

2019-10-03

Yannick opened the meeting by announcing several changes in the JEDI repos that **will require changes to bundles, model and possibly other repos.**

The first set of PRs are to saber (#29) and oops (#416) and involve the mpi functionality in saber (feature/remove_bump_mpi_calls). This merge eliminates direct MPI library calls in saber so that all mpi is now directed through fckit (as it was before the saber merge of last week). If you have had problems compiling saber in the past week because ebuild can't find mpi, then these PRs should solve those problems. In order to support this change, we have already modified the fckit develop branch and **compilation of fckit now requires you to use the jcsda fork of eckit**. See [the JEDI Team discussion board](#) for further information. Most of the bundles have already been modified to accommodate this change. All the containers (singularity, charliecloud, docker; gnu+intel) have been updated with the jcsda fork of eckit, as have the jedi environment modules on Cheyenne and S4. If you don't have the JCSDA fork of eckit installed on your system, it is recommended that you build eckit as part of your bundle. If you have any questions or problems, please contact Mark M or any other member of the JEDI core team.

We have already merged the fckit PR. **We plan to merge the saber and oops PRs on Monday.** The most significant change for model groups is that these PRs require **a small change in the argument list of bump%setup_online()**. Now an fckit communicator is passed as the first argument. For an example of how to accommodate this change, consult [the associated PR in fv3-jedi](#).

Soon after these PRs are merged on Monday, we plan to merge Clementine's feature/add_comm_geom PRs into oops and several model repos (fv3-jedi, mpas-jedi, shallow-water) on Tuesday. This changes the Geometry and Application classes and their constructors to include an MPI comm object. This is necessary to implement more ensemble applications in the future. For an example of the required changes, you can consult [the associated PR in fv3-jedi](#) or one of the other repos mentioned above. Further examples can be found in the qg and I95 model components of oops. This added communicator will be propagated down deeper into the code for use with ensemble DA. Soon this new ensemble functionality will be extended to ufo.

JJ then asked a question about when the Geometry constructors are called to build objects from scratch; in many places the copy constructors are used instead. Yannick confirmed that the creation of a Geometry object from scratch should only happen a limited number of times - most of the time the copy constructor should be used - for example, to create geometry objects for different ensemble members that are typically on different processors.

We then proceeded with a round-table discussion, beginning with the UKMO. There David announced that the UKMO will be directing more resources to JEDI in the coming weeks. They will clarify the details in Yannick's upcoming visit.

Then we proceeded to EMC. Dan reported that he has been working on implementing ebuild for the full fv3-gfs model (with physics). Hamideh has a PR in CRTM that adds the RSS emissivity model for the L-band. She's also working on adding MSS salinity operators to ufo and fv3-jedi.

Rahul is cleaning up after the recent code sprint and he also has PRs in oops and ioda that improves the naming consistency of ObsSpace (previously called ObservationSpace in oops). He also added an ObsSpace constructor that allows the user to pass a config object or handle. This will help enable zero or single obs tests. Rahul and Yannick then briefly discussed the idea of super-obbing. Rahul said it is used by NRL and Yannick said it would be possible to implement this at some point.

Cory is working on porting the GSI ioda workflow to Hera. So, jedi can now create obs and geovals files on Hera.

Jong is working on the Cl6 interface and expects to be done soon.

Next came NRL: Sarah reported that they are having some problems with 3DVar and they are setting up JEDI on a new system with better debugging capability.

Next we turned to Boulder. Clementine has been working on the PR described above. She split a large PR into two pieces. The first piece (which we intend to merge on Tuesday) focuses on adding MPI communicators the Application and Geometry classes in the oops and model repos. The second will focus on ufo, with the intention of adding capability for ensemble H(x).

Xin is nearing another set of PRs for Variational Bias correction. These seek to duplicate the GSI functionality used at EMC. He's now working on amsua and is almost done - he just needs to incorporate surface temperature. He has 12 predictors now, which is a simplified version of GSI. Yannick said that was enough for now and encouraged him to do a PR when ready. Anna added that Emily will also be working on related developments and encouraged Xin to coordinate with her.

Steve H is working on improving the ODC reader in ioda - he has hit a few snags but is overcoming them.

Steve V is working with Steve H on generating large test files for ODC-ioda.

Ming is progressing on the jedi-WRF implementation. The next step is implementing the B matrix and GSI covariances.

Jeff W announced a new team member who will be working at NOAA on JEDI: Sergei Froloff. Over the course of the meeting we also introduced three other new members of JCSDA, namely Dick Dee, Wei Han, and Nan Chen. Nan will be working in Boulder while Dick and Wei will be working mainly from New York state and Madison, WI respectively.

Mark O announced that the first version of the ioda database workflow is now working with GNSSRO data as an initial product. This workflow acquires real-time data from the COSMIC mission and stores it on Amazon S3 along with full-res (c768 gfs) fv3 background files. The files are organized into 1-hour intervals and users can construct longer data windows from these intervals. Currently the S3 database is private so only users who are part of JCSDA's AWS account have access. Mark O added that downloading the data (data egress from S3) can be a bit pricey but is free if the transfer is to an AWS compute instance located in the same region as the S3 database (i.e. the S3 bucket, in AWS terminology).

Chris S asked how the files are being converted from bufr and Mark O responded that they are using code written by Hailing for that purpose.

Hailing has a draft PR in ufo that implements GNSSRO super refraction.

Sarah (L) started at JCSDA several weeks ago and is now reaching out to scientists at NCAR and elsewhere in Boulder to collaborate on aerosol chemistry and associated DA. She plans to enhance CRTM functionality in this area and invites anyone interested in working on this to contact her.

Travis has made good progress on implementing the localization in the LETKF JEDI applications. He has 2 PRs now, but there is still more work to do.

JJ has a PR in mpas-jedi that enhances the diagnostic outputs. Like the previous diagnostics, this uses python scripts, but the new implementation makes use of pandas dataframes to represent and work with the data. Enhanced diagnostics include information on rms values, bias, and other features. The intention is to eventually generalize this and add it also to the DAdiagnostics repo for potential use with other models. If you have any questions or if you'd like to work with them to develop these diagnostics, contact JJ and/or Junmei. Mark O responded that he is interested. He also asked if the diagnostics will be able to process data in ioda format and JJ responded yes.

We then turned back to the UKMO where Marek had joined the meeting. He has been working on a concatenation operator. He has also been trying to sort out some problems that they are still seeing with the saber PR that went in last week. It seems to be having problems calculating the covariance parameters.

Back to Boulder with Maryam, who announced further progress in implementing AWS CodeBuild for JEDI repos. It is now fully implemented for ufo, ioda, and saber, and will soon be fully implemented in oops. CodeBuild tests each repo with both gnu and intel compilers using the same docker containers that are used to build the JEDI Singularity and Charliecloud containers.

Yannick then closed the meeting by announcing a more formal commitment to the weekly meeting structure that we have discussed previously. So, now the round-table discussions such as this are scheduled as a bi-weekly meeting and they will alternate with more focused discussions. This is why the calendar invitation now lists the round-tables as a **JEDI Bi-Weekly Meeting**, alternating with the bi-weekly **JEDI Discussions**. So, we will still have a meeting every week but the theme will vary. There is a loose plan to alternate the focused topic discussions between model-related issues and observation-related issues. So, in other words, there will be approximately one meeting a month devoted to observation-related topics and one devoted to model-related topics. But this pattern is loose - many topics will span both observations and models.

As always, please let us know what topics you would like to discuss, either by [editing the ZenHub board](#) or by contacting a member of the JEDI core team.