

2020-09-03

Yannick opened the meeting by announcing the agenda of a general round-table update.

The only general announcement that he had was for everybody to keep an eye out for code changes that may require some adjustments from the model groups and others. Many of these are minor and are associated with continuing efforts to clean up the JEDI code in preparation for the public release. An example is Yannick's recent [MPI cleaning PR in oops](#). This requires a small change in other repos that should only take a few minutes to make.

There will likely be other similar PRs coming. So, we ask for your help and your patience in making these small adjustments when asked. For example, if you are a member of a model group and you are asked to review a PR in oops or another base repo, please do so as soon as you can so we can get these code changes in.

Then we proceeded with updates from the various JEDI teams.

Before summarizing yesterday's JEDI 1 (Infrastructure) project meeting, Mark M had an announcement. New modules have been installed on the HPC systems including S4, Orion, Discover, Cheyenne and Hera (Hera is still in progress but should be available by the end of the day). The main purpose of this module update is to provide support for C++14 with Intel. But, it also includes updates to hdf5 (needed for ioda-engines), ncep-bufrlib, netcdf, lapack, ncncmp, and a few others. [Instructions for using these new modules have been posted on the JEDI GitHub Team page](#). **Please start using them and let us know if you have any problems.** If there are no problems, we will soon make these the default modules.

Then Mark M proceeded with a summary of the JEDI 1 meeting. We introduced Claude Gibert (new JEDI core team member) and we discussed the new [sol](#) repository that Claude introduced for jcsda. This contains general-purpose python utilities, including tools to work with date/time objects that Claude developed for workflow purposes at ECWMF. In the future we will discuss further how to integrate this with the other JEDI python dependencies including those in the ioda-converters and diagnostics repos.

Then Mark M summarized several other issues that were discussed at the JEDI 1 meeting. Much of the discussion was about the jedi-stack, including finalizing package versions in preparation for the public release and providing support for C++14. After the meeting Mark M introduced new tagged versions of ecbuild (3.3.2.jcsda3) and eckit (11.6.1.jcsda2). These will be incorporated into the jedi-stack for the release. There was also some discussion of current jedi-stack PRs to update ncep-bufrlib and add GSL (needed for ioda-engines). We also decided that odc is no longer needed as part of the default jedi stack. After these and a few other PRs are merged into jedi-stack, we will tag it with v0.4 and use this tagged stack for the release (e.g. in release application containers).

Maryam reported that the pricing model for Codecov has changed and we will not be able to continue with our current codecov PR test functionality unless we pay much more than we do now. Maryam is looking into other options for providing code coverage information, including [coveralls](#) and cdash.

Mark O is working with the MPAS team on a new cmake build system for MPAS that can be more easily integrated with JEDI. He expects to have a PR ready by the end of the week. Mark M wants to prioritize this for the release for use with tutorials and mini-academies.

Ryan then proceeded with the following updates for Hera and JEDI2:

Hera update:

Stack should be built by the end of the day. Need volunteers to test. The big issue was C++14 support here. The sysadmins provided a version of gnu 9.2, which we can link against, but apparently it's too new for Intel 18, which is the default Intel compiler used on Hera. Basically, it can't parse GCC's headers correctly in a few instances, particularly for the Microsoft GSL library. Testing it now.

JEDI2 update:

JEDI2's meeting was last week. At the meeting, Steve gave a demonstration of the basic features of ioda-engines ([see slides here](#)). At the end of the presentation, we referred people to the five ioda-engines examples, located in the repository under the Examples subdirectory. These should give everyone an idea of how to use Groups and ObsGroups to manipulate data.

Steve and I are continuing work preparations for the upcoming JEDI release. Ioda-engines will power the ObsSpaces by then. He and I have both had several discussions with the Obs team planning out the roadmap for migrating both UFO filters and data to the new ioda interface.

We are both working on a converter program that converts the old ioda format to the new ioda format. Ideally, you can just pass in a file as an argument to the program and it updates it with no further interaction.

For the rest of the JEDI2 group, Xin is continuing work on VarBC and plans on finishing in September. Plans for two pull requests, on bias coefficient α and the B matrix for bias. Wojciech is continuing refactoring of OOPS classes for YAML validation. PR merges held up a bit by C++14 support across HPC systems. The UKMO otherwise are working on GNSSRO TL/AD codes, 1DVAR filters, and profile and buddy qc checks.

