DRAFT Addendum to Phase II Final Report for the SBIR project on

An Ensemble Kalman Filtering Approach for Regional Ocean Data Assimilation

Contract N00039-08-C-0017
Amount: $747,336; awarded by Space and Naval Warfare Systems Command

Prepared by

Rui M. Ponte, Principal Investigator
Email: ponte@aer.com; Phone: 781.761.2231

Distribution Statement B: Distribution authorized to US Government Agencies only: contains proprietary information, 30 May 2011. Other requests for this document shall be referred to Space and Naval Warfare Systems Command, 4301 Pacific Highway, OT1 1480D, San Diego, CA 92110-3127.

Submitted to

Space and Naval Warfare Systems Command
4301 Pacific Highway, OT1 1480D
San Diego, CA 92110-3127

Attn: Edward Mozley, PMW 160
Email: edward.mozley@navy.mil; Phone: 858.537.0230

Submitted by

Atmospheric and Environmental Research, Inc.
131 Hartwell Avenue, Lexington, MA 02421-3126
Phone: 781.761.2288; Fax: 781.761.2299

30 May 2011

[This page intentionally left blank.]

DRAFT Addendum to the PHASE II Final Report (period ending 30 May 2011)
CONTRACT NO. N00039-08-C-0017

Administrative Narrative

The Phase II Option was used to extend the Phase II work by transitioning the LETKF-NCOM ocean data assimilation system to the group 7320 computers at Stennis Space Center. The goals of the Phase II Option are stated in the Phase II Final report, and were aimed to address essentially two issues: (1) deploy LETKF-NCOM on the computers at Stennis, and (2) compare LETKF-NCOM results to NCODA-MVOI results in as fair a comparison as is practical.

Various items that detail the administration of the Phase II Option were given in the three quarterly progress reports (QPR), submitted on the 23rd day of September 2010, December 2010 and March 2011. Only the details from prior QPRs that are relevant for this addendum are reproduced below.

The Phase II Option began June 23, 2010, and a kick-off teleconference was held on July 21, 2010. Ed Mozley convened the meeting that included Clark Rowley, Pete Spence Ross Hoffman and Mark Leidner.

Start-up of the project was delayed during late summer 2010 without access to the 7320 computer systems at Stennis Space Center. By late September 2010, Mark Leidner and Ross Hoffman had the hardware and Information Awareness (IA) training to access the 7320 computer systems at Stennis Space Center. The project was largely fallow during the month of October 2010 as Mark Leidner recovered from an unplanned health-related issue. The task of transitioning the LETKF-NCOM ocean data assimilation software began in November. Given the slow start to the Phase II Option, a 6-month, no-cost extension was granted to allow the time for the work to be completed.

As soon as a version of the LETKF-NCOM software was working on the computers at Stennis (Dec 2010), new software development was required to move from the small-sized ensembles (2-6 members) used during the Phase II work, to larger, more realistically-sized ensembles ($>$ 24 members).

Eric Kostelich, Assistant Professor of Mathematics at Arizona State University and co-author of the LETKF software, was hired by AER as a consultant for the last five months of the Phase II Option.

Near the end of the Phase II Option, Mark Leidner made a site visit to NRL at Stennis Space Center (SSC). During the two-day site visit (June 21-22), Mark Leidner and Clark Rowley exchanged technical information about NCOM and LETKF, looking for ways for NCODA and LETKF to work together. The goal of the trip was to

Technical Progress

T–I. Introduction

The simplified sequence of goals during the Phase II Option was: (1) reproduce the same 6-member ensemble assimilation results on Stennis Space Center computers as were generated at AER during the Phase II work, (2) increase the number of ensemble members to a more realistic size (i.e., more than 24) for larger-scale testing and qualification, (3) tune the assimilation system through various cycling data assimilation experiments that cover weeks to months of testing, and (4) compare to existing NCODA-MVOI and NCODA-3DVAR results.

The technical progress reported below shows that tasks 1-2 are complete, task 3 is mostly complete, and task 4 has not started.

