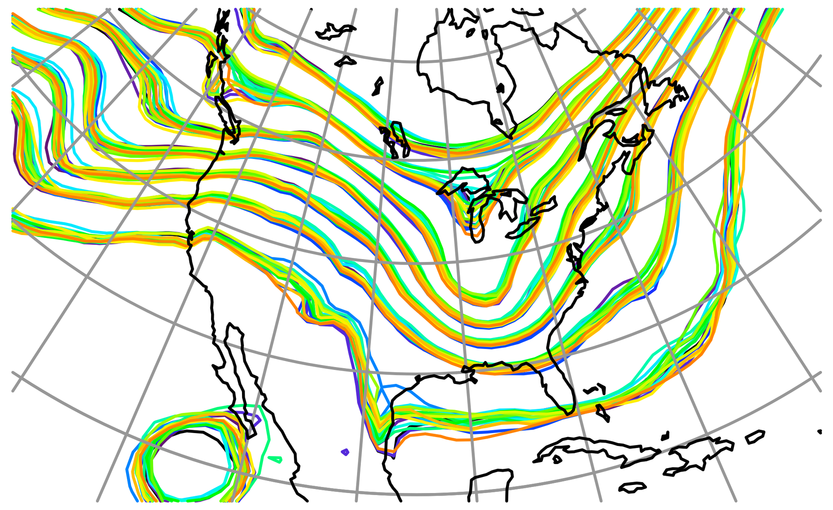

DART: A Community Facility Providing State-of-the-Art, Efficient Ensemble Data Assimilation for Large (Coupled) Geophysical Models

T. Hoar, J. Anderson, N. Collins, H. Kershaw, J. Hendricks, K. Raeder, G. Romine, A. Mizzi, J. Barré, B. Gaubert, L. Madaus, A. Aydogdu, J. Raeder, H. Arango, A. Moore, C. Edwards, E. Curchitser, R. Escudier, R. Dussin, C. Bitz, Y. Zhang, and P. Shrestha

National Center for Atmospheric Research
Boulder, Colorado, USA

`dart@ucar.edu`

Institute for Mathematics Applied to Geosciences

Data Assimilation Research Testbed

1. DART 2.0 Highlights

- ★ Able to handle much larger model states. This is needed for higher-resolution and/or strongly-coupled DA with multiple components. Distributing the model state across all tasks during the entire filter run means no single task must store the entire state at any time.
- ★ One-sided MPI communication allows tasks to request remote data items from other tasks without interrupting their execution or arranging which data items will be needed in advance.
- ★ Computing the forward operators for all ensemble members at the same time leads to code that vectorizes better.
- ★ Native netCDF support eliminates a conversion step that translated between netCDF model files and a DART binary format file. This also reduces the high-water mark for disk requirements.
- ★ Ensemble data can be read and distributed across all tasks on a variable-by-variable basis, reducing the maximum memory requirements.
- ★ Diagnostic state space files are now written in parallel with state-space restart files, resulting in faster I/O and lower memory requirements.
- ★ Support for externally-computed forward observation operators.
- ★ Support for per-observation-type localization radii.

2. DART is ...

The Data Assimilation Research Testbed (DART) is an open source community software facility for ensemble data assimilation developed at the National Center for Atmospheric Research (NCAR). DART works with a wide variety of climate and weather models and observations and has been free and publically available for more than 10 years. Building an interface between DART and a new model does not require an adjoint and generally requires no modifications to the model code. DART works with dozens of models of varying complexity, including (but not limited to):

- weather models, e.g. WRF, COAMPS, COSMO, MPAS Atmosphere,
- components of climate models, e.g. CAM, POP, CLM, WACCM, MPAS Ocean, MITgcm-Ocean, GCCOM, ROMS, JULES, FESOM, CICE5,
- atmospheric chemistry models, e.g. CAM-CHEM, WRF-CHEM,
- ionosphere/thermosphere models, e.g. TIEGCM, GITM,
- low-order and simple research models

DART assimilates a wide variety of observation types including:

- temperature, winds, moisture from NCEP, MADIS, and SSEC,
- total precipitable water, radar observations, radio occultation observations from GPS satellites,
- ocean temperature and salinity from the World Ocean Database,
- land observations such as snow cover fraction, ground water depth, tower fluxes, cosmic ray neutron intensity, and microwave brightness temperature observations.

