

2019-05-02

Yannick opened the meeting by asking users to please check the feature/preqc branches of oops and ufo because he would like to merge them soon, if there are no objections. He plans to add documentation to the JEDI [ReadtheDocs pages](#) demonstrating how these pre-qc filters should be used.

Then there was some discussion of what units we should use for surface pressure. We should make a decision and be consistent. Travis mentioned that the cf convention uses Pascals. Anna mentioned that the models use consistent units but it's not Pascals. It was decided that we should follow the cf convention in the future and ioda files should use Pascals as the unit for Pressure (including surface pressure). We can do a conversion in JEDI if needed but it would be better if the data was stored in Pascals in the ioda files. We should coordinate this change so that it is made across models and ioda at the same time. When Steve returns from vacation next week, we will plan how to do this. Let us know if you would like to participate.

Then Yannick reminded everyone of the upcoming [17th JCSDA Technical Review Meeting and Science Workshop](#), that will be held in Washington DC at the end of this month, May 29-31. Registration is still open until May 17 (though the presentation deadline has passed). Yannick encouraged everyone who has not already registered to do so. There is no registration fee.

EMC/GMAO

Dan has been implementing multi-resolution applications for fv3-jedi that use a different resolution for the inner and outer loops. He has it working now but bump does not yet have the required functionality so he is currently using an identity B-matrix. He's working with Benjamin to implement the required functionality. Yannick mentioned that the ensemble members have the resolution of the inner loop so bump should be able to handle this. Dan responded that the missing piece is allowing for a different resolution for the ensembles and the increment.

Then there was a report from Rahul and Guillaume regarding SOCA. They are busy fixing recent test failures in ctests, adding new operators, and enhancing their cycling runs. Two PRs are in the works. Hamideh has been working on the cool skin operator, with active pull requests for ufo and soca that are currently under review. She has also been working on implementing GMI geovals in SOCA and ufo and comparing them with GSI. Yannick noticed that some of the files had points with invalid values, such as infinity, and these needed to be removed.

Boulder

Maryam announced the initial deployment of a new Continuous Integration (CI) system for JEDI based on [Travis-CI](#). The idea is that whenever there is a pull request on JEDI repositories, Travis CI will automatically initiate tests. There is also an option to manually initiate tests. This is integrated with GitHub with a nice user interface so you can see the results of the tests as they are running. It will be activated in oops first and then it will be applied to other repositories. In a future weekly meeting (TBD), Maryam will demonstrate how it works.

Please let us know if you get emails from Travis-CI. When repositories are linked to Travis-CI, automated emails are sent out, often asking GitHub administrators to authorize Travis-CI access to repositories. You may also get requests to authorize access to personal user accounts on GitHub which you can do if you wish but it should not be required - we're still experimenting and getting to know this Travis-CI interface so please keep us informed (email Maryam) on the information you are receiving and any questions you may have.

Travis is working on cleaning up the SOCA test suite. He has made them more efficient so the total test time is down to about 30 seconds. He also asked the JEDI community about **adding tolerance to the compare.sh script**. In short, many of the JEDI application tests judge pass or fail by writing the output to a text file and comparing it to a reference text file that is included in the repository. This is just done with the unix/linux command-line utility diff. Travis would like to enhance this so that this comparison is done with a certain numerical tolerance and **asked if anyone had any suggestions on how best to implement it or if anyone has worked on it**. Yannick mentioned that this is a long-standing issue and it would be great to have this in.

Xin has been working on his redesign of the ObsAux Class in oops, responding to review comments on his current PR.

Hailing has continued to work on the bending angle GSI operator. The forward operator is almost done and she's working on generating RO geovals.

Mark has been working on developing a new module-based build system for jedi dependencies, currently housed in the jedi-stack repository (based on previous work by Rahul). The motivation is to have a common test and application environment for the containers (Charliecloud and Singularity), the cloud (currently AWS), and HPC systems. The jedi-stack build system is working on AWS and Mark is now using it to change the way the containers are built. Currently, the strategy for selected HPC systems is to provide environment modules but we'd eventually like to also provide the capability to run the containers on HPC systems with no overhead. As a step toward achieving this, Mark has included support for infiniband hardware interconnects in the mpi implementation in the container. This runs now in vagrant but is giving some test failures so it still requires debugging. In any case, **if you have any software that you would like to add to the container, this is a good time to add it** (just email [Mark](#)). Several people mentioned nccmp and valgrind and these are already included in the jedi-stack. Mark expects to have the new containers available by next week.

Steve V has been running the MPAS test suite and noticed that any small change in configuration (such as resolution) requires similar changes to many different yaml files. He was thinking about a good way to automate this and asked if anyone had any thoughts on this. Travis responded that they already use something like this for SOCA. yaml has built-in macros that can be used to define common parameters across multiple files. An example is the &date macros (he thinks they are called anchors in yaml?) that are used to insert a common date for the tests. See the SOCA test files for further examples and /or talk to Travis and other members of the soca team.

Yannick emphasized that editing the yaml files manually is only needed for the tests. In the future, we envision that the yaml files will be generated automatically through GUIs (likely web-based) when users launch applications. We have recently hired a new JCSDA software engineer who will start soon and who will implement such interfaces.

Yali has implemented a terrain height correction for surface pressure into mpas that can be enabled and configured through the yaml files. The initial implementation is done but there are a few outstanding questions regarding the units to use (see above), whether or not to replace the obs values in the file (see below), how the qc filters are impacted by the correction, and getting the test to run on more than 1 MPI task.

This triggered an extended discussion about the more general issue of how to handle bias and other corrections in observational data files. Yannick mentioned that this was discussed by ECMWF participants at the ioda workshop in Monterey last February. In the ECMWF observational data archive, some files have the original observational data stored along with the bias-corrected data. Other files only have corrected data. This has caused immeasurable problems at ECMWF and wastes a lot of time and resources to sort this out. Their strong recommendation was not to make the same mistake: never overwrite the original observational data. The bias-corrected data should be stored alongside the original data (e.g. another column or another file), but should not replace the observed values. We need to discuss how best to implement this but one thing is clear: we always want to save the original data. Chris asked Marek and Steve how this is handled operationally at the Met Office and they said it was done as part of the pre-processing - not in the DA.

UKMO

Steve mentioned that they had an upgrade to their linux operating systems that required changes in the applications so this has been taking up some time this week.

Marek has been working on the RTTOV interface in ufo and expects to issue a PR soon

NRL

Ben asked what Travis CI is and compared it to bamboo in the Atlassian suite. He also asked if there were any tests for GPS ROPP in ufo. Hailing responded that there are currently four tests but Yannick added that these are only enabled if the ROPP repo is available. This requires a license that not all JEDI users/developers have. Hailing is working with Francois to add documentation on the ROPP implementation in JEDI.

Ben also added that NRL does a lot of pre-processing of the observations so he appreciated the earlier discussion and wants to remain in the loop.

John M said they have made progress in refactoring the Neptune model for use with both JEDI and ESMF

After hearing reports from all participants, Yannick opened the floor for further questions

BJ asked if we save the QC flags in ufo for each outer loop. Yannick said yes, this functionality has been added with the recent qcflags feature branches that were merged in to ufo, ioda, and oops earlier this week.

BJ also asked if we call getvalues before the pre-filter. Anna responded yes and Yannick confirmed that getvalues doesn't currently check the qc flags and we need to look at this. We might want to move the data thinning to the constructor of the QC filter class.