Clark Rowley and Mark Leidner discussed approaches to task 4, i.e., making a fair comparison among the MVOI, 3DVAR and LETKF assimilation systems, on several separate phone calls during the Phase II Option. The approaches ranged from off-line comparisons (comparison of results from independent tests; least preferred) to comparisons among versions of NCODA that use MVOI, 3DVAR and LETKF as the analysis technique. This last approach is the most preferred since the same observation preprocessing (including quality control and de-tiding) and validation approach would be applied across all assimilation methods. But the integrated NCODA comparisons require integration of LETKF into the NCODA framework, since a version of NCODA that uses LETKF does not currently exist. (LETKF/NCODA integration was suggested in the work plan for the Phase II Option in the Phase II final report, page 8: *"Integrate the LETKF NCOM into the RELO cycling scripts (i.e., use existing model support codes to surface and lateral boundary condition preparation, etc.)."*.) Without a well-tuned LETKF-NCOM assimilation system during the Phase II Option, however, it was premature to attempt any comparisons to the existing NCODA assimilation systems.

T–II. LETKF Deployment at Stennis Space Center

The preceeding Phase II work used a single-processor version of LETKF-NCOM exclusively. As early as possible in the Phase II Option, it was important to advance the LETKF-NCOM ocean data assimilation system to make use of distributed computing resources so that larger ensembles of forecasts could be used.

First, however, considerable effort was required to compile and test the single-processor version of LETKF at Stennis. The code was developed with the Intel fortran compiler, but an up-to-date Intel compiler was not available at Stennis. Stennis Group 7320 computers make wide use of the Portland Group (PGI) fortran and C compilers. But LETKF's primary author, Eric Kostelich, strongly recommended that we choose the g95 or gfortran compilers instead of the Portland Group compilers. Indeed, we could not get an LETKF executable created with the PGI compilers to execute properly. All of the work reported here made use of gfortran and OpenMPI message-passing software for compiling.

Significant development and testing was required to upgrade the LETKF-NCOM software to parallelize the input and output of NCOM ensemble members. In the single-processor version, all ensemble input or output was handled via a single, large file that contained the full model states for all ensemble members. But for much larger ensembles, this approach was clearly going to produce an i/o bottleneck. Further, with more ensemble members, much more memory would be required to complete the analysis. Without distributing the analysis domain during the solving phase of the analysis, the cpu and memory requirements of a single-processor analysis would make extended testing, not to mention operational use, impractical. A single processor version using a 6-member ensemble takes about 45 minutes to complete one analysis. With the work accomplished during the Phase II Option, a typical analysis using a 36-member ensemble, distributed among 36 processors would finish in 8-12 minutes, depending on the number of observations assimilated.

Among the new software written to accommodate parallel NCOM input to and output from LETKF was software to broadcast NCOM model grid parameters to all of the processors being used. Before any processor can perform its portion of the analysis, it must know everything about the NCOM grid (bathymetry, height of the model's sigma and fixed height surfaces at every grid point, etc). This process was centralized by reading in the NCOM grid parameters only once by the master process and broadcasting these parameters out to all of the slave processes before the solving portion of the analysis begins. For reading and writing the NCOM forecast and initial condition files we were able to reuse software developed during Phase 2, which was previously part of the interface package and is now part of LETKF proper, and then to make use of the procedures built-in to LETKF to communicate forecast and analysis ensemble members.

T–III. Single Observation Testing

Assimilating a single observation yields the finger print of an analysis system. In the case of LETKF, such a finger print is dependent on the background flow. The first five figures of this report present results from single observation tests.

Single observation tests were conducted to quantify the impact that an observation can have on the analysis under controlled conditions. These tests were also conducted to validate the assimilation technique for simple-to-understand cases. Single observation tests were also used to tune the various localization parameters to values appropriate for ocean data assimilation and in particular, the Okinawa Trough NCOM ocean modeling domain in July 2007.

The results presented here use a single observation near the center of the NCOM model domain, approximately 27 N latitude and 126 E longitude. Fig. 1 shows sea surface temperature for the whole modeling domain, including the ensemble mean background (top panel), mean analysis from assimilating the observation (middle panel), and the analysis increment (analysis minus background; bottom panel).

The observation is a single SST measurement, set to be 0.5 C warmer than the background, with an assigned observation error standard deviation of 0.5 C. 0.5 C was chosen since it is a larger but realistic difference between model and observation. As a consequence, the color scale on the the analysis increment panel is -0.5 to +0.5 C to allow for the possibility that the analysis matches the observed value exactly at the observation location (i.e., an increment of +0.5 C).

