii) adjusting (reducing or increasing) the number of parallel forecasts on the fly, (iv) replacing members on the fly to enrich the ensemble with new members, etc.

We demonstrate the efficiency of the system with numerical experiments assimilating real satellites sea surface height and temperature observations in the Red Sea.

1. Introduction

Capabilities in ocean modeling and simulation have witnessed tremendous progress in recent years following the advances in high performance computing (HPC) resources [12], the better understanding of the ocean physics, and the availability of ever increasing amount of in situ and remotely sensed data [14,10]. To take advantage of all sources of information from models and observations, data assimilation, the process by which observations are incorporated into the model, is becoming more and more popular to improve the model forecasting skills along with quantification of uncertainties in its outputs [8]. Data assimilation is now recognized as a crucial component for the development of an ocean operational system.

The celebrated Kalman filter (KF) computes the best (minimum-variance) estimate of a linear dynamical system given available observations [23], and as such provides a readily efficient algorithm for data assimilation and forecasting [18]. Because of its prohibitive computational requirements when implemented with large scale systems and the nonlinear nature of the ocean dynamics, simplified Kalman filters have been introduced for ocean data assimilation [36,40,18]. One of the most promising Kalman filtering schemes

* Corresponding author at: Division of Computer, Electrical and Mathematical Sciences & Engineering, King Abdullah University of Science and Technology, Thuwal, Jeddah 23955-6900, Saudi Arabia.
E-mail address: ibrahim.hoteit@kaust.edu.sa (I. Hoteit).

https://doi.org/10.1016/j.jocs.2018.04.018
1877-7503/

is the ensemble Kalman filter (EnKF), a Monte Carlo approach in which the forecast statistics are estimated from an ensemble of model forecasts [21]. An EnKF assimilation system with a high resolution model and large number of observations is expected to require a large ensemble to provide accurate ocean state estimates [20,17]. Large ensembles should provide more reliable forecast statistics and a smooth forecast covariances for efficient implementation of the filter update steps with the observations.

Increasing the ensemble size would however not only significantly increase the computational load, but would also weaken the robustness of the system and increase the chances of system failure, and thus the workload of the user. Indeed, in case the system crashes, the user will have to manually identify the issue behind its collapse, reconfigure the system and check for consistency before relaunching the jobs. The system failures may be related to a machine problem or may be the result of a dynamical inconsistency between the statistically updated ensemble members and the forecasting model, both of which are unpredictable. The users need therefore to continuously monitor the system execution progress.

In an operational ocean forecasting system, not only huge amount of data need to be processed in a timely manner [33], but the system should also be fault-tolerant in order to recover from failure and deliver real-time responses. In this study, we address these ensemble data assimilation forecasting challenges with an EnKF data assimilation system that we configured for the Red Sea. The system is complex and brings together different components (program executables, data, computational resources). An ensemble of MIT general circulation model (MITgcm) runs are integrated in parallel to provide the forecast statistics for the Data Assimilation Research Testbed (DART) filter to perform the assimilation update with the observations. To overcome the aforementioned problems, and build an efficient fault-tolerant ensemble system we coupled the existing DART-MITgcm assimilation system [39] to a scheduler extension named *Decimate* [26]. The system in [39] was neither fault-tolerant nor scalable to ensembles of thousands of members, hence the use of *Decimate* to remediate those limitations. *Decimate* automatically generates the submission scripts along with the dependencies between the jobs and runs them in an environment where checking and restarting functions just need to be defined by the user. It simplifies the launching and monitoring processes and allows for automatic reconfiguration in case of system failure. This work describes the development of the different components of the assimilation system, their coupling and the parametrization of *Decimate*. First results from a high resolution ensemble assimilation system for the Red Sea are presented and discussed.

The paper is organized as follows. We first give an overview of ensemble data assimilation concept and the DART-MITgcm Red Sea forecasting system in Section 2. Section 3, briefly describes *Decimate* on top of which the DART-MITgcm assimilation system was implemented. Section 4 presents the results of the assimilation experiments that has been conducted in the Red Sea. Finally, a brief summary and discussion is given in Section 5.

2. Ensemble data assimilation and the DART-MITgcm system

2.1. Ensemble data assimilation

We follow a Bayesian filtering formulation of the data assimilation problem in which we aim at sequentially computing the probability distribution of the state vector of the system of interest $\mathbf{x}_k$ at time k conditional on the available measurements $\mathbf{y}_{1:k} \equiv \mathbf{y}_1, \mathbf{y}_2, \cdots, \mathbf{y}_k$ up to time k, that is the posterior probability distribution $p(\mathbf{x}_k|\mathbf{y}_{1:k})$ using Bayes' rule [20]. Given an initial distribution $p(\mathbf{x}_0)$, the measurements $\mathbf{y}_{1:k}$, the state space model