DART provides both state-of-the-art ensemble data assimilation capabilities and an interactive educational platform to researchers and students.

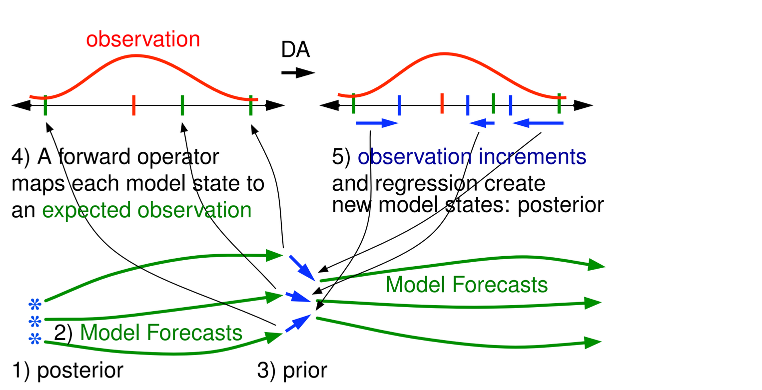

Figure 1: *Schematic for a toy ensemble size of 3.*

J. Anderson, T. Hoar, K. Raeder, H. Liu, N. Collins, R. Torn, and A. Arellano, 2009: *The Data Assimilation Research Testbed: A Community Data Assimilation Facility. BAMS* **90** No. 9 pp. 1283–1296

http://www.image.ucar.edu/DAReS/DART has information about how to download DART, the DART educational materials, and how to contact us.

3. Acknowledgments

The National Center for Atmospheric Research is sponsored by the National Science Foundation. Some computational resources were provided by the Computational and Information Systems Laboratory at NCAR.

4. Motivation for DART 2.0

4.1 Data Layout

One of the challenges for assimilation with truly massive geophysical models is the data layout. Here is a simple example using 4 ensemble members and 4 processing elements (PEs). Each PE is a separate color, a representative observation location is represented by the red star, and some hypothetical area expected to be influenced by the observation is outlined in red.

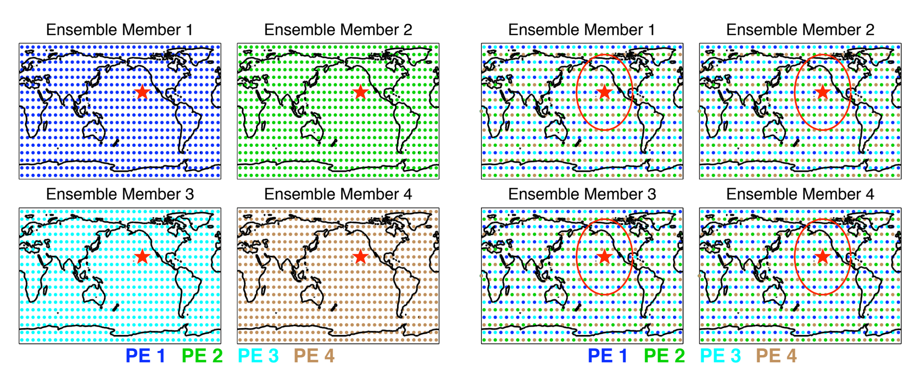

Figure 2: *Two possibilities for data layouts. Left: 'Whole State'. Each ensemble member exists exclusively on a single processing element (PE). Right: 'Distributed'. Each ensemble member is randomly distributed over multiple PEs. Both layouts have advantages and disadvantages and an expensive data transpose is required to go between layouts.*

Both data layouts were used in the first version of DART: one for forward operators and one for the assimilation phase.