The ensemble used for the background began from a 36-member climate ensemble described as the 'nature run' in the the Phase II final report. The 'nature run' of NCOM was a 9-day NCOM forecast with output fields available four times per day (00, 06, 12 and 18 UTC). The 36-member ensemble begins with the 36 model states from the 9-day nature run (9 days times 4 times/day equals 36 NCOM model states). We found that without a few days of cycling data assimilation, the 36-member ensemble retained unwanted characteristics, peculiar to our choice of the starting climate ensemble. After a few days of cycling data assimilation, we found that artifacts in the starting climate ensemble had been washed away through the evolution of the model states during cycling data assimilation. A number of such experiments were conducted to tune the observation standard deviation, and the localization and ensemble inflation parameters of LETKF. At the point in analysis cycling selected as the background time for the single observation tests, a set of model states populates the ensemble that are representative of the analysis uncertainty and that has been reconciled to a time series of observations through cycling data assimilation.

The observation location was also chosen because the background ensemble standard deviation is a bit larger than the average location, 0.5 C, whereas typical values for other observation locations are 0.1-0.3 C. A larger than normal ensemble standard deviation for a particular location indicates larger uncertainty in the background field.

Given that the background uncertainty *and* observation uncertainty are 0.5 C, we would expect the analysis increment at the observation location to be an equal compromise between the background value and the observed value – that is, an increment of 0.25 C. Magnified plots, showing just the analysis patch around the observation (Fig. **??**), demonstrate that the increment is nearly exactly 0.25 C at the observation point (see bottom panel of Fig. **??**). The analysis patch in this case extends 120 km north, south, east and west, for an effective analysis box that is 240 km by 240 km centered on the observation. The horizontal localization also reduces the impact of the observation using a trapezoidal taper that gives full weight to the observation within 9 km north, south, east and west, but then tapers linearly to zero at 120 km. This increment produces a modest change to the background (compare top and middle panels of Fig. **??**). But the pattern of the increment reveals the flow dependence of the background error uncertainty. Because this observation falls on a relatively sharp gradient in sea surface temperature whose position is uncertain among the ensemble members, the influence of this observation follows the temperature gradient both to the west and even further to the east of the observation location. In fact, the flow dependent uncertainty follows the gradient far from the observation location, since the position of the north and east edges of the adjacent pool of warm water is uncertain. This can be seen in the bottom panel of Fig. **??** where contours of the analysis increment (black) are overlaid on top of the mean background sea surface temperature. In this visualization, it is easy to see how the increments follow the temperature gradient on the north and to the east sides of the warm pool.

Fig. **??** shows an west-to-east vertical cross section of potential temperature along 27 N latitude that crosses through the observation point. (See the red line in the top panel of Fig. **??** for the cross section location.) Note that vertical localization restricts the influence to the top 18 meters or so of the ocean, an approximate mixed layer depth for the North Pacific diagnosed from a sampling of nine NCOM model states during early July 2007 (not shown).

Finally, Fig. **??** shows a progression of horizontal localization from very wide influence (1500 km by 1500 km analysis patch, top panel) to the regionally restricted and tapered localization (240 km by 240 km analysis patch with a 111 km trapezoidal taper, bottom panel) that is used throughout all of the other results presented in this report. It is clear that using no taper, i.e., the Heaviside step function (middle panel), will produce artificial and unwanted edges/gradients in the analyzed field.

T–IV. Cycling Assimilation Results

The view of LETKF's impact on the analysis from a single observation in the last section is now scaled up to many observations.

Fig. **??** through Fig. **??** show the same kind of plots as were presented for the single observation results in the previous section, but for a much larger portion of the whole Okinawa Trough modeling domain (see the red boxes in Fig. **??**) and assimilating more than 1,000 AVHRR sea surface temperature retrievals (these are the Multi-Channel AVHRR SSTs, MCSST) plus scattered ship, fixed and drifting buoy measurements.

Apart from some artifacts in the extreme northeast corner of the whole modeling domain, the increments produced by assimilation of many observations is much as we would expect. In some areas, the observations have little impact, because the background is already very close to the observed values. In other areas, the uncertainty in the background field and the difference between the observation and background create large, significant increments. But because neighboring points see the same uncertainty in the background flow, the resulting analysis increments reflect a shared understanding of the uncertainty, which produces increments that vary in space according to the background flow.