$$\mathbf{x}_k = \mathcal{M}_k(\mathbf{x}_{k-1}) + \eta_k \tag{1}$$

$$\mathbf{y}_k = \mathcal{H}_k(\mathbf{x}_k) + \varepsilon_k \tag{2}$$

from which one can obtain the transition distribution $p(\mathbf{x}_k|\mathbf{x}_{k-1})$ and the likelihood $p(\mathbf{y}_k|\mathbf{x}_k)$, the computation can be performed recursively to incorporate the new observation $\mathbf{y}_k$ into the posterior $p(\mathbf{x}_{k-1}|\mathbf{y}_{1:k-1})$ at time $k-1$ to obtain the posterior $p(\mathbf{x}_k|\mathbf{y}_{1:k})$ at time k. $\mathcal{M}_k$ is the dynamical model for advancing the state vector $\mathbf{x}_{k-1}$ from time $k-1$ to time k, and $\mathcal{H}_k$ is the measurement model (or observation operator) at time k. η_k and ε_k respectively refer to independent Gaussian model and observation errors.

Given $p(\mathbf{x}_{k-1}|\mathbf{y}_{1:k-1})$, the Chapman–Kolmogorov equation

$$p(\mathbf{x}_k|\mathbf{y}_{1:k-1}) = \int p(\mathbf{x}_k|\mathbf{x}_{k-1})p(\mathbf{x}_{k-1}|\mathbf{y}_{1:k-1})d\mathbf{x}_{k-1}$$

is used to forecast the state probability distribution at the next time with the dynamical model (1), computing the distribution of $\mathbf{x}_k$ conditional on the observations up to time $k-1$. Bayes' rule is then applied to update the forecast distribution with (the new observation) $\mathbf{y}_k$ to obtain the posterior probability distribution

$$p(\mathbf{x}_k|\mathbf{y}_{1:k}) = \frac{p(\mathbf{y}_k|\mathbf{x}_k)p(\mathbf{x}_k|\mathbf{y}_{1:k-1})}{p(\mathbf{y}_k|\mathbf{y}_{1:k-1})}.$$

This forecast-update cycling process is then repeated whenever a new observation is available.

Bayesian filtering finds applications in many fields including signal processing, meteorology, oceanography, hydrology, finance, motion tracking (of fluids, satellites, airplanes, etc.), among others.

A special case of the Bayesian filter is the Kalman filter, which is designed for linear systems based on orthogonal projections [23,34]. Under the assumption of independent Gaussian model η and observation ε noise, the Kalman filter is optimal in the sense that it computes the best linear unbiased estimate (BLUE). Moreover, due to appealing features, namely easy of implementation, Markovian property (or memory less feature), and sequential process for incorporating the observations, it is widely used in many fields. Nevertheless in oceanography, where the state dimension could be very large (10^7 or more) and the dynamics are strongly nonlinear, a direct implementation of the Kalman filter is not feasible [18,20].

To overcome this problem, Evensen [11] introduced the so-called ensemble Kalman filter (EnKF) as a Monte Carlo implementation of the Kalman filter. In the EnKF, a given (analysis) ensemble of state realization $\mathbf{X}^f = [\mathbf{x}^{a,1}, \mathbf{x}^{a,2}, \cdots, \mathbf{x}^{a,N}]$ is advanced with the dynamical model (1) to compute the forecast ensemble, from which the covariance matrix $\mathbf{P}^f$ used to compute the Kalman Gain is approximated by $\mathbf{P}^{f,e} = \frac{1}{N-1}(X'X'^T)$, where $X' = [\mathbf{x}^{f,1} - \bar{\mathbf{x}}, \mathbf{x}^{f,2} - \bar{\mathbf{x}}, \ldots, \mathbf{x}^{f,N} - \bar{\mathbf{x}}]$ is the ensemble of anomalies and $\bar{\mathbf{x}} = \frac{1}{N}\sum_{i=1}^{N}\mathbf{x}^{f,i}$ the ensemble mean. Once $\mathbf{y}_k$ becomes available, each member of the forecast ensemble is updated using the Kalman filter update step

$$\mathbf{x}_k^{a,i} = \mathbf{x}_k^{f,i} + \mathbf{K}_k\left(\mathbf{y}_k^o + \varepsilon_k^i - \mathbf{H}_k\mathbf{x}_k^{f,i}\right), \quad i = 1, \ldots, N \tag{3}$$

where $\mathbf{K}_k$ is the Kalman Gain

$$\mathbf{K}_k = \left(\mathbf{H}_k\mathbf{P}_k^{f,e}\right)^T\left[\mathbf{H}_k\left(\mathbf{H}_k\mathbf{P}_k^{f,e}\right)^T + \mathbf{R}_k\right]^{-1}, \tag{4}$$

and ε_k^i is sampled from the distribution of the observation error, assumed $\mathcal{N}(0, \mathbf{R}_k)$ [9]. This directly provides an ensemble to start the next assimilation (forecast-analysis) cycle. Perturbing the observations was however shown to introduce noise in the update