Mark O asked about the plan for migrating ioda-converters to the new ioda-engines file format. Steve H said he has a feature branch in ioda to promote a smooth, transparent transition to the new format (PR will be submitted soon). This will be backward-compatible with the old format and will therefore allow for an incremental approach for migrating the converters to the new format. Mark O then asked specifically about the output files used for diagnostics and whether a yaml key could be introduced that would allow the user to select between the old and new formats. Steve H replied that they will still use the old format for now but it should be possible to provide such a yaml key in the future when we begin to migrate the ioda diagnostics to the new formats.

Dan then proceeded with a report from JEDI 3. There was no biweekly meeting this week so Dan just reviewed what the team is working on. In the case of fv3-jedi, the team is working with Mark Potts to interface with UFS. They will mimic a forecast in a constructor for now. Dan also said that many members of the fv3-jedi and other modeling teams are working on static B training and tuning. Since this involves both the JEDI 3 and JEDI 4 projects and since it involves the generic DA algorithms in oops, it was decided that this static B work will be discussed mainly in the JEDI 4 meetings in the future (see Anna's report below). In the case of fv3, this includes work on the variable transforms and on cold starts, including the workflow aspects of generating psi and chi for ensembles.

Marek added that he has been working on adding a local area model (LAM) capability to UM-jedi.

Anna was not able to attend the meeting but she provided the following summary for JEDI 4, delivered by Yannick:

Background error covariances:

- BUMP: Benjamin (on PTO during the last JEDI4 meeting) has a big pull request in saber (remove global arrays), that we should probably merge soon. **Mo delers - please check that this PR works and review**
- um-jedi: Oly is working on background error covariances in um-jedi. Single PE wind transform is done, he is working on the MPI version.
- mpas-jedi: BJ is working on variable transforms for MPAS-jedi and using vertical covariances from BUMP.
- fv3-jedi: Dan is working on static B training for full-resolution GFS (C384 ensemble from Jeff Whitaker). He is adding variable transforms to convert available cold start files to the files that can be used with fv3-jedi.

Ensemble applications:

- Clementine is working on consolidating Ensemble Application (single application to replace EDA, EnsHofX, EnsForecasts). The Ensemble Application will likely use separate yaml files for each ensemble member. Clementine is also working on different thinning for different ensemble members.
- EDA in mpas-jedi: JJ is running EDA experiments with 120km resolution MPAS, exploring RTPP inflation. At the last JEDI4 meeting the group commented that EDA experiments still need active bias correction, QC, and realistic observation perturbations to be more scientifically interesting.

LETKF-type solvers:

- Sergey's PR adding solvers for NOAA LETKF and GETKF is merged into oops.
- Travis is running experiments with LETKF and 3DEnVar in soca; he has scripts for cycling the system with soca (working on laptop and with slurm). Travis was hitting bottlenecks in LETKF, particularly with the search of local observations (brute force was the only option up until last week).
- Helen Kershaw (PSL) implemented search for local observations in LETKF using KD-tree, it is now merged into ioda. This significantly sped up Travis's experiments with LETKF: for ~800K obs local obs search went from 65% of total time to 1% of total time of LETKF. Now the main bottleneck is in the GetValuees. Travis, Helen, Sergey and Anna are investigating what can be done to improve this.

Variational DA:

Yannick is working on cost function refactoring that would allow for time-parallel 4DEnVar and weak-constraint 4DVar, and significantly simplify the cost functions code. This should not require any changes from the model side. Benjamin's PR adding hybrid 4DEnVar got merged to oops.

Yannick then proceed with a brief summary of JEDI 5. Much of the work here is focused on the release, and cleaning up code as mentioned above. He has received many volunteers to participate in the upcoming virtual code sprints for writing tests (Sept 21-25) and documentation (Oct 5-9). Email him if you haven't already and if you are interested in participating.

Then Yannick opened the floor for other general updates.

Mark O reported that the NRT website has been enhanced with new obs sources. Now we are getting GNSSRO data directly from bufr files and the web display offers more detailed plots and information. Soon we expect to have vertical profiles as well. GIIRS data has also been added, getting the data directly on S4. The OBS group is working on hyperspectral plotting tools to enhance the diagnostics further. They are now working to incorporate IASI.

The floor was then open for other comments and questions but there were no takers so the meeting was adjourned.

Next week will be a focused topic meeting: Steve H will lead a tutorial on writing a converter compatible with the new ioda-engines.