

## 2020-02-20

Yannick opened the meeting by reminding people that there will be no meeting next week. Many of the JEDI core team will be in Monterey, CA for the JEDI Academy. We will resume JEDI meetings on March 5. This will stick to the current schedule so that meeting will be a general round-table update.

We started today's update in Boulder with Steve V. He announced progress on running mpas-jedi on Cheyenne with intel compilers (using intel MPI). This has been an issue for a long time and it now appears to be solved with the latest `""jedi/intel-impi""` module (this uses intel 19.0.5). MPAS now builds and all tests pass. Furthermore, this solved another problem that has hindered the MPAS team for some time, namely a hanging issue with Parallel IO when using the gnu/openmpi compilers. Mark M later added that this update also fixed another problem. Previously, all fv3-bundle tests had been passing on Cheyenne with intel-impi with the exception of the geos tests, which were failing in the parallel netcdf writes. With the new `""jedi/intel-impi""` module, all these tests now pass.

Steve V also reported that he has been working on implementing the new unstructured interpolation option for use with MPAS. He hopes this will improve efficiency over bump interpolation.

Yannick asked Dan if the new unstructured interpolation option is now the default for fv3-jedi. Dan had temporarily disconnected but when he reconnected later he confirmed that this is indeed the case. So, for an example of how to use the unstructured interpolation, see fv3-jedi.

Maryam has been working on a practical activity for the upcoming Academy. She has also been improving the retrieval of test data from S3 to Cheyenne for ioda, ufo, and mpas. This is working now and pull requests have been submitted (ioda has already been merged).

Xin reported that VarBC works with fv3 3dVar now without the pre-conditioner. He is now simplifying the code, reducing the number of factories. He also wrote a python script to convert GSI bias correction files to yaml. Ming asked if the resulting yaml files were large. Xin answer that they are text files of about 1 MB. Ming asked if they would be easy to modify. Subsequent discussion clarified that these yaml files are separate from the input yaml files and they are not intended to be edited by hand. Rahul asked if there was a reason that the BC couldn't be part of GeoVals. Yannick thought that this wouldn't be particularly useful at this point.

Mark O has been working on new CMake functionality to install the crtm coefficients and have ufo find them. He is also resolving differences in how the crtm coefficients are used in the tests for ufo and fv3-jedi. He invites input from anyone who is interested in this. He is also modifying how CMake handle optional obs operators like ROPP. He wants to have ufo export these components. Yannick encouraged Mark to include someone in this work from UKMO who is working with RTTOV. Mark O has also been working with the compression algorithms in netcdf for both background and observation files. He says the compression works well and might offer a space saving of up to 70% on netcdf files. Some gfs workflows don't use compression currently.

Mark O also responded to a question that was brought up last week about how stack size limitations might impact the smart pointer usage in C++ that Mark was advocating for last week. It was mentioned that this could be a problem in Fortran. Mark O looked into it a bit more and found that C++ does have a fixed stack size (which you can find and expand with `ulimit -s`) and this could in principle cause problems if you're storing all your data on the stack for efficiency reasons. However, he said that this is usually not a problem and if you're running into it, then you should reconsider your code to see if there is a better way to do what you want to do. Yannick mentioned that Fortran could cause unexpected problems too if you have local C++ variables that allocate memory through Fortran.

Clementine is running EDA on S4 with full fv3 resolution and a comprehensive set of observations she received from Francois. She's using Dan's ensemble members from July with 80 members. She has it running but she has run into some problems with resources, running out of time and/or memory. It was also running more slowly than expected. She compared notes with Steve, who has run an Hofx application on S4 with c192 resolution and 9.5 million observations. They will compare their setups to see if they can solve the problems.

Steve H has been preparing for the Academy, working on an exercise that has participants running 3DVar runs and plotting the results with Dan's diagnostic scripts. He asked for some help from anyone who may be interested in designing good experiments for participants to run. He has also been working on implementing the ObsSpace refactoring mentioned in previous meetings.

Mark M has also been preparing for the academy. With help from Ryan, we now have a web-based JupyterLab interface that facilitates how academy participants access their amazon nodes.

Mark M also asked if we were ready to merge the pull requests (in oops, saber, and fv3-jedi) from Benjamin that incorporate atlas interfaces into bump and saber. Mark mentioned that this is a significant change because it would add atlas as a dependency for all bundles, since saber would now depends on atlas and oops depends on saber. Dan reported that there are still build problems with these branches on Discover when using gnu compilers. If Discover is the only place where these problems occur then it might not be enough to hold up these pull requests. Mark noted that they do pass the gnu-openmpi tests on CodeBuild and Travis (also on his Mac laptop). But, more testing with gnu on HPC systems is needed before we merge these. Yannick urged Dan to post a comment about this in the pull requests so we don't merge before this issue is resolved.