A one-week cycling data assimilation test was performed assimilating temperature and salinity measurements only with a 36-member ensemble, for the period 12Z July 7 to 06Z July 8, 2007. In the framework used in this study, one week is 28, 6-hour analysis/forecast cycles. We chose not to assimilate sea surface height measurements in this test for several reasons: our starting climate ensemble captures the tidal cycle in a completely mixed state of phases; we do not have the de-tiding and de-meaning software for the NCOM model fields required to line up with some sea surface height observations; we felt it was more important to align the ensemble with the 3D model state variables first and perhaps turn on assimilation of sea surface height later, after the ensemble has been reconciled to T and S observations for one week or one month of cycling data assimilation. Sea surface height is the result of the composition and thermodynamic state of the column of water beneath it. Once the 3D fields have been conditioned by and reconciled to available T and S observations, the height field should be in a better condition to begin assimilation of height observations.

Fig. **??** shows the statistical analysis of the fit of the background and analyzed ensemble mean fields to the sea surface temperature observations assimilated (top panel). (The fit to salinity measurements is of limited utility due to the small number of observations and is not shown here.) For each analysis cycle, the fit of the analyzed fields is slightly improved compared to the background, which is one sign of a healthy assimilation system. The occasional spikes in the o-a/o-b stats are possibly due to a quality control issue with the incoming GODAE observation data. This requires further investigation to confirm and address. But apart from what appears to be an occasional QC issue with the observations, the assimilation system seems stable. Further, the analysis of the ensemble variance at the SST observation locations over the course of the 1-week experiment (middle panel) seems stable for both the background and analyzed fields. (The ensemble variance may be starting to slowly diverge during the last two days, or this may be a temporary, small change in the ensemble variance due to changing flow conditions in the modeling domain.) For reference, the number of sea surface temperature observations assimilated at each analysis cycle is shown in the bottom panel of Fig. **??**.

Current Status and Future Activities

The LETKF/NCOM experiments at Stennis are now underway, using a realistic ensemble size and model domain. We conducted preliminary tuning of the LETKF parameters—observation standard deviations, localization, and ensemble inflation. One full week of data assimilation is complete. Analysis of single-obs experiments demonstrate the great potential of LETKF to capture and make use of flow-dependent uncertainty in the background (6-h forecast) ensemble. Technical progress during the option phase includes (i) migration of the LETKF to the Stennis environment; (ii) modification of the I/O and data handling specific for NCOM within LETKF to make full use of the multi-processing capabilities of LETKF; and (iii) demonstration that the system scales well on the Stennis computer system.

Immediate future activities should focus on integrating LETKF into the NCODA system. This work will parallel in some aspects a similar effort that has already integrated 3d-VAR with NCODA. Because LETKF is an ensemble system this will require some extensions, but for the most part the LETKF ensemble mean forecast will be the ”background” provided to the NCODA pre-processing segment and the LETKF ensemble mean analysis will be the ”solution” provided to the NCODA post-processing segment. An interface between the NCODA observation pre-processor output and the LETKF input will need to be forged to assure that observation selection and quality control are handled consistently. But a model for this pathway already exists in 3d-VAR version of NCODA. There will therefore be few changes for an initial demonstration, but there are several issues that will require more careful consideration going forward. First the LETKF estimate of uncertainty is restricted to the subspace spanned by the ensemble. This is precisely what is required for LETKF, but is generally small compared to the true uncertainty. An estimate of the true uncertainty based on the ensemble variability will need to be developed to ensure that the NCODA post-processing assumptions are valid. Second, sources of uncertainty due to the atmospheric forcing and the lateral boundary conditions are not part of the current NCODA design, but are desirable for ensemble data assimilation methods in order to reduce reliance on ad hoc ensemble inflation schemes.

Then, with integration into NCODA complete, experiments running for several weeks to several months will enable fine tuning of LETKF parameters and detailed comparison versus the MVOI and 3d/4d-Var analysis methods. If warranted by results from these experiments, transition to operations should be efficient due to the integration of LETKF with NCODA, which is itself fully operational.

Acknowledgments

AER wishes to thank Tim Campbell and Daniel Goolsby of the Code 7320 Computer support staff at NRL Stennis for keeping the queues running well. AER also appreciates the scientific support provided by Clark Rowley and Pete Spence of the Ocean Dynamics and Prediction Branch.

