2019-01-24

The meeting began a few minutes late because of snow delays in Boulder.

Dan began by announcing that he reverted to the previous fix for the 4D-enVar implementation in JEDI and asked for other modelers in the UKMO and MPAS groups to test it. Yannick said that he has been working on a long-term solution that he will push today. In this branch, he fixed an issue where repeated calls to the Model were re-initializing the GeoVaLs. He eliminated these unnecessary calls in 3DVar and 4Dvar.

Other issues that Dan has been working on include the refactoring of the State and Increment classes in FV3 and keeping pace with recent merges in ioda and ufo.

Guillaume has been working on various issues concerning the vertical resolution, including new runs with a different vertical grid. He asked about how to handle ensemble members that have a different grid, including the specific case where different ensemble members may have a different vertical coordinate. One way to handle this would be re-mapping when the fields are read in but he's asking if there might be a better, more generic way to handle this. Yannick commented that this is an issue with the solver as opposed to the H(x) but no clear solutions were forthcoming.

Guillaume also asked if there is a way to pass the increment back to the model. Yannick mentioned that these were deliberately kept separate. Chris S asked for an elaboration of the problem and Guillaume responded that it has to do with the analysis not being on a different vertical grid. Chris said they do a restaggering for MPAS when they add the increment back in to the State but the issues with the ocean models may be more subtle.

Guillaume also mentioned that the new QC implementation does not work with the marine conventional observation data. He has not had the chance to debug it yet and he agreed to speak with Yannick offline.

Hamideh is working on the H(x) for GMI and she's finding that CRTM is returning spuriously large values. She is now investigating the problem.

The floor then turned to the UK Met Office. Steve S reported that he finished support for variable changes in LFric so now the user can pass a desired array of strings and the appropriate fields will be created. He also asked about the state of the ioda workshop - see below.

Marek has been working on putting the identity operator into ufo and has been testing it with a 3D FGAT application. It's close to working but there are outstanding issue with O-A (observations - analysis). He also had some questions about how to deal with the increments. He said they want to be in a position where the increment is zero if there is no DA. He also identified two steps in the application of the increment, namely reconfiguration of the collocation grid to the model and the adding of the increment to the model state. Yannick and Chris S mentioned that these two steps are lumped together in NEMO and MPAS.

Marek also sought clarification on what changes might be needed for GetValues() in association with a planned change in the way the linearization state is stored. He was worried about last-minute changes that might interfere with achieving their deliverables. Yannick said that we want to move toward a common procedure in GetValues() of doing the variable transforms first and then the interpolation but this has not yet been coded. Yannick assured Marek that these changes should not require substantial changes on the model side.

Marek then asked for feedback on David Rundle's implementation of RTTOV into ufo. Anna recommended that before we merge it, we should define two different obs operators for RTTOV and CRTM. She offered to work with David to define a new C++ class for RTTOV.

Steve S also asked if there was any documentation available on the scope of the QC - what is included and what isn't? Yannick responded that no such documentation yet exists.

Then there was a report from NRL. They are working to get the ioda netcdf file conventions working on their systems and plan to work with Steve H on this.

Nancy then brought up the issue of whether or not we should reschedule the ioda workshop planned for February 11-14.

Yannick announced that a decision on the ioda workshop will be made tomorrow (Friday). An email will be sent out to all concerned individuals when a decision is made.

Based on news stories, it is unlikely that the US government will re-open by tomorrow. If that is the case, then we will indeed have to postpone the ioda workshop in order for NOAA and NASA employees to have sufficient time to re-establish their normal work routine and make travel plans.

JJ reported that we is still working on the splitting of the State and Increment objects and, in particular, whether or not a separate variable list is needed for each of them. Yannick asked why this is the case and JJ answered that it had to do with the availability of variable lists during various tests. It may be possible to reconfigure the tests so that such a separation is unnecessary. Yannick agreed to discuss this with JJ offline.

BJ tested the current 4D-envar implementation with MPAS and it seems to work. However, there was a problem during the second analysis. He suspects this is a problem with MPAS rather than oops and plans to investigate it further.

Junmei reported that she is investigating the quality of the temperature assimilation in MPAS. This has to do with a previously-reported issue that the temperature increments seem to be much larger than those for u and v. She has been conducting sensitivity tests at different resolutions and suspects that the analysis fit to observations might be improved with an improved B-matrix. She is using UKMO observational data sets and a B-matrix generated with bump. Yannick expressed appreciation for this work, saying that it's one of the first times anyone has really looked into the scientific quality of the DA, and encouraging her to continue. He asked about some of the details of the bump implementation and BJ reported that they were using 5 members for the coarse resolution and 20 members for the fine resolution. Yannick encouraged them to check with Benjamin to see if this is sufficient.

Steve V has been benchmarking Mongo DB for possible use as a database tool for ioda. Yannick and Tom cautioned that we need to consider carefully whether or not we want to proceed with this. Though Mongo DB is open source, it is subject to an aggressive, contaminating license. This could potentially make it unusable. Tom mentioned that Mongo DB was dropped by Red Hat because of these licensing issues.

Anna then reported that there are two feature branches for both ioda and ufo that will likely be merged today. The conventional obs files have been updated with the obs errors fixed. Furthermore, the full observation and GeoVaLs files for our April 15, 2018 test date have been completed and they have been made available publicly through the JCSDA public data repository at http://data.jcsda.org. For further details on what these files contain and how to access them, see the note that Anna published today on the JEDI GitHub Team feed. Only the largest, full observation files are on this data repository. Smaller data files used for testing have been included in the ioda and ufo GitHub repositories by means of git-lfs.

Hailing reported on work she has done to create data files for GPSRO for use with testing the ROPP obs operator in ioda and ufo. She has already added one new unit test that will be pushed soon.

Francois has been working on a bug in the adjoint for the GNSSRO implementation and has been investigating in particular with intel compilers on Cheyenne and Discover.

Ming is working with Xin on WRF. In particular, they recently updated the use of the Locations class to be consistent with the latest versions of ioda and ufo.

In addition to this, Xin has developed a TL and AD implementation for WRF that now compiles but has not yet been fully tested.

Mark announced that a new (pseudo) Random Number Generator (RNG) is now merged into the develop branch of oops and is available for use. The C++ implementation can be accessed with **#include oops/util/Random.h** and the Fortran implementation can be accessed through use random_mod. For C++ and Fortran usage examples, see the files Random.h and random.F90 in oops/src/test/util. Developers are encouraged to use this implementation because it is designed to be reproducible for different compilers and platforms, producing a reproducible sequence of pseudo-random numbers for a given random seed. Furthermore, there is a new test in oops to confirm that the RNG is indeed reproducible and to alert the user if this requirement fails. The qg, l95, and FV3 models have already been modified to use this new RNG implementation. Mark will work on associated modifications of ioda and ufo in the coming days.

Mark also announced that work in in progress to set up a node on Amazon Web Services that is equipped with a license to use the intel compilers and MPI library. This should enhance our ability to test JEDI with intel compilers without relying on access to Cheyenne, Discover, and Theia.

Maryam has been familiarizing herself with the boost test framework that is now used in JEDI. This will allow her to replace the boost unit test framework in the coming weeks with equivalent unit test functionality offered by eckit and fckit.

Steve H has been working more with the tools in the ioda_converters repository designed to convert data from a variety of sources into netcdf files that can be readily read by ioda. He is now arranging for these various converters to use a common netcdf writer in order to facilitate future maintenance and implementations. Contact Steve H if there are other data formats that you would like to add to this pipeline.

Clementine is still debugging the block minimizer. There are still issues with dx and dxh.