Then there was some discussion about whether all bundles would now depend on atlas after these PRs are merged. Mark M said yes, since all bundles currently depend on saber through oops. Anna suggested removing the saber dependency from oops and Yannick agreed this would be a good idea. However, even if this is done, the new oops branch still would depend on atlas.

Ming has been planning with colleagues on how to integrate JEDI into two existing projects, namely 3drmta and RRFS (rapid refresh forecast system). In this context, he wants to compare the results of ufo and GSI and he expects to have results from this comparison within the next few months.

Sarah has been coordinating with colleagues on how to further integrate aerosols into JEDI. On her end, she will focus on obs preprocessing.

BJ has been working on the B matrix for MPAS. He was having problems with the linear variable change. It's defined in the model space covariance class but it is being called twice, which is causing problems with the nonlinear minimization of Jb. Yannick said that the minimization of Jb is often not included in operational workflows because it is expensive. As a result, this part of the JEDI code is less mature and may not be well tested. BJ asked if we could add an option not to calculate Jb and Yannick agreed this would be a good idea.

Yannick then pointed out that the single-obs functionality is now working and might be a useful way for BJ to check his B matrix. BJ asked if there was an example test for how to run with a single observation. Steve mentioned that there is one in ufo. Yannick said it would be useful to define a single-observation test that did not need an obs file for input. Rather, we could just generate a single observation from settings in the yaml file. Steve suggested the use of `MakeObs()` for this but Yannick said it could be even simpler than that - just specify the location and value in the yaml file.

We then turned to the meeting participants in google hangouts. Chris H has been working to make the B matrix in the shallow water model more realistic, with help from Dan. He has implemented the multiplication of B by a vector and is now working on the multiplication of the inverse of B by a vector. He has implemented a conjugate gradient method with a Jacobi pre-conditioner. It works but it is slow and he's not sure it will continue to work robustly when he extends the B matrix to include off-diagonal elements. Since the SW model will be used as an example for other models to follow, he wants to make this clear. He has considered developing a more sophisticated conjugate gradient method but he's not sure it would be worth the effort. Yannick agreed - he said that the B matrices in models tend to be poorly conditioned and it could be a substantial amount of work. He questioned whether this would be worthwhile since most operational systems don't even use this diagnostic. Yannick suggested using the GMRESR function in oops and pointed Chris to the qg model for an example.

Dan has been preparing for the Academy and for the atlas code sprint that follows the Academy. In particular, he has been refactoring fv3-jedi to isolate the cubed sphere implementation so we can potentially replace some of this with an atlas interface after the cubed sphere functionality has been added to atlas.

Ryan has also been preparing for the Academy. He said that the Hera system administrators have requested that we do our building and testing in an interactive debug q rather than directly on a login node. He also mentioned that a request has been put in for gcc 9 on Hera and this is still being considered.

In the next report from the UKMO, Marek mentioned that he has been working on the UM-JEDI model, including the grid, the analytic test, and the transform. He and Steve are also preparing lfirc for the upcoming atlas code sprint. Steve S is implementing atlas interfaces into lfirc, following the fv3-jedi implementation. Yannick encouraged them to let us know if they have any problems with the new atlas functionality in the active feature/atlas\_interfaces pull requests mentioned above. If so, we can hold off a little longer with them.

Jake asked about the status of RTTOV in JEDI. Marek said they did the initial implementation nearly a year ago and they got it working then with a single channel but they haven't done much with it since.

Guillaume has been working on refactoring the State and Increment classes in SOCA and on a python resource for the static B. He mentioned that oops is currently breaking the Travis CI test for SOCA. Anna asked him when that was because they just merged in a bugfix recently. Guillaume thought that it was still failing about an hour ago but wasn't sure.

Sergey has been working on devising a strategy for how best to implement advanced LETKF in oops.

Cory reported from EMC and introduced two new members of the JEDI team: Praveen Kumar and Nick Esposito. They'll both be attending the academy next week. They'll be working on the observation processing side of JEDI.

Yannick closed the meeting by again announcing there will be no meeting next week and encouraging all to send us ideas on topics that you would like to see covered in our bi-weekly focused discussions. He also mentioned that the C++ discussion last week was well received and we are likely to have more of them in the future.