4.1.1 Characteristics of Whole State Layout

- If each PE has a complete state, forward operators require no communication.
- Forward operators for all ensemble members are computed in parallel.
- **However**, it is not memory scalable and cannot effectively exploit PE counts greater than the ensemble size.
- **Additionally**, newer architectures have large numbers of PEs with less memory per PE.

4.1.2 Characteristics of the Distributed Layout

- Each processing element (PE) stores all ensemble copies of a subset of the model state.
- All ensemble members for a state variable are on one PE.
- Can compute state mean, variance without communication.
- Forward operators probably require communication.
- All increments are computed in parallel.
- One PE broadcasts observation increments.
- Models frequently use a 'block' style layout as shown here. **However**, load balancing is an issue. PE4 has lots of work, PE1 has none. 'Randomly' distributing the ensemble states is likely to result in a well load-balanced system at the expense of increased communication to compute the forward operators.

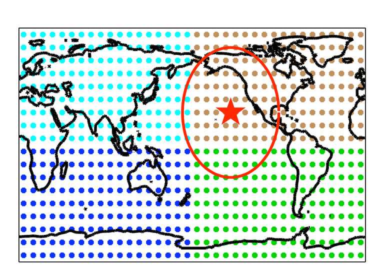

Anderson, J. and N. Collins, 2007: *Scalable Implementations of Ensemble Filter Algorithms for Data Assimilation. J. Atmos. Oceanic Technol.*, **24**, pp. 1452-1463, doi: 10.1175/JTECH2049.1

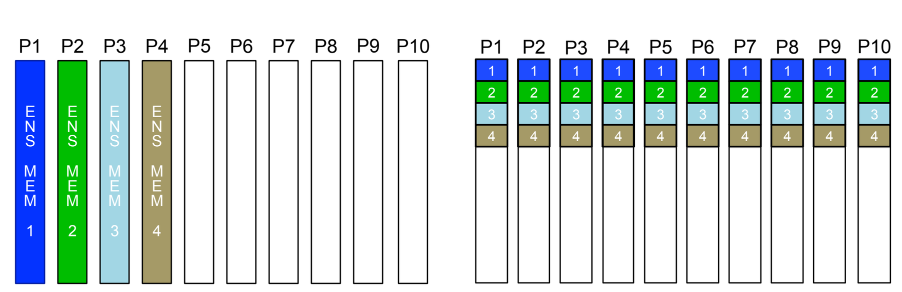

Figure 3: *Two data layouts. Left: whole state vectors are collected on each PE, up to the number of ensemble members. Right: parts of each of the ensemble of state vectors are distributed across all PEs. The more uniform use of memory as in the right panel is a key to efficient computation of forward operators in DART 2.0.*

5. MPI2 one-sided communication

The previous implementation of DART required an all-to-all data transpose of the model state. DART 2.0 replaces local array access with off-core memory retrieval using passive target MPI2 one-sided communication. This removes a previous hard limit on the maximum problem size. The use of one-sided communication allows any task to request data from another task on demand without needing to sychronize.

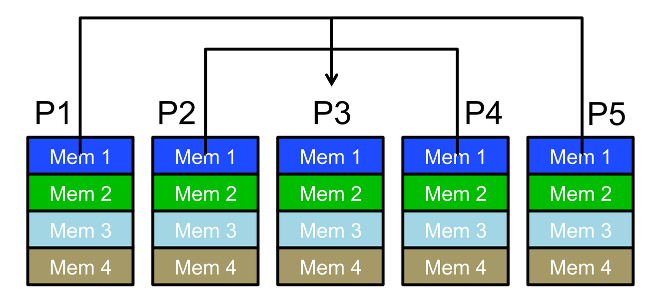

Figure 4: *P3 needs data from other PEs for computations. With MPI2 one-sided communication the other processes are not interrupted. By using more PEs each PE has fewer observations to manage. The forward operator for all ensemble members is computed at once so the code vectorizes better.*

5.0.1 Impact to Code

Most user-visible changes look like the following excerpt from the lorenz_96/model_mod.f90. Note that the new operation *vectorizes*.