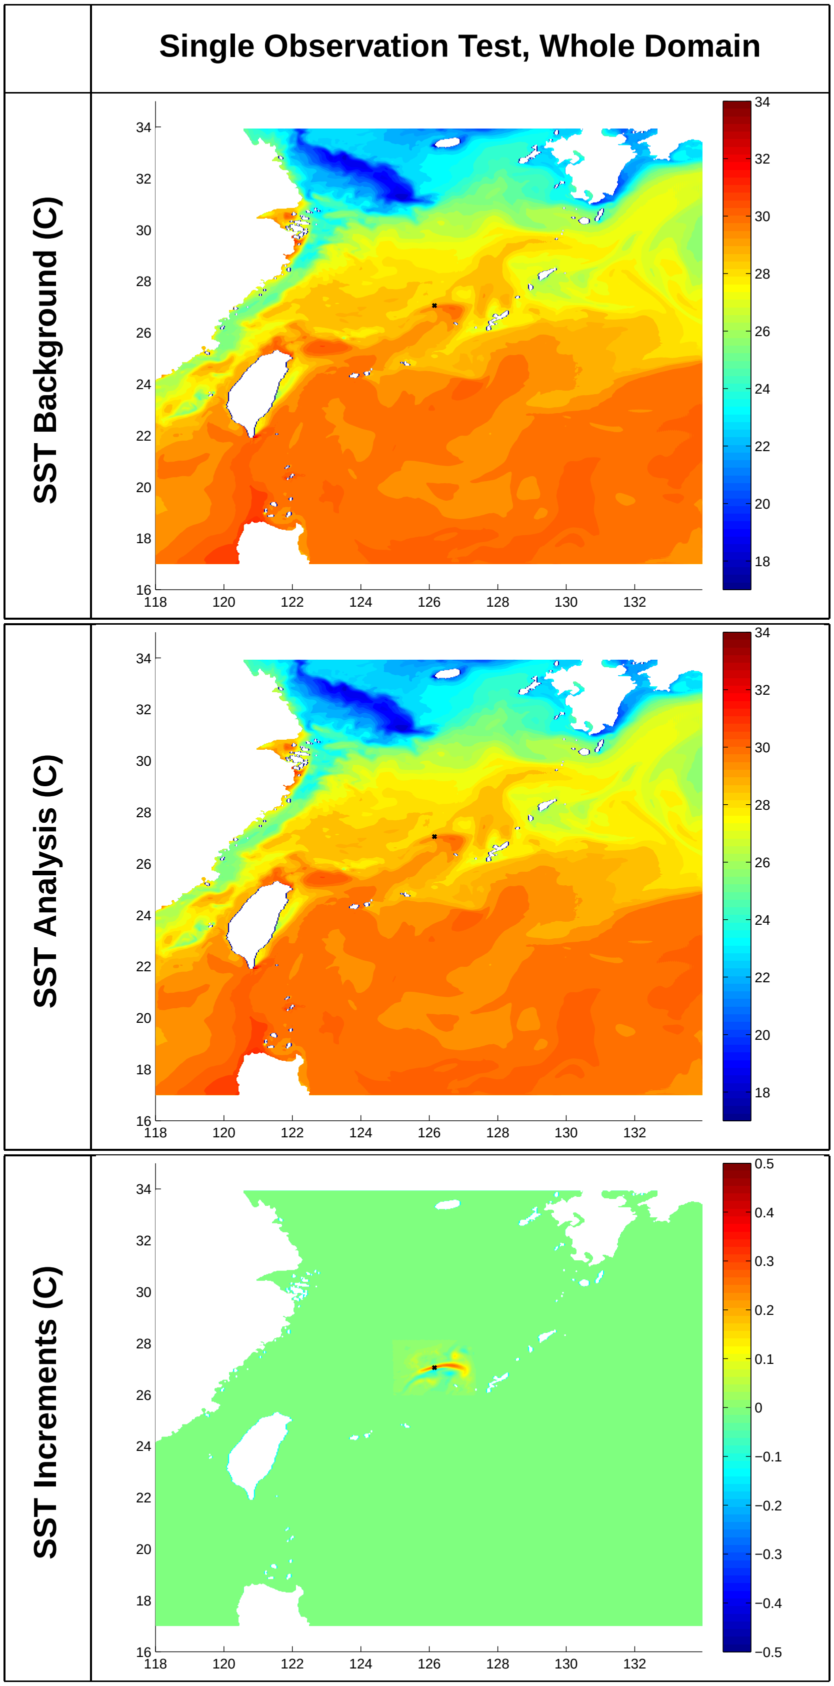
Single Observation Test, Whole Domain
SST Background (C)
SST Analysis (C)
SST Increments (C)
16
18
20
22
24
26
28
30
32
34
118
120
122
124
126
128
130
132
34
32
30
28
26
24
22
20
18
0.5
0.4
0.3
0.2
0.1
0
−0.1
−0.2
−0.3
−0.4
−0.5