

2018-09-27

Mark opened the meeting with several announcements

- Yannick is on vacation, he will be gone next week (Oct 4th) and return to this meeting two weeks from now (Oct 11th).
- Several merges into the develop branches occurred this past week
 - The real April 15th obs and geovals data
 - The GNSSRO results from the code sprint
 - A bugfix for the QG State test
- There has been lots of progress on BUMP
 - The memory leak appears to be fixed
 - The reading in of data from files is improved, but Guillaume is still having trouble with the reading in of length scales
- The above fixes for BUMP are in the "anisotropic correlation" branch of oops
 - It is recommended that everyone use this branch when working with BUMP, particularly if you are having problems with develop
 - The pull request for this branch is expected to be issued soon

JJ reported that he is working on resolving big-endian vs little-endian problems with compiling the MPAS bundle. Files being read by MPAS are constructed in big-endian format, and when compiling on little-endian systems options need to be added to the compilers to be able to read/write these big-endian files. The issue is that `ecbuild/cmake` are detecting a little-endian system and thwarting the attempts to add the "read/write files in big-endian" options for the MPAS code compile. Mark is assisting with resolving this issue. JJ and Mark's work is being tracked in the github `jedi-models` discussion group. Anyone with some thoughts about how to fix this please join in the `jedi-models` discussion group.

Dan is currently working on 4DVar H(x) with anisotropic correlations. He has created new branches in IODA and UFO called `4D_obs` for this work. The focus of this work is to get obs binning (into time slices) operational. Currently, the system places all obs into all time slices. He ran into an issue where the obs binning resulted in zero obs on a process element, and this caught BUMP off guard. Benjamin is currently fixing this. Once this is fixed, Dan should be able to complete the 4D H(x) work.

Dan is also cleaning up the 3DVar code for the radiance obs type. Modifications include restoring the ability to read Tb into CRTM (for diagnostic purposes), being able to specify channels to skip over in the configuration files (CRTM will skip the channels as well as UFO), and moving some hardwired meta data from the CRTM code into configuration files. Ben is helping Dan with this work. Steve H will help with reviewing the changes Dan and Ben have made.

Ben announced that he is close to issuing pull requests for the TL and AD radiance obs operators.

Mariusz asked if there will be major changes to update to the new April 15th radiance obs data. Dan responded that the primary change is that the individual channels are now stored in individual variables in the file.

Hui requested help with handling missing values in the GNSSRO obs operators. A QC step in the GNSSRO obs operators which needs to run after reading in obs/geovals data can create missing values. The issue is that the current system assumes that there are no missing values and does not have the capability to process missing values (such as skipping over missing values during vector dot products). The current set of obs/geovals data does not contain missing values with the exception of the new April 15th data for Radiosondes. This has been worked around by filtering out the missing values when reading in the Radiosonde data. We may be able to leverage this workaround for the GNSSRO issue. Steve H will work with Hui to get this resolved.

Marek reported that they are applying the Increment, State refactoring approach to LFRic. This is based on the work that Dan did for FV3. The Met Office will work on 4DVar H(x) next. They are also working out the issues that are limiting them to aqua-planet simulations.

Marek mentioned that they are now using doxygen build in connection with ufo-bundle and lfric. Mark pointed out that this build can be shut off if desired by editing the UFO top-level `CMakeLists.txt` file.

Rahul is working on the pseudo model discussed last week in which the nonlinear model evolution is replaced by reading a series of states from disk. He is following up on suggestions made by Yannick in last week's meeting. Rahul is also close to getting models/tools working on an Amazon cluster. He has libraries loaded and has FV3 and GSI working. He has not yet run this across nodes but the JEDI core team can help with that. He is building these tools for both ubuntu and CENTOS 7 operating systems because the latter is more conducive to creating HPC clusters on AWS.

Steve H announced the following merges:

- The clobber control for writing an obs output file has been removed. There no longer is a need to place a clobber control in your configuration files. A ZenHub issue has been submitted to remind us to work on this feature later.
- The ODB library build scripts have been moved into the `jedi-tools` repository
 - We are now only supporting ODB2
 - ECMWF's name for ODB2 is `ODB_API` so we should use `ODB_API` from now on to be consistent
 - The wiki instructions for building ODB libraries has been updated accordingly

Steve H mentioned that we will be installing the `ODB_API` libraries into the Singularity image soon so that users of the container will not have to worry about building `ODB_API`.

Steve H and Xin gave an update on the C++ work in IODA.

- We are rewriting `ObsSpace` and `ObsVector` to be totally implemented in C++
 - A Fortran interface will be provided for access by UFO
- We will get the Radiosonde obs operator working first, then follow that with porting all the other obs types over to the new `ObsSpace` and `ObsVector` classes.
- We should have the Radiosonde case working in the next 1 to 2 weeks
- The other obs types are targeted for completion by the end of October
- Xin is working on the C++ implementation for in-memory obs data storage
 - This uses a boost (header-only) package called `MultiIndex`
 - Xin is about a day away from completion of the `MultiIndex` work
- Steve H is working on the `netcdf` file interface
- Steve Vahl is working on the ODB file interface

- Mark is working on the Fortran interface for access by UFO

Mark is reworking the BUMP implementation to switch over to fckit MPI calls. Currently, BUMP is calling the MPI API directly. Mark is working in a fork of fckit and will ultimately issue a pull request to ECMWF. Guillaume pointed out that the fckit MPI barrier call is private, and asked why it is not public. Other fckit MPI calls similar to barrier are public and available. Mark will help get this resolved.

Marek asked about the availability of user documentation for BUMP. Dan responded that Benjamin is planning on writing these documents, perhaps over the next month or two. Mark added that the core team could assist somewhat, but Benjamin will have to write the bulk of the documentation.