Old: (the ensemble state is in array 'x')

```
obs_val = (1.0 - lctnfrac) * x(lower_index) + lctnfrac * x(upper_index)
```

New:

```
x_lower(:) = get_state(lower_index, state_handle)
x_upper(:) = get_state(upper_index, state_handle)
expected_obs(:) = (1.0 - lctnfrac) * x_lower(:) + lctnfrac * x_upper(:)
```

6. Performance of DART 2.0

A series of tests were performed on NCAR's Yellowstone with the Weather Research and Forecasting model (WRF) with 50 ensemble members and 54000 observations. Each ensemble member has a state of 184 million elements, which is approximately 74 GB. Yellowstone has Intel Sandy Bridge processors, a Mellanox InfiniBand, 25 GB of user-usable memory per node, and 8 dual-socket cores for 16 CPUs per node.

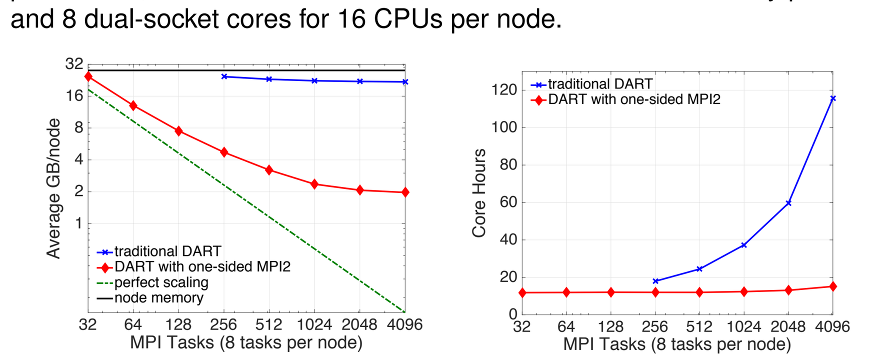

Figure 5: *Left: strong memory scaling for DART 2.0 for a large WRF model. Right: Job core-hour cost almost flat as number of nodes increases. Adding more cores results in a corresponding decrease in wall-clock time.*

6.1 Native netCDF support for I/O

DART 2.0 directly reads and writes netCDF model restart files. Since netCDF is such a common format, this eliminates the need for most models to convert model restarts to a DART format. This reduces the complexity of the cycling scripting and reduces the high-water mark of disk storage required. Models that do not use netCDF will still need conversion routines. Diagnostic files have always been in netCDF and will continue to be so. Ancillary files for things like inflation values are also now in netCDF rather than the old DART restart file format.

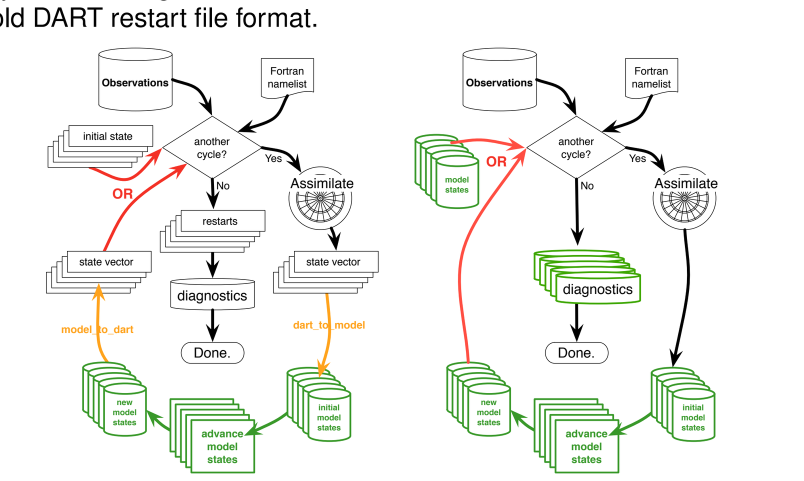

Figure 6: *Left: existing pattern that requires conversion of the model input/output to an intermediate DART format. Right: new pattern for models that use netCDF files for I/O. There is no need for specific DART restart files.*

