

2018-08-23

Agenda:

Our usual weekly meeting schedule has been disrupted the last several weeks by the JCSDA Summer Colloquium and the B-Matrix code sprint.

The agenda for today's meeting is to re-establish our regular weekly schedule and to review progress since we last met.

We started by going around the table in Boulder

Yannick

There are various code changes that will likely be merged into develop this week or next week that you should be aware of. **If you have any issues with any of these branches please let us know, preferably before they are merged:**

(oops) develop-bump: This already has a new MPI implementation developed by Chris and Yannick. It also has the bump implementation developed by Benjamin M. We want to merge this into develop next week if there are no outstanding issues

(oops) develop-bmatrix: This includes code developments implemented during the recent code sprint. Again, if you are aware of any issues with this, please let us know before we merge it

Also - There is another code sprint next week for implementing GNSSRO into jedi. There are several recent code developments that have to do with this, including an **update to ufo** that was just merged today.

Chris has been working on the new MPI capability in oops and on MPI-related problems users have been having with the mpas and Ifric test suites. He has not yet been able to reproduce these problems yet but he is working on it. He suspects that the problems may have to do with calling MPI functions before MPI is initialized. As a preliminary measure to address this, he has added an MPI initialization when runs/applications are initiated.

Rahul asked if serial code can be run without calling mpi_init. The answer is yes, provided that you do not run the executable with mpirun.

Rahul also asked if there is a way to pass an external MPI communicator to eckit. Chris and Yannick responded yes.

Steve is now creating a suite of observational data files that share a common start date: April 15, 2018 00:00 UTC, for use in testing. When completed there will be three distinct sets of data files

1. All of the observations that come from GSI, saved as GeoVaLs objects
2. A single observation - for debugging purposes
3. An intermediate-size set that includes several hundred observations


This should be done in the next few days and is coordinated with the full observation (feature/fullobs) branch of ioda that Xin has been working on

Francois mentioned that he already has a GNSSRO buffer file for that date that he can provide, in a data format used by the models (not yet in ufo standard format)

Xin reported on two developments he has been working on. The first is the multiple-observation capability (fullobs) in ioda. It is nearly ready to be merged but it was decided that it might be best to defer merging this into develop until after the GNSSRO code sprint next week, to avoid any potential unforeseen complications. The second project Xin mentioned is a restructuring of the ioda memory structure: he has a prototype of this already and he is now cleaning it up.

Marius commented that something appears to be missing from the TL/AD observation operators in CRTM. Ben was not sure what he was referring to - everything should be there. They agreed to discuss this further after the meeting.

Anna has been working during last week's code sprint on enhancing the Geometry and State classes in oops. She has also been developing observational operators for ensembles and plans to introduce a new class that can handle ensembles of observation vectors.

Marius then mentioned that Craig Bishop has developed a gig filter  to account for non-Gaussian statistics and asked if this should be implemented. Yannick mentioned that this lies outside the current scope of the project and may be considered in the future

Francois discussed preparations for the GNSSRO code sprint that will occur in Boulder next week. Preliminary work has been completed on a 1D observation operator and merged into ufo. This sets the stage for the 2D operators that will be developed and implemented in the code sprint. During the development of the preliminary operator, they found that ioda was crashing when trying to read a file with a missing field. They implemented a fortran pre-processor to filter out these missing fields. Francois asked whether this should be added as a utility to one of the JEDI repositories and if so, where it should be. Yannick strongly advised against this, encouraging the team to follow an approach that was more like the other observation types, if possible.

In response to questions, Francois emphasized that several branches of ioda will be created in the code sprint and will be visible to all on GitHub. The sprint will focus on implementations of non-local operators specific to GNSSRO but Anna will help to generalize these to other observation types

Mark announced that new tools have been developed for those interested in running JEDI or other JCSDA applications on Amazon Web Services (AWS). In particular, we have now made three Amazon Machine Images (AMIs) publicly available:

- jedi-ubuntu16.04-singularity (us-west-2): This is an ubuntu OS that comes with Singularity pre-installed so you can proceed to pull the JCSDA singularity image and then work in the container
- jedi-amazon2 (us-west-2): This is an Amazon Linux 2 OS that comes with jedi relevant packages pre-installed (boost, netcdf, ecbuild, eckit, fckit, etc), outside of the singularity container
- ufo-bundle (us-east-1): similar to the amazon2 AMI but on an ubuntu OS (outside of the singularity container)

For instructions on how to use, these, consult the AWS section of the JEDI wiki or ask Mark

Then the floor was opened up to external participants.

NCAR

Jake mentioned that there was a change in the mpas code bundle to use temperature as a control variable instead of potential temperature . He says they are still having problems running mpas-bundle on Cheyenne but work continues on addressing them. He mentioned that Junmei created a ZenHub issue (mpas board) associated with the ODB capability in ioda. Steve posted a work-around to this issue on ZenHub. Jake asked if he was using a different branch of oops but Steve assured him that he got the work-around to work using the develop-bump branch of oops, running within the singularity container on his Mac.

A question was raised as to why eckit is included in the container - the reason is just to speed up the build.

BJ brought up an issue with the univariate 3DVar covariance implementation with bump in oops and Yannick suggested that Benjamin M needs to have a look at it.