

2019-05-23

Yannick opened the meeting and announced that the next meeting will take place two weeks from today (June 6th) due to the JCSDA Workshop next week. The agenda for the meeting was to first learn about Travis-CI in the JCSDA repositories, and then collect issues and updates from everyone.

Travis-CI

Maryam gave a presentation (see attached slide) and demo about Travis-CI integration into the JCSDA repos. Travis-CI is a tool that automatically runs tests when updates are issued to a repository. As a start, Travis-CI will be integrated into the OOPS, UFO and IODA repositories. We will let these run for a little while to flush out and fix issues, and then update the remaining repositories as appropriate. Currently, the Travis-CI integration for OOPS and UFO are working and IODA is still under development.

The tests are run automatically whenever a pull request is submitted, and whenever a commit is pushed to an open pull request. A marker appears next to the pull request name (look on the page that loads when you click on the pull requests tab) with the test status: green check mark for pass, red x for fail. A failure can be a build or test failure. You can then click on the pull request and near the bottom click on buttons to get more details about the test status.



Chris H asked if we can get tests to automatically run on HPC systems, especially using the Intel compilers. A discussion ensued which revealed the development plans for Travis-CI include:

- Travis-CI run tests on AWS using Clang, GNU and Intel compilers
 - Once AWS is working, then try for other HPC systems
- Enable the testing of all repos dependent upon the one repo that just triggered a test cycle.
 - Currently, only the OOPS tests are run when OOPS repo is updated (ditto for the other repos), and this feature would fire off all other tests for repos that depend on OOPS such as the model repos

Chris H mentioned that if AWS is testing Clang, GNU and Intel compilers, then his concern about getting tests to run on HPC systems is much reduced.

Yannick reported that Travis S has developed a scheme for keeping tabs on updates to repos that SOCA depends upon, which could be useful while the Travis-CI "run tests on the repos that depend on me" feature is being developed. Travis added that his scheme runs tests periodically (nightly?) and sets up pointers to the branches that last worked with SOCA. That is, if tests pass, these pointers are updated to the latest develop branches; if tests fail, the pointers are left at the last known good branches.

Yannick then asked for any issues or updates from the floor. He added that a BUMP PR that will affect model results is under review and we would like to get that merged in the next day or two. Please review and make the appropriate changes for your model. After this BUMP PR is merged, the BUMP code will move to its own repo which should help us manage updates more smoothly.

Chris H has been focused on getting a working software stack on Theia. This issue goes beyond the JEDI stack so it has consumed a significant amount of time. Chris may be able to get some leverage or guidance from the jedi-stack work.

NCAR

Ming is making progress with WRF integration. He has forward stepping working, but he has encountered an issue with GeoVaLs being accessible to the TL/AD process.

Met Office

Marek has an RTTOV update PR under review. It was determined that Marek's PR depends on Anna's use_pressure PR. It was then decided to merge the BUMP PR first, the use_pressure PR second, and then merge the RTTOV PR.

Marek opened a discussion concerning the integration of LFRic into the JEDI DA flows. He is getting pulled in different directions with integration requirements from JEDI and from groups outside JCSDA. Several options were offered, and it was decided to continue the discussion offline.

Steve S asked about the shallow water (SW) toy model status. He is interested in the MPI split communicator solution in the context of ensemble forecasting. Chris H reported that model forecast is working, 3DVar is work in progress, and TL/AD is working but not integrated yet. Chris mentioned that he will need help when he gets to the ensemble flow. Yannick added that the split communicator was added in now for testing purposes, with the plan being to add in the ensemble flow later.

EMC

Guillaume reported success with running 3DVar using coupled models. He raised two issues: the need for regression tests using realistic cases (beyond unit testing), and the need to get the change of variables working. Yannick added that the change of variables work will be done after the upcoming JCSDA Workshop and JEDI Academy (meaning the second half of June).

Guillaume added that SOCA is dealing with integration issues similar to the one Marek raised and they have two new programs, which run outside JEDI, that help with their situation. There needs to be some IO updates to JEDI for passing increment data to/from these programs, which can be implemented as options to JEDI. The SOCA example may be useful for Marek's solution to the LFRic case.

Hamideh reported progress with GMI assimilation.

Jong is in the thick of CICE6 integration. One unfortunate issue that has emerged is the extensive use of global variables in CICE6.

Cory is working with Marius and Dan on an aerosol assimilation poster for next week's JCSDA Workshop.

NRL

Sarah reported that they have two versions of JEDI NEPTUNE: the initial integration which can only run using one process element, and a refactored integration, which is under development, that will allow runs using multiple process elements, as well as other features.

Next meeting will be in two weeks on June 6th.