7. ROMS with DART 2.0

Andrew Moore (ammoore@ucsc.edu)
Hernan Arango (arango@marine.rutgers.edu)
Chris Edwards (cedwards@ucsc.edu)

DART-ROMS is being developed as part of a multi-institute project funded by NOPP with the primary aim of coupling ROMS to COAMPS. The coupled ROMS-COAMPS model will capitalize on the data assimilation capabilities of the two models: firstly through DART-COAMPS, and secondly via the 4D-Var capabilities of the two models. While the primary focus of the NOPP project is to understand intraseasonal variability in the Indian Ocean, DART-ROMS is undergoing initial development and testing using an existing tried-and-trusted test case, namely the circulation along the North American west coast. In the example presented here, (Figures 7,8) a low-resolution version (30 km) of the ROMS test case is represented and is forced in uncoupled mode by surface fluxes derived from COAMPS fields. Real observations of temperature and salinity from a variety of satellite and in situ platforms were assimilated into ROMS using the prototype DART-ROMS driver.

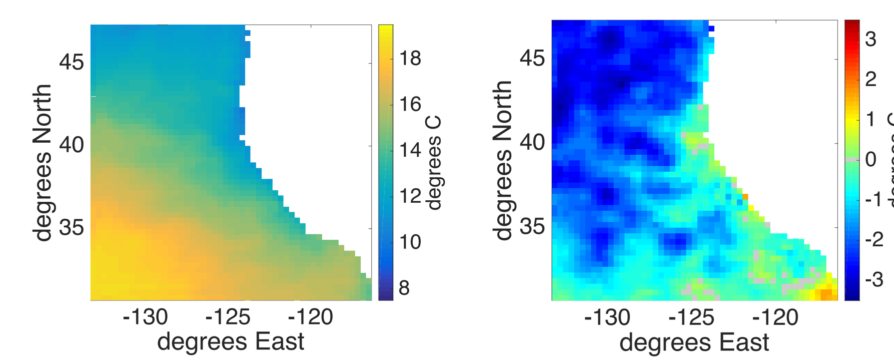

Figure 7: *Left: Ensemble mean temperatures of the upper 10 meters. Right: Ensemble mean of the innovations for the temperature.*

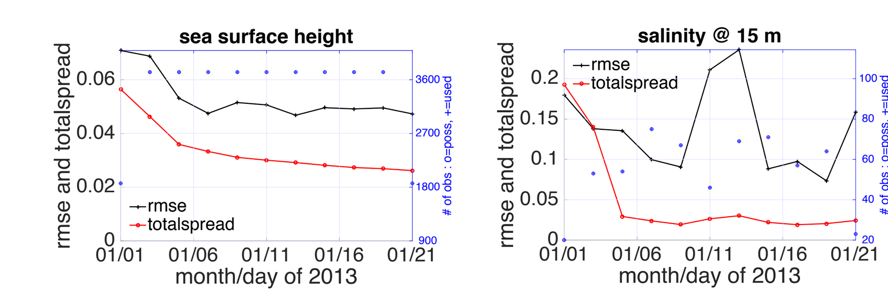

Figure 8: *Left: Root-mean-squared-error (RMSE) and totalspread (the square root of the pooled variance of the observation and the ensemble) for the sea surface height. 'Error' is defined to be the difference between the observation and the ensemble mean. Right: RMSE and totalspread for all salinity observations between 10 m and 20 m deep.*

8. TerrSysMP

Prabhakar Shrestha (pshrestha@uni-bonn.de)

TerrSysMP was developed within the Collaborative Research Center TR32 to improve understanding of patterns of energy, water and carbon fluxes and feedbacks in the terrestrial system. TerrSysMP is a modular platform, which uses an external coupler (OASIS3/OASIS3-MCT) and three component models for the atmosphere (Consortium for Small-Scale Modeling; COSMO), land surface (the NCAR Community Land Model; CLM 3.5) and groundwater (ParFlow). The test configuration for a weakly coupled DA experiment with TerrSysMP/DART is a periodic 30km domain. OSSE's are being carried out as a preliminary step toward strongly coupled DA.

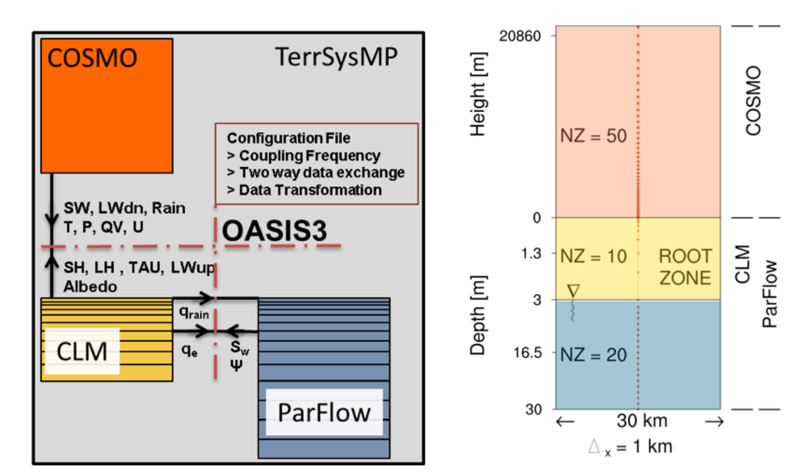

Figure 9: *Left: Schematic of the TerrSysMP system. Right: A vertical cross-section of the model domain. Of particular interest for strongly coupled DA is the localization across the model boundaries.*

9. FESOM Marmara Sea OSSE

Ali Aydogdu (ali.aydogdu@cmcc.it)

An unstructured mesh ocean model (FESOM) is interfaced with DART for an ensemble data assimilation system in the Turkish Straits System. Synthetic temperature and salinity observations are assimilated along the tracks of the ferries in the eastern Marmara Sea with a six-hour assimilation cycle. There are 30 ensemble members using a horizontal and vertical localization halfwidths of ≈3 km and ≈25 m respectively. The observation temperature and salt error variances are 0.5 °C and 0.25 psu, respectively.

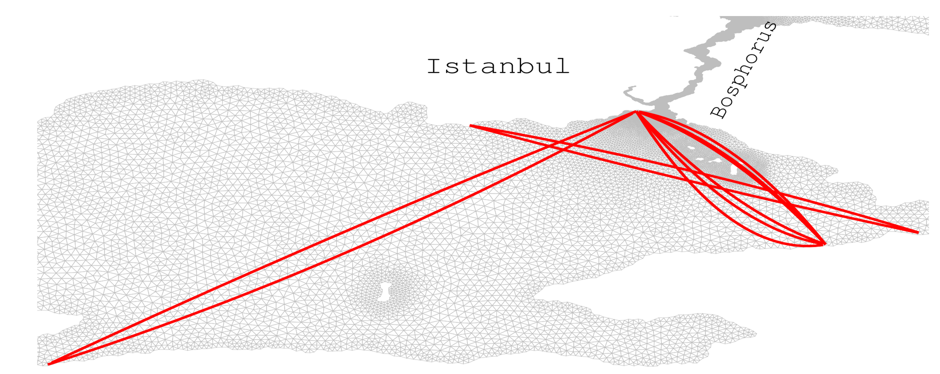

Figure 10: *The unstructured mesh and the ferry tracks for a small portion of the study area. Synthetic observations were created for every location of the ferries at 1 minute intervals. All observation locations were at 5m depth.*

Figure 11 compares the RMS of salinity difference (psu) between the prior ensemble mean and the nature run in the first 10 m of the water column after the 16^{th} assimilation cycle. The experiment with assimilation shows an improvement in the eastern and central basin.

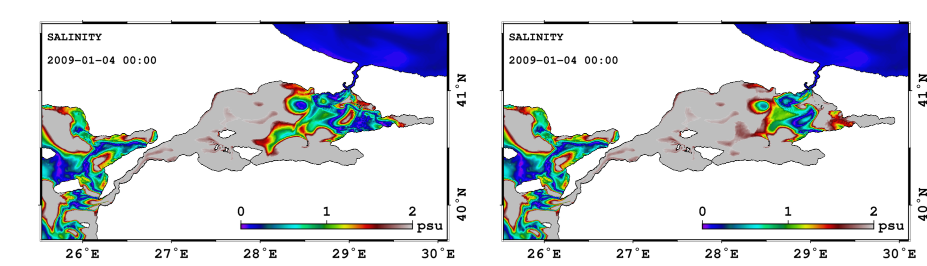

Figure 11: *Left: Assimilation Run. RMSE of the ensemble mean and the true state of the salinity after 16 assimilation cycles (4 days). Right: No Assimilation Run. RMSE of the ensemble mean and the true state of the salinity at the same assimilation cycle. Note that gray denotes a larger difference.*

10. CICE5 with DART 2.0

Cecilia Bitz (bitz@atmos.washington.edu)
Yongfei Zhang (yz4362@atmos.washington.edu)

Perfect model observing system simulation experiments (OSSEs) are conducted to investigate data assimilation and post-processing methods via DART/CICE5[1]. The goal is to improve sea ice thickness (SIT), which is a crucial state for sea ice forecast. Experiments assimilating synthetic observations of sea ice concentration (SIC), sea ice age (AGE), and SIT are compared with a free run. Figure 12 shows that SIC-only assimilation performs well along the ice edges but not in the central Arctic. The joint assimilation of SIC and SIT removes SIT error almost everywhere; however, SIT satellite observations are not readily available. Promisingly, assimilating AGE together with SIC improves SIT simulation.

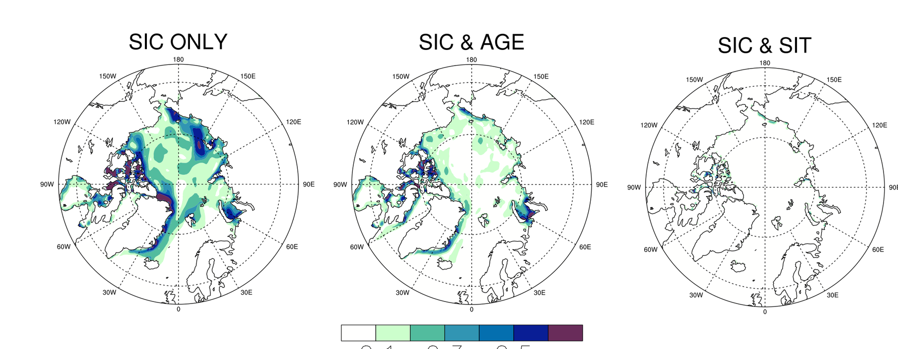

Figure 12: *Left: RMSE of the SIT when assimilating SIC only. Middle: RMSE of the SIT when assimilating both SIC and AGE. Right: RMSE of the SIT when assimilating both SIC and SIT. The units are meters. Typical sea ice thickness values are 1-3 meters.*

1. The Los Alamos Sea Ice Model is often referred to as the **C**ommunity **Ice** **C**od**E** (CICE).